

Configure and manage

StorageGRID 11.8

NetApp May 10, 2024

This PDF was generated from https://docs.netapp.com/us-en/storagegrid-118/admin/index.html on May 10, 2024. Always check docs.netapp.com for the latest.

Table of Contents

Configure and manage a StorageGRID system	1
Administer StorageGRID	1
Manage objects with ILM	. 319
System hardening	. 438
Configure StorageGRID for FabricPool	. 445

Configure and manage a StorageGRID system

Administer StorageGRID

Administer StorageGRID: Overview

Use these instructions to configure and administer a StorageGRID system.

About these instructions

The primary tasks for configuring and administering StorageGRID allow you to:

- Use the Grid Manager to set up groups and users
- · Create tenant accounts to allow S3 and Swift client applications to store and retrieve objects
- Configure and manage StorageGRID networks
- Configure AutoSupport
- Manage node settings

Before you begin

- You have a general understanding of the StorageGRID system.
- You have fairly detailed knowledge of Linux command shells, networking, and server hardware setup and configuration.

Get started with Grid Manager

Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	119
Microsoft Edge	119
Mozilla Firefox	119

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

Sign in to the Grid Manager

You access the Grid Manager sign-in page by entering the fully qualified domain name (FQDN) or IP address of an Admin Node into the address bar of a supported web browser.

Overview

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 aren't exactly the same:

- Alarm acknowledgments (legacy system) made on one Admin Node aren't 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.

Connect to HA group

If Admin Nodes are included in a high availability (HA) group, you connect using the virtual IP address of the HA group or a fully qualified domain name that maps to the virtual IP address. The primary Admin Node should be selected as the group's primary interface, so that when you access the Grid Manager, you access it on the primary Admin Node unless the primary Admin Node is not available. See Manage high availability groups.

Use SSO

The sign-in steps are slightly different if single sign-on (SSO) has been configured.

Sign in to Grid Manager on first Admin Node

Before you begin

- You have your login credentials.
- You are using a supported web browser.
- · Cookies are enabled in your web browser.
- You belong to a user group that has at least one permission.
- · You have the URL for the Grid Manager:

https://FQDN or Admin Node IP/

You can use the fully qualified domain name, the IP address of an Admin Node, or the virtual IP address of an HA group of Admin Nodes.

To access the Grid Manager on a port other than the default port for HTTPS (443), include the port number in the URL:

https://FQDN_or_Admin_Node_IP:port/



SSO is not available on the restricted Grid Manager port. You must use port 443.

Steps

1. Launch a supported web browser.

- 2. In the browser's address bar, enter the URL for the Grid Manager.
- 3. If you are prompted with a security alert, install the certificate using the browser's installation wizard. See Manage security certificates.
- 4. Sign in to the Grid Manager.

The sign-in screen that appears depends on whether single sign-on (SSO) has been configured for StorageGRID.

Not using SSO

- a. Enter your username and password for the Grid Manager.
- b. Select Sign In.

NetApp StorageGRID [®]
Grid Manager
Username
1
Password
Sign in
Tenant sign in NetApp support NetApp.com

Using SSO

- If StorageGRID is using SSO and this is the first time you have accessed the URL on this browser:
 - a. Select **Sign in**. You can leave the 0 in the Account field.

NetApp StorageGRID [®]
Sign in
Account
0
Sign in
NetApp support NetApp.com

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

Sign in with your organizational account		
someone@example.com		
Password		
Sign in		

- If StorageGRID is using SSO and you have previously accessed the Grid Manager or a tenant account:
 - a. Enter **0** (the account ID for the Grid Manager) or select **Grid Manager** if it appears in the list of recent accounts.

~	Sign in		
	3		
R	ecent		
	Grid Manager	V	
A	ccount		
	0		
	Sign in		
Ň	etApp support NetApp.com		

When you are signed in, the home page of the Grid Manager appears, which includes the dashboard. To learn what information is provided, see View and manage the dashboard.

StorageGRID dashboa	nrd					Actions ~
 You have 4 notifications: 1 ● 3 ▲ 						
Overview Performance Storage	ILM Nodes					
Health status @	Data space usage 2.11 MB (0%) of 3.	breakdown @ 09 TB used overall				c
License 1	Site name 🗢 Data Center 2	Data storage usage 0%	 Used space 682.53 KB 	¢	Total space 926.62 GB	÷
License	Data Center 3 Data Center 1	0% 0%	646.12 KB 779.21 KB		926.62 GB 1.24 TB	
Total objects in the grid 🛛	Metadata allowed 3.62 MB (0%) of 25	d space usage breakd	own 🥑			C.
0	Data Center 1 has th the grid.	e highest metadata space	e usage and it deter	mines the	e metadata space a	vailable in
	Site name 🗢	Metadata space usage	Used space	¢	Allowed space	• ^
	Data Center 3	0%	2.71 MB		19.32 GB	*

Sign into another Admin Node

Follow these steps to sign in to another Admin Node.

Not using SSO

Steps

- 1. In the browser's address bar, enter the fully qualified domain name or IP address of the other Admin Node. Include the port number as required.
- 2. Enter your username and password for the Grid Manager.
- 3. Select Sign In.

Using SSO

If StorageGRID is using SSO and you have signed in to one Admin Node, you can access other Admin Nodes without having to sign in again.

Steps

- 1. Enter the fully qualified domain name or IP address of the other Admin Node in the browser's address bar.
- 2. If your SSO session has expired, enter your credentials again.

Sign out of the Grid Manager

When you are done working with the Grid Manager, you must sign out to ensure that unauthorized users can't access the StorageGRID system. Closing your browser might not sign you out of the system, based on browser cookie settings.

Steps

1. Select your user name in the top-right corner.



2. Select Sign out.

Option	Description
SSO not in use	You are signed out of the Admin Node.
	The Grid Manager sign in page is displayed.
	Note: If you signed into more than one Admin Node, you must sign out of each node.
SSO enabled	You are signed out of all Admin Nodes you were accessing. The StorageGRID sign in page is displayed. Grid Manager is listed as the default in the Recent Accounts drop-down, and the Account ID field shows 0.
	Note: If SSO is enabled and you are also signed in to the Tenant Manager, you must also sign out of the tenant account to sign out of SSO.

Change your password

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

Before you begin

You are signed in to the Grid Manager using a supported web browser.

About this task

If you sign in to StorageGRID as a federated user or if single sign-on (SSO) is enabled, you can't change your password in Grid Manager. Instead, you must change your password in the external identity source, for

Steps

- 1. From the Grid Manager header, select *your name* > Change password.
- 2. Enter your current password.
- 3. Type a new password.

Your password must contain at least 8 and no more than 32 characters. Passwords are case-sensitive.

- 4. Re-enter the new password.
- 5. Select Save.

View StorageGRID license information

You can view the license information for your StorageGRID system, such as the maximum storage capacity of your grid, whenever necessary.

Before you begin

• You are signed in to the Grid Manager using a supported web browser.

About this task

If there is an issue with the software license for this StorageGRID system, the Health status card on the dashboard includes a License status icon and a **License** link. The number indicates the number of license-related issues.

Health status 🥥		
	License 1	
	License	

Steps

- 1. Access the License page by doing one of the following:
 - Select MAINTENANCE > System > License.
 - From the Health status card on the dashboard, select the License status icon or the License link.

This link appears only if there is an issue with the license.

- 2. View the read-only details for the current license:
 - StorageGRID system ID, which is the unique identification number for this StorageGRID installation

- License serial number
- License type, either Perpetual or Subscription
- · Licensed storage capacity of the grid
- Supported storage capacity
- License end date. N/A appears for a perpetual license.
- Support end date

This date is read from the current license file and might be out of date if you extended or renewed the support service contract after obtaining the license file. To update this value, see Update StorageGRID license information. You can also view the actual contract end date using Active IQ.

· Contents of the license text file

Update StorageGRID license information

You must update the license information for your StorageGRID system any time the terms of your license change. For example, you must update the license information if you purchase additional storage capacity for your grid.

Before you begin

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

Steps

- 1. Select MAINTENANCE > System > License.
- 2. In the Update license section, select **Browse**.
- 3. Locate and select the new license file (.txt).

The new license file is validated and displayed.

- 4. Enter the provisioning passphrase.
- 5. Select Save.

Use the API

Use the Grid Management API

You can perform system management tasks using the Grid Management REST API instead of the Grid Manager user interface. For example, you might want to use the API to automate operations or to create multiple entities, such as users, more quickly.

Top-level resources

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

• /grid: Access is restricted to Grid Manager users and is based on the configured group permissions.

- /org: Access is restricted to users who belong to a local or federated LDAP group for a tenant account. For details, see Use a tenant account.
- /private: Access is restricted to Grid Manager users and is based on the configured group permissions. The private APIs are subject to change without notice. StorageGRID private endpoints also ignore the API version of the request.

Issue API requests

The Grid Management API uses the Swagger open source API platform. Swagger provides an intuitive user interface that allows developers and non-developers to perform real-time operations in StorageGRID with the API.

The Swagger user interface provides complete details and documentation for each API operation.

Before you begin

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



Any API operations you perform using the API Docs webpage are live operations. Be careful not to create, update, or delete configuration data or other data by mistake.

Steps

1. From the Grid Manager header, select the help icon and select **API documentation**.



2. To perform an operation with the private API, select **Go to private API documentation** on the StorageGRID Management API page.

The private APIs are subject to change without notice. StorageGRID private endpoints also ignore the API version of the request.

3. Select the desired operation.

When you expand an API operation, you can see the available HTTP actions, such as GET, PUT, UPDATE, and DELETE.

4. Select an HTTP action to see the request details, including the endpoint URL, a list of any required or optional parameters, an example of the request body (when required), and the possible responses.

GET	/grid/groups Lists Grid Administrator Groups	
arameters		Try it out
ame	Description	
(pe tring query)	filter by group type Available values : local, federated	
mit nteger query)	maximum number of results Default value : 25	
narker tring query)	marker-style pagination offset (value is Group's URN) marker - marker-style pagination offset (value	
ncludeMarker oolean query)	if set, the marker element is also returned	
rder tring query)	pagination order (desc requires marker) Available values : asc, desc	
Responses		Response content type application/json v
ode Des	cription	
00 SUG Exa	mple Value Model	
{	"responseTime": "2021-03-29T14:22:19.673Z", "status": "success", "apiversion": "3.3", "serversion": "5.1ce	

- 5. Determine if the request requires additional parameters, such as a group or user ID. Then, obtain these values. You might need to issue a different API request first to get the information you need.
- 6. Determine if you need to modify the example request body. If so, you can select **Model** to learn the requirements for each field.
- 7. Select Try it out.
- 8. Provide any required parameters, or modify the request body as required.
- 9. Select Execute.

10. Review the response code to determine if the request was successful.

Grid Management API operations

The Grid Management API organizes the available operations into the following sections.



This list only includes operations available in the public API.

- **accounts**: Operations to manage storage tenant accounts, including creating new accounts and retrieving storage usage for a given account.
- **alarms**: Operations to list current alarms (legacy system), and return information about the health of the grid, including the current alerts and a summary of node connection states.
- alert-history: Operations on resolved alerts.
- alert-receivers: Operations on alert notification receivers (email).
- alert-rules: Operations on alert rules.
- alert-silences: Operations on alert silences.
- alerts: Operations on alerts.
- audit: Operations to list and update the audit configuration.
- auth: Operations to perform user session authentication.

The Grid Management API supports the Bearer Token Authentication Scheme. To sign in, you provide a username and password in the JSON body of the authentication request (that is, POST /api/v3/authorize). If the user is successfully authenticated, a security token is returned. This token must be provided in the header of subsequent API requests ("Authorization: Bearer *token*"). The token expires after 16 hours.



If single sign-on is enabled for the StorageGRID system, you must perform different steps to authenticate. See "Authenticating in to the API if single sign-on is enabled."

See "Protecting against Cross-Site Request Forgery" for information about improving authentication security.

- client-certificates: Operations to configure client certificates so that StorageGRID can be accessed securely using external monitoring tools.
- **config**: Operations related to the product release and versions of the Grid Management API. You can list the product release version and the major versions of the Grid Management API supported by that release, and you can disable deprecated versions of the API.
- deactivated-features: Operations to view features that might have been deactivated.
- dns-servers: Operations to list and change configured external DNS servers.
- drive-details: Operations on drives for specific storage appliance models.
- endpoint-domain-names: Operations to list and change S3 endpoint domain names.
- erasure-coding: Operations on erasure-coding profiles.
- expansion: Operations on expansion (procedure-level).
- expansion-nodes: Operations on expansion (node-level).
- expansion-sites: Operations on expansion (site-level).

- grid-networks: Operations to list and change the Grid Network List.
- grid-passwords: Operations for grid password management.
- **groups**: Operations to manage local Grid Administrator Groups and to retrieve federated Grid Administrator Groups from an external LDAP server.
- **identity-source**: Operations to configure an external identity source and to manually synchronize federated group and user information.
- ilm: Operations on information lifecycle management (ILM).
- in-progress-procedures: Retrieves the maintenance procedures that are currently in progress.
- license: Operations to retrieve and update the StorageGRID license.
- logs: Operations for collecting and downloading log files.v
- **metrics**: Operations on StorageGRID metrics including instant metric queries at a single point in time and range metric queries over a range of time. The Grid Management API uses the Prometheus systems monitoring tool as the backend data source. For information about constructing Prometheus queries, see the Prometheus web site.



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

- node-details: Operations on node details.
- node-health: Operations on node health status.
- node-storage-state: Operations on node storage status.
- ntp-servers: Operations to list or update external Network Time Protocol (NTP) servers.
- objects: Operations on objects and object metadata.
- recovery: Operations for the recovery procedure.
- recovery-package: Operations to download the Recovery Package.
- regions: Operations to view and create regions.
- s3-object-lock: Operations on global S3 Object Lock settings.
- server-certificate: Operations to view and update Grid Manager server certificates.
- snmp: Operations on the current SNMP configuration.
- storage-watermarks: Storage node watermarks.
- traffic-classes: Operations for traffic classification policies.
- untrusted-client-network: Operations on the untrusted Client Network configuration.
- users: Operations to view and manage Grid Manager users.

Grid Management API versioning

The Grid Management API uses versioning to support non-disruptive upgrades.

For example, this Request URL specifies version 4 of the API.

https://hostname_or_ip_address/api/v4/authorize

The major version of the API is bumped when changes are made that are *not compatible* with older versions. The minor version of the API is bumped when changes are made that *are compatible* with older versions.

Compatible changes include the addition of new endpoints or new properties.

The following example illustrates how the API version is bumped based on the type of changes made.

Type of change to API	Old version	New version
Compatible with older versions	2.1	2.2
Not compatible with older versions	2.1	3.0
	3.0	4.0

When you install StorageGRID software for the first time, only the most recent version of the API is enabled. However, when you upgrade to a new feature release of StorageGRID, you continue to have access to the older API version for at least one StorageGRID feature release.



You can configure the supported versions. See the **config** section of the Swagger API documentation for the Grid Management API for more information. You should deactivate support for the older version after updating all API clients to use the newer version.

Outdated requests are marked as deprecated in the following ways:

- The response header is "Deprecated: true"
- The JSON response body includes "deprecated": true
- A deprecated warning is added to nms.log. For example:

Received call to deprecated v2 API at POST "/api/v2/authorize"

Determine which API versions are supported in the current release

Use the GET /versions API request to return a list of the supported API major versions. This request is located in the **config** section of the Swagger API documentation.

```
GET https://{{IP-Address}}/api/versions
{
    "responseTime": "2023-06-27T22:13:50.750Z",
    "status": "success",
    "apiVersion": "4.0",
    "data": [
        2,
        3,
        4
    ]
}
```

Specify an API version for a request

You can specify the API version using a path parameter (/api/v4) or a header (Api-Version: 4). If you provide both values, the header value overrides the path value.

```
curl https://[IP-Address]/api/v4/grid/accounts
```

curl -H "Api-Version: 4" https://[IP-Address]/api/grid/accounts

Protect against Cross-Site Request Forgery (CSRF)

You can help protect against Cross-Site Request Forgery (CSRF) attacks against StorageGRID by using CSRF tokens to enhance authentication that uses cookies. The Grid Manager and Tenant Manager automatically enable this security feature; other API clients can choose whether to enable it when they sign in.

An attacker that can trigger a request to a different site (such as with an HTTP form POST) can cause certain requests to be made using the signed-in user's cookies.

StorageGRID helps protect against CSRF attacks by using CSRF tokens. When enabled, the contents of a specific cookie must match the contents of either a specific header or a specific POST body parameter.

To enable the feature, set the csrfToken parameter to true during authentication. The default is false.

```
curl -X POST --header "Content-Type: application/json" --header "Accept:
application/json" -d "{
   \"username\": \"MyUserName\",
   \"password\": \"MyPassword\",
   \"cookie\": true,
   \"corfToken\": true
}" "https://example.com/api/v3/authorize"
```

When true, a GridCsrfToken cookie is set with a random value for sign-ins to the Grid Manager, and the AccountCsrfToken cookie is set with a random value for sign-ins to the Tenant Manager.

If the cookie is present, all requests that can modify the state of the system (POST, PUT, PATCH, DELETE) must include one of the following:

- The X-Csrf-Token header, with the value of the header set to the value of the CSRF token cookie.
- For endpoints that accept a form-encoded body: A csrfToken form-encoded request body parameter.

See the online API documentation for additional examples and details.



Requests that have a CSRF token cookie set will also enforce the "Content-Type: application/json" header for any request that expects a JSON request body as an additional protection against CSRF attacks.

Use the API if single sign-on is enabled (Active Directory)

If you have configured and enabled single sign-on (SSO) and you use Active Directory as the SSO provider, you must issue a series of API requests to obtain an authentication token that is valid for the Grid Management API or the Tenant Management API.

Sign in to the API if single sign-on is enabled

These instructions apply if you are using Active Directory as the SSO identity provider.

Before you begin

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

About this task

To obtain an authentication token, you can use one of the following examples:

- The storagegrid-ssoauth.py Python script, which is located in the StorageGRID installation files directory (./rpms for Red Hat Enterprise Linux, ./debs for Ubuntu or Debian, and ./vsphere for VMware).
- An example workflow of curl requests.

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



The example curl workflow does not protect the password from being seen by other users.

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

Steps

- 1. Select one of the following methods to obtain an authentication token:
 - Use the storagegrid-ssoauth.py Python script. Go to step 2.
 - Use curl requests. Go to step 3.
- 2. If you want to use the storagegrid-ssoauth.py script, pass the script to the Python interpreter and run the script.

When prompted, enter values for the following arguments:

- The SSO method. Enter ADFS or adfs.
- The SSO username
- The domain where StorageGRID is installed
- The address for StorageGRID
- The tenant account ID, if you want to access the Tenant Management API.

The StorageGRID authorization token is provided in the output. You can now use the token for other requests, similar to how you would use the API if SSO was not being used.

- 3. If you want to use curl requests, use the following procedure.
 - a. Declare the variables needed to sign in.

```
export SAMLUSER='my-sso-username'
export SAMLPASSWORD='my-password'
export SAMLDOMAIN='my-domain'
export TENANTACCOUNTID='12345'
export STORAGEGRID_ADDRESS='storagegrid.example.com'
export AD_FS_ADDRESS='adfs.example.com'
```



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

b. To receive a signed authentication URL, issue a POST request to /api/v3/authorize-saml, and remove the additional JSON encoding from the response.

This example shows a POST request for a signed authentication URL for TENANTACCOUNTID. The results will be passed to python -m json.tool to remove the JSON encoding.

```
curl -X POST "https://$STORAGEGRID_ADDRESS/api/v3/authorize-saml" \
    -H "accept: application/json" -H "Content-Type: application/json" \
    --data "{\"accountId\": \"$TENANTACCOUNTID\"}" | python -m
json.tool
```

The response for this example includes a signed URL that is URL-encoded, but it does not include the additional JSON-encoding layer.

```
{
    "apiVersion": "3.0",
    "data":
    "https://adfs.example.com/adfs/ls/?SAMLRequest=fZHLbsIwEEV%2FJTuv7...
sSl%2BfQ33cvfwA%3D&RelayState=12345",
    "responseTime": "2018-11-06T16:30:23.355Z",
    "status": "success"
}
```

c. Save the SAMLRequest from the response for use in subsequent commands.

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

d. Get a full URL that includes the client request ID from AD FS.

One option is to request the login form using the URL from the previous response.

```
curl "https://$AD_FS_ADDRESS/adfs/ls/?SAMLRequest=
$SAMLREQUEST&RelayState=$TENANTACCOUNTID" | grep 'form method="post"
id="loginForm"'
```

The response includes the client request ID:

```
<form method="post" id="loginForm" autocomplete="off"
novalidate="novalidate" onKeyPress="if (event && event.keyCode == 13)
Login.submitLoginRequest();" action="/adfs/ls/?
SAMLRequest=fZHRToMwFIZfhb...UJikvo77sXPw%3D%3D&RelayState=12345&clie
nt-request-id=00000000-0000-0000-ee02-008000000de" >
```

e. Save the client request ID from the response.

export SAMLREQUESTID='00000000-0000-0000-ee02-008000000de'

f. Send your credentials to the form action from the previous response.

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

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



If multi-factor authentication (MFA) is enabled for your SSO system, the form post will also contain the second password or other credentials.

```
HTTP/1.1 302 Found
Content-Length: 0
Content-Type: text/html; charset=utf-8
Location:
https://adfs.example.com/adfs/ls/?SAMLRequest=fZHRToMwFIZfhb...UJikvo
77sXPw%3D%3D&RelayState=12345&client-request-id=00000000-0000-
ee02-008000000de
Set-Cookie: MSISAuth=AAEAADAvsHpXk6ApV...pmP0aEiNtJvWY=; path=/adfs;
HttpOnly; Secure
Date: Tue, 06 Nov 2018 16:55:05 GMT
```

g. Save the MSISAuth cookie from the response.

export MSISAuth='AAEAADAvsHpXk6ApV...pmP0aEiNtJvWY='

h. Send a GET request to the specified location with the cookies from the authentication POST.

```
curl "https://$AD_FS_ADDRESS/adfs/ls/?SAMLRequest=
$SAMLREQUEST&RelayState=$TENANTACCOUNTID&client-request-
id=$SAMLREQUESTID" \
--cookie "MSISAuth=$MSISAuth" --include
```

The response headers will contain AD FS session information for later logout usage, and the response body contains the SAMLResponse in a hidden form field.

```
HTTP/1.1 200 OK
Cache-Control: no-cache, no-store
Pragma: no-cache
Content-Length: 5665
Content-Type: text/html; charset=utf-8
Expires: -1
Server: Microsoft-HTTPAPI/2.0
P3P: ADFS doesn't have P3P policy, please contact your site's admin
for more details
Set-Cookie:
SamlSession=a3dpbnRlcnMtUHJpbWFyeS1BZG1pbi0xNzgmRmFsc2Umcng4NnJDZmFKV
XFxVWx3bkl1MnFuUSUzZCUzZCYmJiYmXzE3MjAyZTA5LThmMDqtNDRkZC04Yzq5LTQ3ND
UxYzA3ZjkzYw==; 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://storagegrid.example.com:443/api/saml-response">
  <input type="hidden" name="SAMLResponse"</pre>
value="PHNhbWxwOlJlc3BvbnN...1scDpSZXNwb25zZT4=" /><input</pre>
type="hidden" name="RelayState" value="12345" />
```

i. Save the SAMLResponse from the hidden field:

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

j. Using the saved SAMLResponse, make a StorageGRID/api/saml-response request to generate a StorageGRID authentication token.

For RelayState, use the tenant account ID or use 0 if you want to sign in to the Grid Management API.

```
curl -X POST "https://$STORAGEGRID_ADDRESS:443/api/saml-response" \
    -H "accept: application/json" \
    --data-urlencode "SAMLResponse=$SAMLResponse" \
    --data-urlencode "RelayState=$TENANTACCOUNTID" \
    | python -m json.tool
```

The response includes the authentication token.

```
{
    "apiVersion": "3.0",
    "data": "56eb07bf-21f6-40b7-af0b-5c6cacfb25e7",
    "responseTime": "2018-11-07T21:32:53.486Z",
    "status": "success"
}
```

k. Save the authentication token in the response as MYTOKEN.

export MYTOKEN="56eb07bf-21f6-40b7-af0b-5c6cacfb25e7"

You can now use MYTOKEN for other requests, similar to how you would use the API if SSO was not being used.

Sign out of the API if single sign-on is enabled

If single sign-on (SSO) has been enabled, you must issue a series of API requests to sign out of the Grid Management API or the Tenant Management API. These instructions apply if you are using Active Directory as the SSO identity provider

About this task

If required, you can sign out of the StorageGRID API by logging out from your organization's single logout page. Or, you can trigger single logout (SLO) from StorageGRID, which requires a valid StorageGRID bearer token.

Steps

1. To generate a signed logout request, pass `cookie "sso=true" to the SLO API:

```
curl -k -X DELETE "https://$STORAGEGRID_ADDRESS/api/v3/authorize" \
-H "accept: application/json" \
-H "Authorization: Bearer $MYTOKEN" \
--cookie "sso=true" \
| python -m json.tool
```

A logout URL is returned:

```
{
    "apiVersion": "3.0",
    "data":
    "https://adfs.example.com/adfs/ls/?SAMLRequest=fZDNboMwEIRfhZ...HcQ%3D%3
D",
    "responseTime": "2018-11-20T22:20:30.839Z",
    "status": "success"
}
```

2. Save the logout URL.

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

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

```
curl --include "$LOGOUT REQUEST"
```

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

```
HTTP/1.1 302 Found
Location: https://$STORAGEGRID_ADDRESS:443/api/saml-
logout?SAMLResponse=fVLLasMwEPwVo7ss%...%23rsa-sha256
Set-Cookie: MSISSignoutProtocol=U2FtbA==; expires=Tue, 20 Nov 2018
22:35:03 GMT; path=/adfs; HttpOnly; Secure
```

4. Delete the StorageGRID bearer token.

Deleting the StorageGRID bearer token works the same way as without SSO. If `cookie "sso=true" is not provided, the user is logged out of StorageGRID without affecting the SSO state.

```
curl -X DELETE "https://$STORAGEGRID_ADDRESS/api/v3/authorize" \
-H "accept: application/json" \
-H "Authorization: Bearer $MYTOKEN" \
--include
```

A 204 No Content response indicates the user is now signed out.

HTTP/1.1 204 No Content

Use the API if single sign-on is enabled (Azure)

If you have configured and enabled single sign-on (SSO) and you use Azure as the SSO provider, you can use two example scripts to obtain an authentication token that is valid for the Grid Management API or the Tenant Management API.

Sign in to the API if Azure single sign-on is enabled

These instructions apply if you are using Azure as the SSO identity provider

Before you begin

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

About this task

To obtain an authentication token, you can use the following example scripts:

- The storagegrid-ssoauth-azure.py Python script
- The storagegrid-ssoauth-azure.js Node.js script

Both scripts are located in the StorageGRID installation files directory (./rpms for Red Hat Enterprise Linux, ./debs for Ubuntu or Debian, and ./vsphere for VMware).

To write your own API integration with Azure, see the storagegrid-ssoauth-azure.py script. The Python script makes two requests to StorageGRID directly (first to get the SAMLRequest, and later to get the authorization token), and also calls the Node.js script to interact with Azure to perform the SSO operations.

SSO operations can be executed using a series of API requests, but doing so is not straightforward. The Puppeteer Node.js module is used to scrape the Azure SSO interface.

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

Steps

- 1. Install the required dependencies, as follows:
 - a. Install Node.js (see https://nodejs.org/en/download/).
 - b. Install the required Node.js modules (puppeteer and jsdom):

npm install -g <module>

2. Pass the Python script to the Python interpreter to run the script.

The Python script will then call the corresponding Node.js script to perform the Azure SSO interactions.

- 3. When prompted, enter values for the following arguments (or pass them in using parameters):
 - $\circ\,$ The SSO email address used to sign in to Azure
 - The address for StorageGRID
 - $\circ\,$ The tenant account ID, if you want to access the Tenant Management API
- 4. When prompted, enter the password and be prepared to provide an MFA authorization to Azure if

requested.

```
c:\Users\user\Documents\azure_sso>py storagegrid-azure-ssoauth.py --sso-email-address user@my-domain.com
--sg-address storagegrid.examp.e.com --tenant-account-id-0
Enter the user's SSO password:
Watch for and approve a 2FA authorization request
StorageGRID Auth Token: {'responseTime': '2021-10-04T21:30:48.807Z', 'status': 'success', 'apiVersion':
 '3.4', 'data': '4807d93e-a3df-48f2-9680-906cd255979e'}
```



The script assumes MFA is done using Microsoft Authenticator. You might need to modify the script to support other forms of MFA (such as entering a code received in a text message).

The StorageGRID authorization token is provided in the output. You can now use the token for other requests, similar to how you would use the API if SSO was not being used.

Use the API if single sign-on is enabled (PingFederate)

If you have configured and enabled single sign-on (SSO) and you use PingFederate as the SSO provider, you must issue a series of API requests to obtain an authentication token that is valid for the Grid Management API or the Tenant Management API.

Sign in to the API if single sign-on is enabled

These instructions apply if you are using PingFederate as the SSO identity provider

Before you begin

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

About this task

To obtain an authentication token, you can use one of the following examples:

- The storagegrid-ssoauth.py Python script, which is located in the StorageGRID installation files directory (./rpms for Red Hat Enterprise Linux, ./debs for Ubuntu or Debian, and ./vsphere for VMware).
- An example workflow of curl requests.

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



The example curl workflow does not protect the password from being seen by other users.

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

Steps

1. Select one of the following methods to obtain an authentication token:

• Use the storagegrid-ssoauth.py Python script. Go to step 2.

- Use curl requests. Go to step 3.
- 2. If you want to use the storagegrid-ssoauth.py script, pass the script to the Python interpreter and run the script.

When prompted, enter values for the following arguments:

- The SSO method. You can enter any variation of "pingfederate" (PINGFEDERATE, pingfederate, and so on).
- The SSO username
- The domain where StorageGRID is installed. This field is not used for PingFederate. You can leave it blank or enter any value.
- The address for StorageGRID
- The tenant account ID, if you want to access the Tenant Management API.

python3 storagegrid-ssoauth.py
sso method: pingfederate
saml user: my-sso-username
saml_domain:
sg_address: storagegrid.example.com
tenant_account_id: 12345
Enter the user's SAML password:

StorageGRID Auth Token: 56eb07bf-21f6-40b7-afob-5c6cacfb25e7

The StorageGRID authorization token is provided in the output. You can now use the token for other requests, similar to how you would use the API if SSO was not being used.

- 3. If you want to use curl requests, use the following procedure.
 - a. Declare the variables needed to sign in.

```
export SAMLUSER='my-sso-username'
export SAMLPASSWORD='my-password'
export TENANTACCOUNTID='12345'
export STORAGEGRID_ADDRESS='storagegrid.example.com'
```

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

b. To receive a signed authentication URL, issue a POST request to /api/v3/authorize-saml, and remove the additional JSON encoding from the response.

This example shows a POST request for a signed authentication URL for TENANTACCOUNTID. The results will be passed to python -m json.tool to remove the JSON encoding.

```
curl -X POST "https://$STORAGEGRID_ADDRESS/api/v3/authorize-saml" \
    -H "accept: application/json" -H "Content-Type: application/json" \
    --data "{\"accountId\": \"$TENANTACCOUNTID\"}" | python -m
json.tool
```

The response for this example includes a signed URL that is URL-encoded, but it does not include the additional JSON-encoding layer.

```
{
   "apiVersion": "3.0",
   "data": "https://my-pf-baseurl/idp/SSO.saml2?...",
   "responseTime": "2018-11-06T16:30:23.355Z",
   "status": "success"
}
```

c. Save the SAMLRequest from the response for use in subsequent commands.

export SAMLREQUEST="https://my-pf-baseurl/idp/SSO.saml2?..."

d. Export the response and cookie, and echo the response:

```
RESPONSE=$(curl -c - "$SAMLREQUEST")
```

```
echo "$RESPONSE" | grep 'input type="hidden" name="pf.adapterId"
id="pf.adapterId"'
```

e. Export the 'pf.adapterId' value, and echo the response:

```
export ADAPTER='myAdapter'
```

```
echo "$RESPONSE" | grep 'base'
```

f. Export the 'href' value (remove the trailing slash /), and echo the response:

```
export BASEURL='https://my-pf-baseurl'
```

echo "\$RESPONSE" | grep 'form method="POST"'

g. Export the 'action' value:

export SSOPING='/idp/.../resumeSAML20/idp/SSO.ping'

h. Send cookies along with credentials:

```
curl -b <(echo "$RESPONSE") -X POST "$BASEURL$SSOPING" \
--data "pf.username=$SAMLUSER&pf.pass=
$SAMLPASSWORD&pf.ok=clicked&pf.cancel=&pf.adapterId=$ADAPTER"
--include</pre>
```

i. Save the SAMLResponse from the hidden field:

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

j. Using the saved SAMLResponse, make a StorageGRID/api/saml-response request to generate a StorageGRID authentication token.

For RelayState, use the tenant account ID or use 0 if you want to sign in to the Grid Management API.

```
curl -X POST "https://$STORAGEGRID_ADDRESS:443/api/saml-response" \
    -H "accept: application/json" \
    --data-urlencode "SAMLResponse=$SAMLResponse" \
    --data-urlencode "RelayState=$TENANTACCOUNTID" \
    | python -m json.tool
```

The response includes the authentication token.

```
{
    "apiVersion": "3.0",
    "data": "56eb07bf-21f6-40b7-af0b-5c6cacfb25e7",
    "responseTime": "2018-11-07T21:32:53.486Z",
    "status": "success"
}
```

k. Save the authentication token in the response as MYTOKEN.

export MYTOKEN="56eb07bf-21f6-40b7-af0b-5c6cacfb25e7"

You can now use MYTOKEN for other requests, similar to how you would use the API if SSO was not being used.

Sign out of the API if single sign-on is enabled

If single sign-on (SSO) has been enabled, you must issue a series of API requests to sign out of the Grid Management API or the Tenant Management API.

These instructions apply if you are using PingFederate as the SSO identity provider

About this task

If required, you can sign out of the StorageGRID API by logging out from your organization's single logout page. Or, you can trigger single logout (SLO) from StorageGRID, which requires a valid StorageGRID bearer token.

Steps

1. To generate a signed logout request, pass `cookie "sso=true" to the SLO API:

```
curl -k -X DELETE "https://$STORAGEGRID_ADDRESS/api/v3/authorize" \
-H "accept: application/json" \
-H "Authorization: Bearer $MYTOKEN" \
--cookie "sso=true" \
| python -m json.tool
```

A logout URL is returned:

```
{
    "apiVersion": "3.0",
    "data": "https://my-ping-
url/idp/SLO.saml2?SAMLRequest=fZDNboMwEIRfhZ...HcQ%3D%3D",
    "responseTime": "2021-10-12T22:20:30.839Z",
    "status": "success"
}
```

2. Save the logout URL.

```
export LOGOUT_REQUEST='https://my-ping-
url/idp/SLO.saml2?SAMLRequest=fZDNboMwEIRfhZ...HcQ%3D%3D'
```

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

```
curl --include "$LOGOUT REQUEST"
```

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

```
HTTP/1.1 302 Found
Location: https://$STORAGEGRID_ADDRESS:443/api/saml-
logout?SAMLResponse=fVLLasMwEPwVo7ss%...%23rsa-sha256
Set-Cookie: PF=QoKs...SgCC; Path=/; Secure; HttpOnly; SameSite=None
```

4. Delete the StorageGRID bearer token.

Deleting the StorageGRID bearer token works the same way as without SSO. If `cookie "sso=true" is not provided, the user is logged out of StorageGRID without affecting the SSO state.

```
curl -X DELETE "https://$STORAGEGRID_ADDRESS/api/v3/authorize" \
-H "accept: application/json" \
-H "Authorization: Bearer $MYTOKEN" \
--include
```

A 204 No Content response indicates the user is now signed out.

HTTP/1.1 204 No Content

Deactivate features with the API

You can use the Grid Management API to completely deactivate certain features in the StorageGRID system. When a feature is deactivated, no one can be assigned permissions to perform the tasks related to that feature.

About this task

The Deactivated Features system allows you to prevent access to certain features in the StorageGRID system. Deactivating a feature is the only way to prevent the root user or users who belong to admin groups with **Root access** permission from being able to use that feature.

To understand how this functionality might be useful, consider the following scenario:

Company A is a service provider who leases the storage capacity of their StorageGRID system by creating tenant accounts. To protect the security of their leaseholders' objects, Company A wants to ensure that its own employees can never access any tenant account after the account has been deployed.

Company A can accomplish this goal by using the Deactivate Features system in the Grid Management API. By completely deactivating the **Change tenant root password** feature in the Grid Manager (both the UI and the API), Company A 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.

Steps

- 1. Access the Swagger documentation for the Grid Management API. See Use the Grid Management API.
- 2. Locate the Deactivate Features endpoint.

3. To deactivate a feature, such as Change tenant root password, send a body to the API like this:

```
{ "grid": {"changeTenantRootPassword": true} }
```

When the request is complete, the Change tenant root password feature is disabled. The **Change tenant root password** management permission no longer appears in the user interface, and any API request that attempts to change the root password for a tenant will fail with "403 Forbidden."

Reactivate deactivated features

By default, you can use the Grid Management API to reactivate a feature that has been deactivated. However, if you want to prevent deactivated features from ever being reactivated, you can deactivate the **activateFeatures** feature itself.



The **activateFeatures** feature can't be reactivated. If you decide to deactivate this feature, be aware that you will permanently lose the ability to reactivate any other deactivated features. You must contact technical support to restore any lost functionality.

Steps

- 1. Access the Swagger documentation for the Grid Management API.
- 2. Locate the Deactivate Features endpoint.
- 3. To reactivate all features, send a body to the API like this:

{ "grid": null }

When this request is complete, all features, including the Change tenant root password feature, are reactivated. The **Change tenant root password** management permission now appears in the user interface, and any API request that attempts to change the root password for a tenant will succeed, assuming the user has the **Root access** or **Change tenant root password** management permission.



The previous example causes *all* deactivated features to be reactivated. If other features have been deactivated that should remain deactivated, you must explicitly specify them in the PUT request. For example, to reactivate the Change tenant root password feature and continue to deactivate the Alarm acknowledgment feature, send this PUT request:

{ "grid": { "alarmAcknowledgment": true } }

Control access to StorageGRID

Control StorageGRID access: Overview

You control who can access StorageGRID and which tasks users can perform by creating or importing groups and users and assigning permissions to each group. Optionally, you can enable single sign-on (SSO), create client certificates, and change grid passwords.

Control access to the Grid Manager

You determine who can access the Grid Manager and the Grid Management API by importing groups and users from an identity federation service or by setting up local groups and local users.

Using identity federation makes setting up groups and users faster, and it allows users to sign in to StorageGRID using familiar credentials. You can configure identity federation if you use Active Directory, OpenLDAP, or Oracle Directory Server.



Contact technical support if you want to use another LDAP v3 service.

You determine which tasks each user can perform by assigning different permissions to each group. For example, you might want users in one group to be able to manage ILM rules and users in another group to perform maintenance tasks. A user must belong to at least one group to access the system.

Optionally, you can configure a group to be read-only. Users in a read-only group can only view settings and features. They can't make any changes or perform any operations in the Grid Manager or Grid Management API.

Enable single sign-on

The StorageGRID system supports single sign-on (SSO) using the Security Assertion Markup Language 2.0 (SAML 2.0) standard. After you configure and enable SSO, all users must be authenticated by an external identity provider before they can access the Grid Manager, the Tenant Manager, the Grid Management API, or the Tenant Management API. Local users can't sign in to StorageGRID.

Change provisioning passphrase

The provisioning passphrase is required for many installation and maintenance procedures, and for downloading the StorageGRID Recovery Package. The passphrase is also required to download backups of the grid topology information and encryption keys for the StorageGRID system. You can change the passphrase as required.

Change node console passwords

Each node in your grid has a unique node console password, which you need to log in to the node as "admin" using SSH, or to the root user on a VM/physical console connection. As needed, you can change the node console password for each node.

Change the provisioning passphrase

Use this procedure to change the StorageGRID provisioning passphrase. The passphrase is required for recovery, expansion, and maintenance procedures. The passphrase is also required to download Recovery Package backups that include the grid topology information, grid node console passwords, and encryption keys for the StorageGRID system.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have Maintenance or Root access permissions.
- You have the current provisioning passphrase.

About this task

The provisioning passphrase is required for many installation and maintenance procedures, and for downloading the Recovery Package. The provisioning passphrase is not listed in the Passwords.txt file. Make sure to document the provisioning passphrase and keep it in a safe and secure location.

Steps

- 1. Select CONFIGURATION > Access control> Grid passwords.
- 2. Under Change provisioning passphrase, select Make a change
- 3. Enter your current provisioning passphrase.
- 4. Enter the new passphrase. The passphrase must contain at least 8 and no more than 32 characters. Passphrases are case-sensitive.
- 5. Store the new provisioning passphrase in a secure location. It is required for installation, expansion, and maintenance procedures.
- 6. Re-enter the new passphrase, and select **Save**.

The system displays a green success banner when the provisioning passphrase change is complete.

Provisioning passphrase successfully changed. Go to the Recovery Package to download a new Recovery Package.

- 7. Select Recovery Package.
- 8. Enter the new provisioning passphrase to download the new Recovery Package.



After changing the provisioning passphrase, you must immediately download a new Recovery Package. The Recovery Package file allows you to restore the system if a failure occurs.

Change node console passwords

Each node in your grid has a unique node console password, which you need to log in to the node. Use these steps to change each unique node console password for each node in your grid.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.
- You have the current provisioning passphrase.

About this task

Use the node console password to log in to a node as "admin" using SSH, or to the root user on a VM/physical console connection. The change node console password process creates new passwords for each node in your grid and stores the passwords in an updated Passwords.txt file in the recovery package. The passwords are listed in the Password column in the Passwords.txt file.



There are separate SSH access passwords for the SSH keys used for communication between nodes. The SSH access passwords aren't changed by this procedure.

Access the wizard

Steps

- 1. Select CONFIGURATION > Access control > Grid passwords.
- 2. Under Change node console passwords, select Make a change.

Enter the provisioning passphrase

Steps

- 1. Enter the provisioning passphrase for your grid.
- 2. Select Continue.

Download the current recovery package

Before changing node console passwords, download the current recovery package. You can use the passwords in this file if the password change process fails for any node.

Steps

1. Select Download recovery package.

2. Copy the recovery package file (.zip) to two safe, secure, and separate locations.



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

3. Select Continue.

4. When the confirmation dialog appears, select **Yes** if you are ready to start changing the node console passwords.

You can't cancel this process after it starts.

Change node console passwords

When the node console password process starts, a new recovery package is generated that includes the new passwords. Then, the passwords are updated on each node.

Steps

1. Wait for the new recovery package to be generated, which might take a few minutes.

2. Select Download new recovery package.

- 3. When the download completes:
 - a. Open the .zip file.
 - b. Confirm that you can access the contents, including the Passwords.txt file, which contains the new node console passwords.
 - c. Copy the new recovery package file (.zip) to two safe, secure, and separate locations.



Don't overwrite the old recovery package.

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

- 4. Select the checkbox to indicate you have downloaded the new recovery package and verified the content.
- 5. Select **Change node console passwords** and wait for all nodes to be updated with the new passwords. This might take a few minutes.

If passwords are changed for all nodes, a green success banner appears. Go to the next step.
If there is an error during the update process, a banner message lists the number of nodes that failed to have their passwords changed. The system will automatically retry the process on any node that failed to have its password changed. If the process ends with some nodes still not having a changed password, the **Retry** button appears.

If the password update failed for one or more nodes:

- a. Review the error messages listed in the table.
- b. Resolve the issues.
- c. Select Retry.



Retrying only changes the node console passwords on the nodes that failed during previous password change attempts.

- 6. After node console passwords have been changed for all nodes, delete the first recovery package you downloaded.
- 7. Optionally, use the **Recovery package** link to download an additional copy of the new recovery package.

Use identity federation

Using identity federation makes setting up groups and users faster, and it allows users to sign in to StorageGRID using familiar credentials.

Configure identity federation for Grid Manager

You can configure identity federation in the Grid Manager if you want admin groups and users to be managed in another system such as Active Directory, Azure Active Directory (Azure AD), OpenLDAP, or Oracle Directory Server.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- · You have specific access permissions.
- You are using Active Directory, Azure AD, OpenLDAP, or Oracle Directory Server as the identity provider.



If you want to use an LDAP v3 service that is not listed, contact technical support.

- If you plan to use OpenLDAP, you must configure the OpenLDAP server. See Guidelines for configuring an OpenLDAP server.
- If you plan to enable single sign-on (SSO), you have reviewed the requirements and considerations for single sign-on.
- If you plan to use Transport Layer Security (TLS) for communications with the LDAP server, the identity provider is using TLS 1.2 or 1.3. See Supported ciphers for outgoing TLS connections.

About this task

You can configure an identity source for the Grid Manager if you want to import groups from another system such as Active Directory, Azure AD, OpenLDAP, or Oracle Directory Server. You can import the following types of groups:

• Admin groups. The users in admin groups can sign in to the Grid Manager and perform tasks, based on the management permissions assigned to the group.

• Tenant user groups for tenants that don't use their own identity source. Users in tenant groups can sign in to the Tenant Manager and perform tasks, based on the permissions assigned to the group in the Tenant Manager. See Create tenant account and Use a tenant account for details.

Enter the configuration

Steps

- 1. Select **CONFIGURATION > Access control > Identity federation**.
- 2. Select Enable identity federation.
- 3. In the LDAP service type section, select the type of LDAP service you want to configure.

LDAP service type			
Select the type of LDAP service	you want to configure.		
			No. market
Active Directory	Azure	OpenLDAP	Other

Select Other to configure values for an LDAP server that uses Oracle Directory Server.

- 4. If you selected **Other**, complete the fields in the LDAP Attributes section. Otherwise, go to the next step.
 - **User Unique Name**: The name of the attribute that contains the unique identifier of an LDAP user. This attribute is equivalent to sAMAccountName for Active Directory and uid for OpenLDAP. If you are configuring Oracle Directory Server, enter uid.
 - **User UUID**: The name of the attribute that contains the permanent unique identifier of an LDAP user. This attribute is equivalent to <code>objectGUID</code> for Active Directory and <code>entryUUID</code> for OpenLDAP. If you are configuring Oracle Directory Server, enter <code>nsuniqueid</code>. Each user's value for the specified attribute must be a 32-digit hexadecimal number in either 16-byte or string format, where hyphens are ignored.
 - **Group Unique Name**: The name of the attribute that contains the unique identifier of an LDAP group. This attribute is equivalent to sAMAccountName for Active Directory and cn for OpenLDAP. If you are configuring Oracle Directory Server, enter cn.
 - **Group UUID**: The name of the attribute that contains the permanent unique identifier of an LDAP group. This attribute is equivalent to <code>objectGUID</code> for Active Directory and <code>entryUUID</code> for OpenLDAP. If you are configuring Oracle Directory Server, enter <code>nsuniqueid</code>. Each group's value for the specified attribute must be a 32-digit hexadecimal number in either 16-byte or string format, where hyphens are ignored.
- 5. For all LDAP service types, enter the required LDAP server and network connection information in the Configure LDAP server section.
 - Hostname: The fully qualified domain name (FQDN) or IP address of the LDAP server.
 - $\circ~\textbf{Port}:$ The port used to connect to the LDAP server.



The default port for STARTTLS is 389, and the default port for LDAPS is 636. However, you can use any port as long as your firewall is configured correctly.

• Username: The full path of the distinguished name (DN) for the user that will connect to the LDAP

server.

For Active Directory, you can also specify the Down-Level Logon Name or the User Principal Name.

The specified user must have permission to list groups and users and to access the following attributes:

- sAMAccountName or uid
- objectGUID, entryUUID, or nsuniqueid
- cn
- memberOf or isMemberOf
- Active Directory: objectSid, primaryGroupID, userAccountControl, and userPrincipalName
- Azure: accountEnabled and userPrincipalName
- Password: The password associated with the username.



If you change the password in the future, you must update it on this page.

 Group Base DN: The full path of the distinguished name (DN) for an LDAP subtree you want to search for groups. In the Active Directory example (below), all groups whose Distinguished Name is relative to the base DN (DC=storagegrid,DC=example,DC=com) can be used as federated groups.



The **Group unique name** values must be unique within the **Group Base DN** they belong to.

• **User Base DN**: The full path of the distinguished name (DN) of an LDAP subtree you want to search for users.



The User unique name values must be unique within the User Base DN they belong to.

• **Bind username format** (optional): The default username pattern StorageGRID should use if the pattern can't be determined automatically.

Providing **Bind username format** is recommended because it can allow users to sign in if StorageGRID is unable to bind with the service account.

Enter one of these patterns:

- UserPrincipalName pattern (Active Directory and Azure): [USERNAME]@example.com
- Down-level logon name pattern (Active Directory and Azure): example \ [USERNAME]
- Distinguished name pattern: CN=[USERNAME], CN=Users, DC=example, DC=com

Include [USERNAME] exactly as written.

6. In the Transport Layer Security (TLS) section, select a security setting.

• **Use STARTTLS**: Use STARTTLS to secure communications with the LDAP server. This is the recommended option for Active Directory, OpenLDAP, or Other, but this option is not supported for Azure.

- **Use LDAPS**: The LDAPS (LDAP over SSL) option uses TLS to establish a connection to the LDAP server. You must select this option for Azure.
- **Do not use TLS**: The network traffic between the StorageGRID system and the LDAP server will not be secured. This option is not supported for Azure.



Using the **Do not use TLS** option is not supported if your Active Directory server enforces LDAP signing. You must use STARTTLS or LDAPS.

- 7. If you selected STARTTLS or LDAPS, choose the certificate used to secure the connection.
 - **Use operating system CA certificate**: Use the default Grid CA certificate installed on the operating system to secure connections.
 - Use custom CA certificate: Use a custom security certificate.

If you select this setting, copy and paste the custom security certificate into the CA certificate text box.

Test the connection and save the configuration

After entering all values, you must test the connection before you can save the configuration. StorageGRID verifies the connection settings for the LDAP server and the bind username format, if you provided one.

Steps

- 1. Select Test connection.
- 2. If you did not provide a bind username format:
 - A "Test connection successful" message appears if the connection settings are valid. Select **Save** to save the configuration.
 - A "test connection could not be established" message appears if the connection settings are invalid. Select **Close**. Then, resolve any issues and test the connection again.
- 3. If you provided a bind username format, enter the username and password of a valid federated user.

For example, enter your own username and password. Don't include any special characters in the username, such as @ or /.

Test Connection	×
To test the connection and the bind username format, enter the username and password of a federated user. For example, enter your ow federated username and password. The test values are not saved.	n
Test username	
myusername	
The username of a federated user.	
Test password	
	0
Cancel Test Connectio	n

- A "Test connection successful" message appears if the connection settings are valid. Select **Save** to save the configuration.
- An error message appears if the connection settings, bind username format, or test username and password are invalid. Resolve any issues and test the connection again.

Force synchronization with the identity source

The StorageGRID system periodically synchronizes federated groups and users from the identity source. You can force synchronization to start if you want to enable or restrict user permissions as quickly as possible.

Steps

- 1. Go to the Identity federation page.
- 2. Select **Sync server** at the top of the page.

The synchronization process might take some time depending on your environment.



The **Identity federation synchronization failure** alert is triggered if there is an issue synchronizing federated groups and users from the identity source.

Disable identity federation

You can temporarily or permanently disable identity federation for groups and users. When identity federation is disabled, there is no communication between StorageGRID and the identity source. However, any settings you have configured are retained, allowing you to easily reenable identity federation in the future.

About this task

Before you disable identity federation, you should be aware of the following:

- Federated users will be unable to sign in.
- Federated users who are currently signed in will retain access to the StorageGRID system until their session expires, but they will be unable to sign in after their session expires.
- Synchronization between the StorageGRID system and the identity source will not occur, and alerts or alarms will not be raised for accounts that have not been synchronized.
- The Enable identity federation checkbox is disabled if single sign-on (SSO) is set to Enabled or Sandbox Mode. The SSO Status on the Single Sign-on page must be Disabled before you can disable identity federation. See Disable single sign-on.

Steps

- 1. Go to the Identity federation page.
- 2. Uncheck the Enable identity federation checkbox.

Guidelines for configuring an OpenLDAP server

If you want to use an OpenLDAP server for identity federation, you must configure specific settings on the OpenLDAP server.



For identity sources that aren't ActiveDirectory or Azure, StorageGRID will not automatically block S3 access to users who are disabled externally. To block S3 access, delete any S3 keys for the user or remove the user from all groups.

Memberof and refint overlays

The memberof and refint overlays should be enabled. For more information, see the instructions for reverse group membership maintenance in the

OpenLDAP documentation: Version 2.4 Administrator's Guide.

Indexing

You must configure the following OpenLDAP attributes with the specified index keywords:

- olcDbIndex: objectClass eq
- olcDbIndex: uid eq,pres,sub
- olcDbIndex: cn eq,pres,sub
- olcDbIndex: entryUUID eq

In addition, ensure the fields mentioned in the help for Username are indexed for optimal performance.

See the information about reverse group membership maintenance in the OpenLDAP documentation: Version 2.4 Administrator's Guide.

Manage admin groups

You can create admin groups to manage the security permissions for one or more admin users. Users must belong to a group to be granted access to the StorageGRID system.

Before you begin

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

Create an admin group

Admin groups allow you to determine which users can access which features and operations in the Grid Manager and the Grid Management API.

Access the wizard

Steps

- 1. Select CONFIGURATION > Access control > Admin groups.
- 2. Select Create group.

Choose a group type

You can create a local group or import a federated group.

- Create a local group if you want to assign permissions to local users.
- Create a federated group to import users from the identity source.

Local group

Steps

- 1. Select Local group.
- 2. Enter a display name for the group, which you can update later as required. For example, "Maintenance Users" or "ILM Administrators."
- 3. Enter a unique name for the group, which you can't update later.
- 4. Select Continue.

Federated group

Steps

- 1. Select Federated group.
- 2. Enter the name of the group you want to import, exactly as it appears in the configured identity source.
 - For Active Directory and Azure, use the sAMAccountName.
 - For OpenLDAP, use the CN (Common Name).
 - For another LDAP, use the appropriate unique name for the LDAP server.
- 3. Select Continue.

Manage group permissions

Steps

- 1. For **Access mode**, select whether users in the group can change settings and perform operations in the Grid Manager and the Grid Management API or whether they can only view settings and features.
 - **Read-write** (default): Users can change settings and perform the operations allowed by their management permissions.
 - Read-only: Users can only view settings and features. They can't make any changes or perform any
 operations in the Grid Manager or Grid Management API. Local read-only users can change their own
 passwords.



If a user belongs to multiple groups and any group is set to **Read-only**, the user will have read-only access to all selected settings and features.

2. Select one or more admin group permissions.

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

3. If you are creating a local group, select **Continue**. If you are creating a federated group, select **Create** group and Finish.

Add users (local groups only)

Steps

1. Optionally, select one or more local users for this group.

If you have not yet created local users, you can save the group without adding users. You can add this

group to the user on the Users page. See Manage users for details.

2. Select Create group and Finish.

View and edit admin groups

You can view details for existing groups, modify a group, or duplicate a group.

- To view basic information for all groups, review the table on the Groups page.
- To view all details for a specific group or to edit a group, use the **Actions** menu or the details page.

Task	Actions menu	Details page
View group details	 a. Select the checkbox for the group. b. Select Actions > View group details. 	Select the group name in the table.
Edit display name (local groups only)	 a. Select the checkbox for the group. b. Select Actions > Edit group name. c. Enter the new name. d. Select Save changes. 	 a. Select the group name to display the details. b. Select the edit icon c. Enter the new name. d. Select Save changes.
Edit access mode or permissions	 a. Select the checkbox for the group. b. Select Actions > View group details. c. Optionally, change the group's Access mode. d. Optionally, select or clear admin group permissions. e. Select Save changes. 	 a. Select the group name to display the details. b. Optionally, change the group's Access mode. c. Optionally, select or clear admin group permissions. d. Select Save changes.

Duplicate a group

Steps

- 1. Select the checkbox for the group.
- 2. Select Actions > Duplicate group.
- 3. Complete the Duplicate group wizard.

Delete a group

You can delete an admin group when you want to remove the group from the system, and remove all permissions associated with the group. Deleting an admin group removes any users from the group, but does not delete the users.

Steps

- 1. From the Groups page, select the checkbox for each group you want to remove.
- 2. Select Actions > Delete group.
- 3. Select Delete groups.

Admin group permissions

When creating admin user groups, you select one or more permissions to control access to specific features of the Grid Manager. You can then assign each user to one or more of these admin groups to determine which tasks that user can perform.

You must assign at least one permission to each group; otherwise, users belonging to that group will not be able to sign in to the Grid Manager or the Grid Management API.

By default, any user who belongs to a group that has at least one permission can perform the following tasks:

- Sign in to the Grid Manager
- · View the dashboard
- · View the Nodes pages
- Monitor grid topology
- · View current and resolved alerts
- View current and historical alarms (legacy system)
- · Change their own password (local users only)
- · View certain information provided on the Configuration and Maintenance pages

Interaction between permissions and Access mode

For all permissions, the group's **Access mode** setting determines whether users can change settings and perform operations or whether they can only view the related settings and features. If a user belongs to multiple groups and any group is set to **Read-only**, the user will have read-only access to all selected settings and features.

The following sections describe the permissions you can assign when creating or editing an admin group. Any functionality not explicitly mentioned requires the **Root access** permission.

Root access

This permission provides access to all grid administration features.

Acknowledge alarms (legacy)

This permission provides access to acknowledge and respond to alarms (legacy system). All signed-in users can view current and historical alarms.

If you want a user to monitor grid topology and acknowledge alarms only, you should assign this permission.

Change tenant root password

This permission provides access to the **Change root password** option on the Tenants page, allowing you to control who can change the password for the tenant's local root user. This permission is also used for migrating S3 keys when the S3 key import feature is enabled. Users who don't have this permission can't see the

Change root password option.



To grant access to the Tenants page, which contains the **Change root password** option, also assign the **Tenant accounts** permission.

Grid topology page configuration

This permission provides access to the Configuration tabs on the SUPPORT > Tools > Grid topology page.

ILM

This permission provides access to the following ILM menu options:

- Rules
- Policies
- · Erasure coding
- Regions
- Storage pools



Users must have the **Other grid configuration** and **Grid topology page configuration** permissions to manage storage grades.

Maintenance

Users must have the Maintenance permission to use these options:

- CONFIGURATION > Access control:
 - · Grid passwords
- CONFIGURATION > Network:
 - S3 endpoint domain names
- MAINTENANCE > Tasks:
 - $\circ\,$ Decommission
 - Expansion
 - Object existence check
 - Recovery
- MAINTENANCE > System:
 - Recovery package
 - · Software update
- SUPPORT > Tools:
 - Logs

Users who don't have the Maintenance permission can view, but not edit, these pages:

• MAINTENANCE > Network:

DNS servers

- Grid Network
- NTP servers
- MAINTENANCE > System:
 - License
- CONFIGURATION > Network:
 - S3 endpoint domain names
- CONFIGURATION > Security:
 - · Certificates
- CONFIGURATION > Monitoring:
 - Audit and syslog server

Manage alerts

This permission provides access to options for managing alerts. Users must have this permission to manage silences, alert notifications, and alert rules.

Metrics query

This permission provides access to:

- SUPPORT > Tools > Metrics page
- Custom Prometheus metrics queries using the Metrics section of the Grid Management API
- · Grid Manager dashboard cards that contain metrics

Object metadata lookup

This permission provides access to the **ILM > Object metadata lookup** page.

Other grid configuration

This permission provides access to additional grid configuration options.



To see these additional options, users must also have the **Grid topology page configuration** permission.

- ILM:
 - Storage grades
- CONFIGURATION > System:
 - Storage options
- SUPPORT > Alarms (legacy):
 - · Custom events
 - Global alarms
 - · Legacy email setup
- SUPPORT > Other:
 - Link cost

Storage appliance administrator

This permission provides:

- Access to the E-Series SANtricity System Manager on storage appliances through the Grid Manager.
- The ability to perform troubleshooting and maintenance tasks on the Manage drives tab for appliances that support these operations.

Tenant accounts

This permission provides the ability to:

- · Access the Tenants page, where you can create, edit, and remove tenant accounts
- · View existing traffic classification policies
- · View Grid Manager dashboard cards that contain tenant details

Manage users

You can view local and federated users. You can also create local users and assign them to local admin groups to determine which Grid Manager features these users can access.

Before you begin

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

Create a local user

You can create one or more local users and assign each user to one or more local groups. The group's permissions control which Grid Manager and Grid Management API features the user can access.

You can create local users only. Use the external identity source to manage federated users and groups.

The Grid Manager includes one predefined local user, named "root." You can't remove the root user.



If single sign-on (SSO) is enabled, local users can't sign in to StorageGRID.

Access the wizard

Steps

- 1. Select CONFIGURATION > Access control > Admin users.
- 2. Select Create user.

Enter user credentials

Steps

- 1. Enter the user's full name, a unique username, and a password.
- 2. Optionally, select Yes if this user should not have access to the Grid Manager or Grid Management API.
- 3. Select Continue.

Assign to groups

Steps

1. Optionally, assign the user to one or more groups to determine the user's permissions.

If you have not yet created groups, you can save the user without selecting groups. You can add this user to a group on the Groups page.

If a user belongs to multiple groups, the permissions are cumulative. See Manage admin groups for details.

2. Select Create user and select Finish.

View and edit local users

÷.

You can view details for existing local and federated users. You can modify a local user to change the user's full name, password, or group membership. You can also temporarily prevent a user from accessing the Grid Manager and the Grid Management API.

You can edit local users only. Use the external identity source to manage federated users.

- To view basic information for all local and federated users, review the table on the Users page.
- To view all details for a specific user, edit a local user, or change a local user's password, use the **Actions** menu or the details page.

Any edits are applied the next time the user signs out and then signs back in to the Grid Manager.

Local users can change their own passwords using the **Change password** option in the Grid Manager banner.

Task	Actions menu	Details page
View user details	a. Select the checkbox for the user.b. Select Actions > View user details.	Select the user's name in the table.
Edit full name (local users only)	 a. Select the checkbox for the user. b. Select Actions > Edit full name. c. Enter the new name. d. Select Save changes. 	 a. Select the user's name to display the details. b. Select the edit icon . c. Enter the new name. d. Select Save changes.

Task	Actions menu	Details page
Deny or allow StorageGRID access	 a. Select the checkbox for the user. b. Select Actions > View user details. c. Select the Access tab. d. Select Yes to prevent the user from signing in to the Grid Manager or the Grid Management API, or select No to allow the user to sign in. e. Select Save changes. 	 a. Select the user's name to display the details. b. Select the Access tab. c. Select Yes to prevent the user from signing in to the Grid Manager or the Grid Management API, or select No to allow the user to sign in. d. Select Save changes.
Change password (local users only)	 a. Select the checkbox for the user. b. Select Actions > View user details. c. Select the Password tab. d. Enter a new password. e. Select Change password. 	a. Select the user's name to display the details.b. Select the Password tab.c. Enter a new password.d. Select Change password.
Change groups (local users only)	 a. Select the checkbox for the user. b. Select Actions > View user details. c. Select the Groups tab. d. Optionally, select the link after a group name to view the group's details in a new browser tab. e. Select Edit groups to select different groups. f. Select Save changes. 	 a. Select the user's name to display the details. b. Select the Groups tab. c. Optionally, select the link after a group name to view the group's details in a new browser tab. d. Select Edit groups to select different groups. e. Select Save changes.

Duplicate a user

You can duplicate an existing user to create a new user with the same permissions.

Steps

- 1. Select the checkbox for the user.
- 2. Select Actions > Duplicate user.
- 3. Complete the Duplicate user wizard.

Delete a user

You can delete a local user to permanently remove that user from the system.



You can't delete the root user.

Steps

1. From the Users page, select the checkbox for each user you want to remove.

- 2. Select Actions > Delete user.
- 3. Select Delete user.

Use single sign-on (SSO)

Configure single sign-on

When single sign-on (SSO) is enabled, users can only access the Grid Manager, the Tenant Manager, the Grid Management API, or the Tenant Management API if their credentials are authorized using the SSO sign-in process implemented by your organization. Local users can't sign in to StorageGRID.

How single sign-on works

The StorageGRID system supports single sign-on (SSO) using the Security Assertion Markup Language 2.0 (SAML 2.0) standard.

Before enabling single sign-on (SSO), review how the StorageGRID sign-in and sign-out processes are affected when SSO is enabled.

Sign in when SSO is enabled

When SSO is enabled and you sign in to StorageGRID, you are redirected to your organization's SSO page to validate your credentials.

Steps

1. Enter the fully qualified domain name or IP address of any StorageGRID Admin Node in a web browser.

The StorageGRID Sign in page appears.

• If this is the first time you have accessed the URL on this browser, you are prompted for an account ID:

NetApp StorageGRID [®]
Sign in
Account
0
Sign in
NetApp support NetApp.com

 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:

Tenant Manage	ər
Recent	
S3 tenant	۲
Account	
62984032838045582045	



The StorageGRID Sign in page is not shown when you enter the complete URL for a tenant account (that is, a fully qualified domain name or IP address followed by /?accountId=20-digit-account-id). Instead, you are immediately redirected to your organization's SSO sign-in page, where you can sign in with your SSO credentials.

- 2. Indicate whether you want to access the Grid Manager or the Tenant Manager:
 - To access the Grid Manager, leave the **Account ID** field blank, enter **0** as the account ID, or select **Grid Manager** if it appears in the list of recent accounts.
 - To access the Tenant Manager, enter the 20-digit tenant account ID or select a tenant by name if it appears in the list of recent accounts.
- 3. Select Sign in

StorageGRID redirects you to your organization's SSO sign-in page. For example:

nue rel	omeone@example.o	com	
sword	assword		

4. Sign in with your SSO credentials.

If your SSO credentials are correct:

- a. The identity provider (IdP) provides an authentication response to StorageGRID.
- b. StorageGRID validates the authentication response.
- c. If the response is valid and you belong to a federated group with StorageGRID access permissions, you are signed in to the Grid Manager or the Tenant Manager, depending on which account you selected.



If the service account is inaccessible, you can still sign in, as long as you are an existing user that belongs to a federated group with StorageGRID access permissions.

5. Optionally, access other Admin Nodes, or access the Grid Manager or the Tenant Manager, if you have adequate permissions.

You don't need to reenter your SSO credentials.

Sign out when SSO is enabled

When SSO is enabled for StorageGRID, what happens when you sign out depends on what you are signed in to and where you are signing out from.

Steps

- 1. Locate the **Sign out** link in the top-right corner of the user interface.
- 2. Select Sign out.

The StorageGRID Sign in page appears. The **Recent Accounts** drop-down is updated to include **Grid Manager** or the name of the tenant, so you can access these user interfaces more quickly in the future.

If you are signed in to	And you sign out from	You are signed out of
Grid Manager on one or more Admin Nodes	Grid Manager on any Admin Node	Grid Manager on all Admin Nodes Note: If you use Azure for SSO, it might take a few minutes to be signed out of all Admin Nodes.
Tenant Manager on one or more Admin Nodes	Tenant Manager on any Admin Node	Tenant Manager on all Admin Nodes
Both Grid Manager and Tenant Manager	Grid Manager	The Grid Manager only. You must also sign out of the Tenant Manager to sign out of SSO.
	Tenant Manager	The Tenant Manager only. You must also sign out of the Grid Manager to sign out of SSO.



The table summarizes what happens when you sign out if you are using a single browser session. If you are signed in to StorageGRID across multiple browser sessions, you must sign out of all browser sessions separately.

Requirements and considerations for single sign-on

Before enabling single sign-on (SSO) for a StorageGRID system, review the requirements and considerations.

Identity provider requirements

StorageGRID supports the following SSO identity providers (IdP):

- Active Directory Federation Service (AD FS)
- Azure Active Directory (Azure AD)
- PingFederate

You must configure identity federation for your StorageGRID system before you can configure an SSO identity provider. The type of LDAP service you use for identity federation controls which type of SSO you can implement.

Configured LDAP service type	Options for SSO identity provider
Active Directory	Active Directory
	• Azure
	PingFederate
Azure	Azure

AD FS requirements

You can use any of the following versions of AD FS:

- Windows Server 2022 AD FS
- Windows Server 2019 AD FS
- Windows Server 2016 AD FS



Windows Server 2016 should be using the KB3201845 update, or higher.

Additional requirements

- Transport Layer Security (TLS) 1.2 or 1.3
- Microsoft .NET Framework, version 3.5.1 or higher

Considerations for Azure

If you use Azure as the SSO type and users have user principal names that don't use the sAMAccountName as the prefix, login issues can occur if StorageGRID loses its connection with the LDAP server. To allow users to sign in, you must restore the connection to the LDAP server.

Server certificate requirements

By default, StorageGRID uses a management interface certificate on each Admin Node to secure access to the Grid Manager, the Tenant Manager, the Grid Management API, and the Tenant Management API. When you configure relying party trusts (AD FS), enterprise applications (Azure), or service provider connections (PingFederate) for StorageGRID, you use the server certificate as the signature certificate for StorageGRID requests.

If you have not already configured a custom certificate for the management interface, you should do so now. When you install a custom server certificate, it is used for all Admin Nodes, and you can use it in all StorageGRID relying party trusts, enterprise applications, or SP connections.



Using an Admin Node's default server certificate in a relying party trust, enterprise application, or SP connection is not recommended. If the node fails and you recover it, a new default server certificate is generated. Before you can sign in to the recovered node, you must update the relying party trust, enterprise application, or SP connection with the new certificate.

You can access an Admin Node's server certificate by logging in to the command shell of the node and going to the /var/local/mgmt-api directory. A custom server certificate is named custom-server.crt. The node's default server certificate is named server.crt.

Port requirements

Single sign-on (SSO) is not available on the restricted Grid Manager or Tenant Manager ports. You must use the default HTTPS port (443) if you want users to authenticate with single sign-on. See Control access at external firewall.

Confirm federated users can sign in

Before you enable single sign-on (SSO), you must confirm that at least one federated user can sign in to the Grid Manager and in to the Tenant Manager for any existing tenant accounts.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.
- · You have already configured identity federation.

Steps

1. If there are existing tenant accounts, confirm that none of the tenants is using its own identity source.



When you enable SSO, an identity source configured in the Tenant Manager is overridden by the identity source configured in the Grid Manager. Users belonging to the tenant's identity source will no longer be able to sign in unless they have an account with the Grid Manager identity source.

- a. Sign in to the Tenant Manager for each tenant account.
- b. Select ACCESS MANAGEMENT > Identity federation.
- c. Confirm that the **Enable identity federation** checkbox is not selected.
- d. If it is, confirm that any federated groups that might be in use for this tenant account are no longer required, clear the checkbox, and select **Save**.
- 2. Confirm that a federated user can access the Grid Manager:
 - a. From Grid Manager, select CONFIGURATION > Access control > Admin groups.
 - b. Ensure that at least one federated group has been imported from the Active Directory identity source and that it has been assigned the Root access permission.
 - c. Sign out.
 - d. Confirm you can sign back in to the Grid Manager as a user in the federated group.
- 3. If there are existing tenant accounts, confirm that a federated user who has Root access permission can sign in:
 - a. From the Grid Manager, select TENANTS.
 - b. Select the tenant account, and select Actions > Edit.
 - c. On the Enter details tab, select Continue.
 - d. If the Use own identity source checkbox is selected, uncheck the box and select Save.

Edit the tenant Enter details 2 Select permissions
Select permissions for this tenant account. Allow platform services Use own identity source Allow S3 Select

The Tenant page appears.

- e. Select the tenant account, select Sign in, and sign in to the tenant account as the local root user.
- f. From the Tenant Manager, select ACCESS MANAGEMENT > Groups.
- g. Ensure that at least one federated group from the Grid Manager has been assigned the Root access permission for this tenant.
- h. Sign out.
- i. Confirm you can sign back in to the tenant as a user in the federated group.

Related information

- Requirements and considerations for single sign-on
- Manage admin groups
- Use a tenant account

Use sandbox mode

You can use sandbox mode to configure and test single sign-on (SSO) before enabling it for all StorageGRID users. After SSO has been enabled, you can return to sandbox mode whenever you need to change or retest the configuration.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.
- You have configured identity federation for your StorageGRID system.

• For the identity federation **LDAP service type**, you selected either Active Directory or Azure, based on the SSO identity provider you plan to use.

Configured LDAP service type	Options for SSO identity provider
Active Directory	Active Directory
	Azure DingEnderete
	• Fingrederate
Azure	Azure

About this task

When SSO is enabled and a user attempts to sign in to an Admin Node, StorageGRID sends an authentication request to the SSO identity provider. In turn, the SSO identity provider sends an authentication response back to StorageGRID, indicating whether the authentication request was successful. For successful requests:

- The response from Active Directory or PingFederate includes a universally unique identifier (UUID) for the user.
- The response from Azure includes a User Principal Name (UPN).

To allow StorageGRID (the service provider) and the SSO identity provider to communicate securely about user authentication requests, you must configure certain settings in StorageGRID. Next, you must use the SSO identity provider's software to create a relying party trust (AD FS), Enterprise Application (Azure) or Service Provider (PingFederate) for each Admin Node. Finally, you must return to StorageGRID to enable SSO.

Sandbox mode makes it easy to perform this back-and-forth configuration and to test all of your settings before you enable SSO. When you are using sandbox mode, users can't sign in using SSO.

Access sandbox mode

Steps

1. Select CONFIGURATION > Access control > Single sign-on.

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

Single Sign-on
You can enable single sign-on (SSO) if you want an external identity provider (IdP) to authorize all user access to StorageGRID. To start, enable identity federation and confirm that at least one federated user has Root Access permission to the Grid Manager and to the Tenant Manager for any existing tenant accounts. Next, select Sandbox Mode to configure, save, and then test your SSO settings. After verifying the connections, select Enabled and click Save to start using SSO.
SSO status 🥹 💿 Disabled 💿 Sandbox Mode 💿 Enabled
Save



If the SSO Status options don't appear, confirm you have configured the identity provider as the federated identity source. See Requirements and considerations for single sign-on.

2. Select Sandbox Mode.

The Identity Provider section appears.

Enter identity provider details

Steps

- 1. Select the SSO type from the drop-down list.
- 2. Complete the fields in the Identity Provider section based on the SSO type you selected.

Active Directory

a. Enter the **Federation service name** for the identity provider, exactly as it appears in Active Directory Federation Service (AD FS).



To locate the federation service name, go to Windows Server Manager. Select **Tools** > **AD FS Management**. From the Action menu, select **Edit Federation Service Properties**. The Federation Service Name is shown in the second field.

- b. Specify which TLS certificate will be used to secure the connection when the identity provider sends SSO configuration information in response to StorageGRID requests.
 - **Use operating system CA certificate**: Use the default CA certificate installed on the operating system to secure the connection.
 - Use custom CA certificate: Use a custom CA certificate to secure the connection.

If you select this setting, copy the text of the custom certificate and and paste it in the **CA Certificate** text box.

• **Do not use TLS**: Do not use a TLS certificate to secure the connection.



If you change the CA certificate, immediately restart the mgmt-api service on the Admin Nodes and test for a successful SSO into the Grid Manager.

- c. In the Relying Party section, specify the **Relying party identifier** for StorageGRID. This value controls the name you use for each relying party trust in AD FS.
 - For example, if your grid has only one Admin Node and you don't anticipate adding more Admin Nodes in the future, enter SG or StorageGRID.
 - If your grid includes more than one Admin Node, include the string [HOSTNAME] in the identifier. For example, SG-[HOSTNAME]. This generates a table that shows the relying party identifier for each Admin Node in your system, based on the node's hostname.



You must create a relying party trust for each Admin Node in your StorageGRID system. Having a relying party trust for each Admin Node ensures that users can securely sign in to and out of any Admin Node.

d. Select Save.

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



Azure

- a. Specify which TLS certificate will be used to secure the connection when the identity provider sends SSO configuration information in response to StorageGRID requests.
 - **Use operating system CA certificate**: Use the default CA certificate installed on the operating system to secure the connection.
 - Use custom CA certificate: Use a custom CA certificate to secure the connection.

If you select this setting, copy the text of the custom certificate and and paste it in the **CA Certificate** text box.

• **Do not use TLS**: Do not use a TLS certificate to secure the connection.



If you change the CA certificate, immediately restart the mgmt-api service on the Admin Nodes and test for a successful SSO into the Grid Manager.

- b. In the Enterprise Application section, specify the **Enterprise application name** for StorageGRID. This value controls the name you use for each enterprise application in Azure AD.
 - For example, if your grid has only one Admin Node and you don't anticipate adding more Admin Nodes in the future, enter SG or StorageGRID.
 - If your grid includes more than one Admin Node, include the string [HOSTNAME] in the identifier. For example, SG-[HOSTNAME]. This generates a table that shows an enterprise application name for each Admin Node in your system, based on the node's hostname.



You must create an enterprise application for each Admin Node in your StorageGRID system. Having an enterprise application for each Admin Node ensures that users can securely sign in to and out of any Admin Node.

- c. Follow the steps in Create enterprise applications in Azure AD to create an enterprise application for each Admin Node listed in the table.
- d. From Azure AD, copy the federation metadata URL for each enterprise application. Then, paste this URL into the corresponding **Federation metadata URL** field in StorageGRID.
- e. After you have copied and pasted a federation metadata URL for all Admin Nodes, select **Save**.

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



PingFederate

- a. Specify which TLS certificate will be used to secure the connection when the identity provider sends SSO configuration information in response to StorageGRID requests.
 - **Use operating system CA certificate**: Use the default CA certificate installed on the operating system to secure the connection.
 - Use custom CA certificate: Use a custom CA certificate to secure the connection.

If you select this setting, copy the text of the custom certificate and and paste it in the **CA Certificate** text box.

• **Do not use TLS**: Do not use a TLS certificate to secure the connection.



If you change the CA certificate, immediately restart the mgmt-api service on the Admin Nodes and test for a successful SSO into the Grid Manager.

b. In the Service Provider (SP) section, specify the **SP connection ID** for StorageGRID. This value controls the name you use for each SP connection in PingFederate.

- For example, if your grid has only one Admin Node and you don't anticipate adding more Admin Nodes in the future, enter SG or StorageGRID.
- If your grid includes more than one Admin Node, include the string [HOSTNAME] in the identifier. For example, SG-[HOSTNAME]. This generates a table that shows the SP connection ID for each Admin Node in your system, based on the node's hostname.



You must create an SP connection for each Admin Node in your StorageGRID system. Having an SP connection for each Admin Node ensures that users can securely sign in to and out of any Admin Node.

c. Specify the federation metadata URL for each Admin Node in the Federation metadata URL field.

Use the following format:

```
https://<Federation Service
Name>:<port>/pf/federation_metadata.ping?PartnerSpId=<SP Connection
ID>
```

d. Select Save.

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



Configure relying party trusts, enterprise applications, or SP connections

When the configuration is saved, the Sandbox mode confirmation notice appears. This notice confirms that sandbox mode is now enabled and provides overview instructions.

StorageGRID can remain in sandbox mode as long as required. However, when **Sandbox Mode** is selected on the Single Sign-on page, SSO is disabled for all StorageGRID users. Only local users can sign in.

Follow these steps to configure relying party trusts (Active Directory), complete enterprise applications (Azure), or configure SP connections (PingFederate).

Active Directory

Steps

- 1. Go to Active Directory Federation Services (AD FS).
- 2. Create one or more relying party trusts for StorageGRID, using each relying party identifier shown in the table on the StorageGRID Single Sign-on page.

You must create one trust for each Admin Node shown in the table.

For instructions, go to Create relying party trusts in AD FS.

Azure

Steps

- 1. From the Single sign-on page for the Admin Node you are currently signed in to, select the button to download and save the SAML metadata.
- 2. Then, for any other Admin Nodes in your grid, repeat these steps:
 - a. Sign in to the node.
 - b. Select CONFIGURATION > Access control > Single sign-on.
 - c. Download and save the SAML metadata for that node.
- 3. Go to the Azure Portal.
- 4. Follow the steps in Create enterprise applications in Azure AD to upload the SAML metadata file for each Admin Node into its corresponding Azure enterprise application.

PingFederate

Steps

- 1. From the Single sign-on page for the Admin Node you are currently signed in to, select the button to download and save the SAML metadata.
- 2. Then, for any other Admin Nodes in your grid, repeat these steps:
 - a. Sign in to the node.
 - b. Select CONFIGURATION > Access control > Single sign-on.
 - c. Download and save the SAML metadata for that node.
- 3. Go to PingFederate.
- Create one or more service provider (SP) connections for StorageGRID. Use the SP connection ID for each Admin Node (shown in the table on the StorageGRID Single Sign-on page) and the SAML metadata you downloaded for that Admin Node.

You must create one SP connection for each Admin Node shown in the table.

Test SSO connections

Before you enforce the use of single sign-on for your entire StorageGRID system, you should confirm that single sign-on and single logout are correctly configured for each Admin Node.

Active Directory

Steps

1. From the StorageGRID Single Sign-on page, locate the link in the Sandbox mode message.

The URL is derived from the value you entered in the Federation service name field.

Sandbox mode	
Sandbox mode is currently enabled. Use this mode to configure relying party trusts and to confirm that single sign-on (SSO) and single logout (SLO) are correctly configured for the StorageGRID system.	
1. Use Active Directory Federation Services (AD FS) to create relying party trusts for StorageGRID. Create one trust for each Admin Node, using the relying party identifier(s) shown below.	
2. Go to your identity provider's sign-on page: https://ad2016.saml.sgws/adfs/ls/idpinitiatedsignon.htm	
3. From this page, sign in to each StorageGRID relying party trust. If the SSO operation is successful, StorageGRID displays a page with a success message. Otherwise, an error message is displayed.	
When you have confirmed SSO for each of the relying party trusts and you are ready to enforce the use of SSO for StorageGRID, change the SSO Status to Enabled, and click Save.	

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

You are not signed in.	
 Sign in to this site. Sign in to one of the following sites: 	
SG-DC1-ADM1	v
Sign in	

- 4. Enter your federated username and password.
 - If the SSO sign-in and logout operations are successful, a success message appears.

Single sign-on authentication and logout test completed successfully.

- If the SSO operation is unsuccessful, an error message appears. Fix the issue, clear the browser's cookies, and try again.
- 5. Repeat these steps to verify the SSO connection for each Admin Node in your grid.

Azure

Steps

- 1. Go to the Single sign-on page in the Azure portal.
- 2. Select Test this application.
- 3. Enter the credentials of a federated user.
 - If the SSO sign-in and logout operations are successful, a success message appears.

Single sign-on authentication and logout test completed successfully.

- If the SSO operation is unsuccessful, an error message appears. Fix the issue, clear the browser's cookies, and try again.
- 4. Repeat these steps to verify the SSO connection for each Admin Node in your grid.

PingFederate

Steps

1. From the StorageGRID Single Sign-on page, select the first link in the Sandbox mode message.

Select and test one link at a time.

- 2. Enter the credentials of a federated user.
 - If the SSO sign-in and logout operations are successful, a success message appears.

Single sign-on authentication and logout test completed successfully.

- If the SSO operation is unsuccessful, an error message appears. Fix the issue, clear the browser's cookies, and try again.
- 3. Select the next link to verify the SSO connection for each Admin Node in your grid.

If you see a Page Expired message, select the **Back** button in your browser and resubmit your credentials.

Enable single sign-on

When you have confirmed you can use SSO to sign in to each Admin Node, you can enable SSO for your entire StorageGRID system.



When SSO is enabled, all users must use SSO to access the Grid Manager, the Tenant Manager, the Grid Management API, and the Tenant Management API. Local users can no longer access StorageGRID.

Steps

- 1. Select CONFIGURATION > Access control > Single sign-on.
- 2. Change the SSO Status to Enabled.
- 3. Select Save.
- 4. Review the warning message, and select **OK**.

Single sign-on is now enabled.



If you are using the Azure Portal and you access StorageGRID from the same computer you use to access Azure, ensure that the Azure Portal user is also an authorized StorageGRID user (a user in a federated group that has been imported into StorageGRID) or log out of the Azure Portal before attempting to sign in to StorageGRID.

Create relying party trusts in AD FS

You must use Active Directory Federation Services (AD FS) to create a relying party trust for each Admin Node in your system. You can create relying party trusts using PowerShell commands, by importing SAML metadata from StorageGRID, or by entering the data manually.

Before you begin

- You have configured single sign-on for StorageGRID and you selected AD FS as the SSO type.
- Sandbox mode is selected on the Single sign-on page in Grid Manager. See Use sandbox mode.
- You know the fully qualified domain name (or the IP address) and the relying party identifier for each Admin Node in your system. You can find these values in the Admin Nodes detail table on the StorageGRID Single Sign-on page.



You must create a relying party trust for each Admin Node in your StorageGRID system. Having a relying party trust for each Admin Node ensures that users can securely sign in to and out of any Admin Node.

- You have experience creating relying party trusts in AD FS, or you have access to the Microsoft AD FS documentation.
- You are using the AD FS Management snap-in, and you belong to the Administrators group.
- If you are creating the relying party trust manually, you have the custom certificate that was uploaded for the StorageGRID management interface, or you know how to log in to an Admin Node from the command shell.

About this task

These instructions apply to Windows Server 2016 AD FS. If you are using a different version of AD FS, you will notice slight differences in the procedure. See the Microsoft AD FS documentation if you have questions.

Create a relying party trust using Windows PowerShell

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

Steps

- 1. From the Windows start menu, right-select the PowerShell icon, and select **Run as Administrator**.
- 2. At the PowerShell command prompt, enter the following command:

`Add-AdfsRelyingPartyTrust -Name "*Admin_Node_Identifer*" -MetadataURL "https:// *Admin_Node_FQDN*/api/samI-metadata"

• For *Admin_Node_Identifier*, enter the Relying Party Identifier for the Admin Node, exactly as it appears on the Single Sign-on page. For example, SG-DC1-ADM1.

- For Admin_Node_FQDN, enter the fully qualified domain name for the same Admin Node. (If necessary, you can use the node's IP address instead. However, if you enter an IP address here, be aware that you must update or recreate this relying party trust if that IP address ever changes.)
- 3. From Windows Server Manager, select **Tools > AD FS Management**.

The AD FS management tool appears.

4. Select **AD FS > Relying Party Trusts**.

The list of relying party trusts appears.

- 5. Add an Access Control Policy to the newly created relying party trust:
 - a. Locate the relying party trust you just created.
 - b. Right-click the trust, and select Edit Access Control Policy.
 - c. Select an Access Control Policy.
 - d. Select Apply, and select OK
- 6. Add a Claim Issuance Policy to the newly created Relying Party Trust:
 - a. Locate the relying party trust you just created.
 - b. Right-click the trust, and select Edit claim issuance policy.
 - c. Select Add rule.
 - d. On the Select Rule Template page, select **Send LDAP Attributes as Claims** from the list, and select **Next**.
 - e. On the Configure Rule page, enter a display name for this rule.

For example, ObjectGUID to Name ID or UPN to Name ID.

- f. For the Attribute Store, select Active Directory.
- g. In the LDAP Attribute column of the Mapping table, type **objectGUID** or select **User-Principal-Name**.
- h. In the Outgoing Claim Type column of the Mapping table, select Name ID from the drop-down list.
- i. Select **Finish**, and select **OK**.
- 7. Confirm that the metadata was imported successfully.
 - a. Right-click the relying party trust to open its properties.
 - b. Confirm that the fields on the Endpoints, Identifiers, and Signature tabs are populated.

If the metadata is missing, confirm that the Federation metadata address is correct, or enter the values manually.

- Repeat these steps to configure a relying party trust for all of the Admin Nodes in your StorageGRID system.
- 9. When you are done, return to StorageGRID and test all relying party trusts to confirm they are configured correctly. See Use Sandbox mode for instructions.

Create a relying party trust by importing federation metadata

You can import the values for each relying party trust by accessing the SAML metadata for each Admin Node.

Steps

- 1. In Windows Server Manager, select **Tools**, and then select **AD FS Management**.
- 2. Under Actions, select Add Relying Party Trust.
- 3. On the Welcome page, choose Claims aware, and select Start.
- 4. Select Import data about the relying party published online or on a local network.
- 5. In **Federation metadata address (host name or URL)**, type the location of the SAML metadata for this Admin Node:

https://Admin_Node_FQDN/api/saml-metadata

For *Admin_Node_FQDN*, enter the fully qualified domain name for the same Admin Node. (If necessary, you can use the node's IP address instead. However, if you enter an IP address here, be aware that you must update or recreate this relying party trust if that IP address ever changes.)

6. Complete the Relying Party Trust wizard, save the relying party trust, and close the wizard.



When entering the display name, use the Relying Party Identifier for the Admin Node, exactly as it appears on the Single Sign-on page in the Grid Manager. For example, SG-DC1-ADM1.

- 7. Add a claim rule:
 - a. Right-click the trust, and select Edit claim issuance policy.
 - b. Select Add rule:
 - c. On the Select Rule Template page, select **Send LDAP Attributes as Claims** from the list, and select **Next**.
 - d. On the Configure Rule page, enter a display name for this rule.

For example, ObjectGUID to Name ID or UPN to Name ID.

- e. For the Attribute Store, select Active Directory.
- f. In the LDAP Attribute column of the Mapping table, type **objectGUID** or select **User-Principal-Name**.
- g. In the Outgoing Claim Type column of the Mapping table, select Name ID from the drop-down list.
- h. Select Finish, and select OK.
- 8. Confirm that the metadata was imported successfully.
 - a. Right-click the relying party trust to open its properties.
 - b. Confirm that the fields on the Endpoints, Identifiers, and Signature tabs are populated.

If the metadata is missing, confirm that the Federation metadata address is correct, or enter the values manually.

- Repeat these steps to configure a relying party trust for all of the Admin Nodes in your StorageGRID system.
- 10. When you are done, return to StorageGRID and test all relying party trusts to confirm they are configured correctly. See Use Sandbox mode for instructions.

Create a relying party trust manually

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

Steps

- 1. In Windows Server Manager, select Tools, and then select AD FS Management.
- 2. Under Actions, select Add Relying Party Trust.
- 3. On the Welcome page, choose Claims aware, and select Start.
- 4. Select Enter data about the relying party manually, and select Next.
- 5. Complete the Relying Party Trust wizard:
 - a. Enter a display name for this Admin Node.

For consistency, use the Relying Party Identifier for the Admin Node, exactly as it appears on the Single Sign-on page in the Grid Manager. For example, SG-DC1-ADM1.

- b. Skip the step to configure an optional token encryption certificate.
- c. On the Configure URL page, select the **Enable support for the SAML 2.0 WebSSO protocol** checkbox.
- d. Type the SAML service endpoint URL for the Admin Node:

https://Admin_Node_FQDN/api/saml-response

For *Admin_Node_FQDN*, enter the fully qualified domain name for the Admin Node. (If necessary, you can use the node's IP address instead. However, if you enter an IP address here, be aware that you must update or recreate this relying party trust if that IP address ever changes.)

e. On the Configure Identifiers page, specify the Relying Party Identifier for the same Admin Node:

Admin_Node_Identifier

For *Admin_Node_Identifier*, enter the Relying Party Identifier for the Admin Node, exactly as it appears on the Single Sign-on page. For example, SG-DC1-ADM1.

f. Review the settings, save the relying party trust, and close the wizard.

The Edit Claim Issuance Policy dialog box appears.



If the dialog box does not appear, right-click the trust, and select **Edit claim issuance policy**.

- 6. To start the Claim Rule wizard, select Add rule:
 - a. On the Select Rule Template page, select **Send LDAP Attributes as Claims** from the list, and select **Next**.
 - b. On the Configure Rule page, enter a display name for this rule.

For example, ObjectGUID to Name ID or UPN to Name ID.

- c. For the Attribute Store, select Active Directory.
- d. In the LDAP Attribute column of the Mapping table, type **objectGUID** or select **User-Principal-Name**.
- e. In the Outgoing Claim Type column of the Mapping table, select Name ID from the drop-down list.
- f. Select Finish, and select OK.

- 7. Right-click the relying party trust to open its properties.
- 8. On the **Endpoints** tab, configure the endpoint for single logout (SLO):
 - a. Select Add SAML.
 - b. Select Endpoint Type > SAML Logout.
 - c. Select **Binding > Redirect**.
 - d. In the **Trusted URL** field, enter the URL used for single logout (SLO) from this Admin Node:

https://Admin_Node_FQDN/api/saml-logout

For *Admin_Node_FQDN*, enter the Admin Node's fully qualified domain name. (If necessary, you can use the node's IP address instead. However, if you enter an IP address here, be aware that you must update or recreate this relying party trust if that IP address ever changes.)

- e. Select **OK**.
- 9. On the **Signature** tab, specify the signature certificate for this relying party trust:
 - a. Add the custom certificate:
 - If you have the custom management certificate you uploaded to StorageGRID, select that certificate.
 - If you don't have the custom certificate, log in to the Admin Node, go the /var/local/mgmt-api directory of the Admin Node, and add the custom-server.crt certificate file.

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. Select Apply, and select OK.

The Relying Party properties are saved and closed.

- 10. Repeat these steps to configure a relying party trust for all of the Admin Nodes in your StorageGRID system.
- 11. When you are done, return to StorageGRID and test all relying party trusts to confirm they are configured correctly. See Use sandbox mode for instructions.

Create enterprise applications in Azure AD

You use Azure AD to create an enterprise application for each Admin Node in your system.

Before you begin

- You have started configuring single sign-on for StorageGRID and you selected **Azure** as the SSO type.
- Sandbox mode is selected on the Single sign-on page in Grid Manager. See Use sandbox mode.
- You have the **Enterprise application name** for each Admin Node in your system. You can copy these values from the Admin Node details table on the StorageGRID Single Sign-on page.



You must create an enterprise application for each Admin Node in your StorageGRID system. Having an enterprise application for each Admin Node ensures that users can securely sign in to and out of any Admin Node.

- You have experience creating enterprise applications in Azure Active Directory.
- You have an Azure account with an active subscription.
- You have one of the following roles in the Azure account: Global Administrator, Cloud Application Administrator, or owner of the service principal.

Access Azure AD

Steps

- 1. Log in to the Azure Portal.
- 2. Navigate to Azure Active Directory.
- 3. Select Enterprise applications.

Create enterprise applications and save StorageGRID SSO configuration

To save the SSO configuration for Azure in StorageGRID, you must use Azure to create an enterprise application for each Admin Node. You will copy the federation metadata URLs from Azure and paste them into the corresponding **Federation metadata URL** fields on the StorageGRID Single Sign-on page.

Steps

- 1. Repeat the following steps for each Admin Node.
 - a. In the Azure Enterprise applications pane, select New application.
 - b. Select Create your own application.
 - c. For the name, enter the **Enterprise application name** you copied from the Admin Node details table on the StorageGRID Single Sign-on page.
 - d. Leave the **Integrate any other application you don't find in the gallery (Non-gallery)** radio button selected.
 - e. Select Create.
 - f. Select the **Get started** link in the **2. Set up single sign on** box, or select the **Single sign-on** link in the left margin.
 - g. Select the **SAML** box.
 - h. Copy the App Federation Metadata Url, which you can find under Step 3 SAML Signing Certificate.
 - i. Go to the StorageGRID Single Sign-on page, and paste the URL in the **Federation metadata URL** field that corresponds to the **Enterprise application name** you used.
- 2. After you have pasted a federation metadata URL for each Admin Node and made all other needed changes to the SSO configuration, select **Save** on the StorageGRID Single Sign-on page.

Download SAML metadata for every Admin Node

After the SSO configuration is saved, you can download a SAML metadata file for each Admin Node in your StorageGRID system.

Steps

- 1. Repeat these steps for each Admin Node.
 - a. Sign in to StorageGRID from the Admin Node.
 - b. Select CONFIGURATION > Access control > Single sign-on.

- c. Select the button to download the SAML metadata for that Admin Node.
- d. Save the file, which you will upload into Azure AD.

Upload SAML metadata to each enterprise application

After downloading a SAML metadata file for each StorageGRID Admin Node, perform the following steps in Azure AD:

Steps

- 1. Return to the Azure Portal.
- 2. Repeat these steps for each enterprise application:



You might need to refresh the Enterprise applications page to see applications you previously added in the list.

- a. Go to the Properties page for the enterprise application.
- b. Set Assignment required to No (unless you want to separately configure assignments).
- c. Go to the Single sign-on page.
- d. Complete the SAML configuration.
- e. Select the **Upload metadata file** button and select the SAML metadata file you downloaded for the corresponding Admin Node.
- f. After the file loads, select **Save** and then select **X** to close the pane. You are returned to the Set up Single Sign-On with SAML page.
- 3. Follow the steps in Use sandbox mode to test each application.

Create service provider (SP) connections in PingFederate

You use PingFederate to create a service provider (SP) connection for each Admin Node in your system. To speed up the process, you will import the SAML metadata from StorageGRID.

Before you begin

- You have configured single sign-on for StorageGRID and you selected **Ping Federate** as the SSO type.
- Sandbox mode is selected on the Single sign-on page in Grid Manager. See Use sandbox mode.
- You have the **SP connection ID** for each Admin Node in your system. You can find these values in the Admin Nodes detail table on the StorageGRID Single Sign-on page.
- You have downloaded the SAML metadata for each Admin Node in your system.
- · You have experience creating SP connections in PingFederate Server.
- You have the Administrator's Reference Guide for PingFederate Server. The PingFederate documentation provides detailed step-by-step instructions and explanations.
- You have the Admin permission for PingFederate Server.

About this task

These instructions summarize how to configure PingFederate Server version 10.3 as an SSO provider for StorageGRID. If you are using another version of PingFederate, you might need to adapt these instructions.
Refer to the PingFederate Server documentation for detailed instructions for your release.

Complete prerequisites in PingFederate

Before you can create the SP connections you will use for StorageGRID, you must complete prerequisite tasks in PingFederate. You will use information from these prerequisites when you configure the SP connections.

Create data store

If you haven't already, create a data store to connect PingFederate to the AD FS LDAP server. Use the values you used when configuring identity federation in StorageGRID.

- Type: Directory (LDAP)
- LDAP Type: Active Directory
- Binary Attribute Name: Enter objectGUID on the LDAP Binary Attributes tab exactly as shown.

Create password credential validator

If you haven't already, create a password credential validator.

- Type: LDAP Username Password Credential Validator
- Data store: Select the data store you created.
- Search base: Enter information from LDAP (for example, DC=saml, DC=sgws).
- Search filter: sAMAccountName=\${username}
- Scope: Subtree

Create IdP adapter instance

If you haven't already, create an IdP adapter instance.

Steps

- 1. Go to Authentication > Integration > IdP Adapters.
- 2. Select Create New Instance.
- 3. On the Type tab, select HTML Form IdP Adapter.
- 4. On the IdP Adapter tab, select Add a new row to 'Credential Validators'.
- 5. Select the password credential validator you created.
- 6. On the Adapter Attributes tab, select the **username** attribute for **Pseudonym**.
- 7. Select Save.

Create or import signing certificate

If you haven't already, create or import the signing certificate.

Steps

- 1. Go to Security > Signing & Decryption Keys & Certificates.
- 2. Create or import the signing certificate.

Create an SP connection in PingFederate

When you create an SP connection in PingFederate, you import the SAML metadata you downloaded from StorageGRID for the Admin Node. The metadata file contains many of the specific values you need.



You must create an SP connection for each Admin Node in your StorageGRID system, so that users can securely sign in to and out of any node. Use these instructions to create the first SP connection. Then, go to Create additional SP connections to create any additional connections you need.

Choose SP connection type

Steps

- 1. Go to **Applications > Integration > SP Connections**.
- 2. Select Create Connection.
- 3. Select **Do not use a template for this connection**.
- 4. Select Browser SSO Profiles and SAML 2.0 as the protocol.

Import SP metadata

Steps

- 1. On the Import Metadata tab, select File.
- Choose the SAML metadata file you downloaded from the StorageGRID Single sign-on page for the Admin Node.
- 3. Review the Metadata Summary and the information provided on the General Info tab.

The Partner's Entity ID and the Connection Name are set to the StorageGRID SP connection ID. (for example, 10.96.105.200-DC1-ADM1-105-200). The Base URL is the IP of the StorageGRID Admin Node.

4. Select Next.

Configure IdP Browser SSO

Steps

- 1. From the Browser SSO tab, select Configure Browser SSO.
- 2. On the SAML profiles tab, select the **SP-initiated SSO**, **SP-initial SLO**, **IdP-initiated SSO**, and **IdP-initiated SLO** options.
- 3. Select Next.
- 4. On the Assertion Lifetime tab, make no changes.
- 5. On the Assertion Creation tab, select **Configure Assertion Creation**.
 - a. On the Identity Mapping tab, select Standard.
 - b. On the Attribute Contract tab, use the **SAML_SUBJECT** as the Attribute Contract and the unspecified name format that was imported.
- 6. For Extend the Contract, select **Delete** to remove the urn:oid, which is not used.

Map adapter instance

Steps

- 1. On the Authentication Source Mapping tab, select Map New Adapter Instance.
- 2. On the Adapter instance tab, select the adapter instance you created.
- 3. On the Mapping Method tab, select Retrieve Additional Attributes From a Data Store.
- 4. On the Attribute Source & User Lookup tab, select Add Attribute Source.
- 5. On the Data Store tab, provide a description and select the data store you added.
- 6. On the LDAP Directory Search tab:
 - Enter the **Base DN**, which should exactly match the value you entered in StorageGRID for the LDAP server.
 - For the Search Scope, select Subtree.
 - For the Root Object Class, search for and add either of these attributes: objectGUID or userPrincipalName.
- 7. On the LDAP Binary Attribute Encoding Types tab, select Base64 for the objectGUID attribute.
- 8. On the LDAP Filter tab, enter **sAMAccountName=\${username}**.
- 9. On the Attribute Contract Fulfillment tab, select **LDAP (attribute)** from the Source drop-down and select either **objectGUID** or **userPrincipalName** from the Value drop-down.
- 10. Review and then save the attribute source.
- 11. On the Failsave Attribute Source tab, select Abort the SSO Transaction.
- 12. Review the summary and select **Done**.
- 13. Select Done.

Configure protocol settings

Steps

- 1. On the SP Connection > Browser SSO > Protocol Settings tab, select Configure Protocol Settings.
- 2. On the Assertion Consumer Service URL tab, accept the default values, which were imported from the StorageGRID SAML metadata (**POST** for Binding and /api/saml-response for Endpoint URL).
- 3. On the SLO Service URLs tab, accept the default values, which were imported from the StorageGRID SAML metadata (**REDIRECT** for Binding and /api/saml-logout for Endpoint URL.
- 4. On the Allowable SAML Bindings tab, clear **ARTIFACT** and **SOAP**. Only **POST** and **REDIRECT** are required.
- 5. On the Signature Policy tab, leave the **Require Authn Requests to be Signed** and **Always Sign Assertion** checkboxes selected.
- 6. On the Encryption Policy tab, select None.
- 7. Review the summary and select **Done** to save the protocol settings.
- 8. Review the summary and select **Done** to save the Browser SSO settings.

Configure credentials

Steps

1. From the SP Connection tab, select Credentials.

- 2. From the Credentials tab, select Configure Credentials.
- 3. Select the signing certificate you created or imported.
- 4. Select Next to go to Manage Signature Verification Settings.
 - a. On the Trust Model tab, select Unanchored.
 - b. On the Signature Verification Certificate tab, review the signing certificate information, which was imported from the StorageGRID SAML metadata.
- 5. Review the summary screens and select **Save** to save the SP connection.

Create additional SP connections

You can copy the first SP connection to create the SP connections you need for each Admin Node in your grid. You upload new metadata for each copy.



The SP connections for different Admin Nodes use identical settings, with the exception of the Partner's Entity ID, Base URL, Connection ID, Connection Name, Signature Verification, and SLO Response URL.

Steps

- 1. Select **Action** > **Copy** to create a copy of the initial SP connection for each additional Admin Node.
- 2. Enter the Connection ID and Connection Name for the copy, and select Save.
- 3. Choose the metadata file corresponding to the Admin Node:
 - a. Select Action > Update with Metadata.
 - b. Select Choose File and upload the metadata.
 - c. Select Next.
 - d. Select Save.
- 4. Resolve the error due to the unused attribute:
 - a. Select the new connection.
 - b. Select Configure Browser SSO > Configure Assertion Creation > Attribute Contract.
 - c. Delete the entry for **urn:oid**.
 - d. Select Save.

Disable single sign-on

You can disable single sign-on (SSO) if you no longer want to use this functionality. You must disable single sign-on before you can disable identity federation.

Before you begin

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

Steps

1. Select CONFIGURATION > Access control > Single sign-on.

The Single Sign-on page appears.

- 2. Select the **Disabled** option.
- 3. Select Save.

A warning message appears indicating that local users will now be able to sign in.

4. Select OK.

The next time you sign in to StorageGRID, the StorageGRID Sign in page appears and you must enter the username and password for a local or federated StorageGRID user.

Temporarily disable and reenable single sign-on for one Admin Node

You might not be able to sign in to the Grid Manager if the single sign-on (SSO) system goes down. In this case, you can temporarily disable and reenable SSO for one Admin Node. To disable and then reenable SSO, you must access the node's command shell.

Before you begin

- · You have specific access permissions.
- You have the Passwords.txt file.
- You know the password for the local root user.

About this task

After you disable SSO for one Admin Node, you can sign in to the Grid Manager as the local root user. To secure your StorageGRID system, you must use the node's command shell to reenable SSO on the Admin Node as soon as you sign out.



Disabling SSO for one Admin Node does not affect the SSO settings for any other Admin Nodes in the grid. The **Enable SSO** checkbox on the Single Sign-on page in the Grid Manager remains selected, and all existing SSO settings are maintained unless you update them.

Steps

- 1. Log in to an Admin Node:
 - a. Enter the following command: ssh admin@Admin_Node_IP
 - b. Enter the password listed in the Passwords.txt file.
 - c. Enter the following command to switch to root: su -
 - d. Enter the password listed in the ${\tt Passwords.txt}$ file.

When you are logged in as root, the prompt changes from \$ to #.

2. Run the following command:disable-saml

A message indicates that the command applies to this Admin Node only.

3. Confirm that you want to disable SSO.

A message indicates that single sign-on is disabled on the node.

4. From a web browser, access the Grid Manager on the same Admin Node.

The Grid Manager sign-in page is now displayed because SSO has been disabled.

- 5. Sign in with the username root and the local root user's password.
- 6. If you disabled SSO temporarily because you needed to correct the SSO configuration:
 - a. Select **CONFIGURATION > Access control > Single sign-on**.
 - b. Change the incorrect or out-of-date SSO settings.
 - c. Select Save.

Selecting Save from the Single Sign-on page automatically reenables SSO for the entire grid.

- 7. If you disabled SSO temporarily because you needed to access the Grid Manager for some other reason:
 - a. Perform whatever task or tasks you need to perform.
 - b. Select Sign out, and close the Grid Manager.
 - c. Reenable SSO on the Admin Node. You can perform either of the following steps:
 - Run the following command: enable-saml

A message indicates that the command applies to this Admin Node only.

Confirm that you want to enable SSO.

A message indicates that single sign-on is enabled on the node.

- Reboot the grid node: reboot
- 8. From a web browser, access the Grid Manager from the same Admin Node.
- 9. Confirm that the StorageGRID Sign in page appears and that you must enter your SSO credentials to access the Grid Manager.

Use grid federation

What is grid federation?

You can use grid federation to clone tenants and replicate their objects between two StorageGRID systems for disaster recovery.

What is a grid federation connection?

A grid federation connection is a bidirectional, trusted, and secure connection between Admin and Gateway Nodes in two StorageGRID systems.

Workflow for grid federation

The workflow diagram summarize the steps for configuring a grid federation connection between two grids.



Considerations and requirements for grid federation connections

- Both grids used for grid federation must be running StorageGRID 11.7 or later.
- A grid can have one or more grid federation connections to other grids. Each grid federation connection is independent of any other connections. For example, if Grid 1 has one connection with Grid 2 and a second connection with Grid 3, there is no implied connection between Grid 2 and Grid 3.
- Grid federation connections are bidirectional. After the connection is established, you can monitor and manage the connection from either grid.
- At least one grid federation connection must exist before you can use account clone or cross-grid replication.

Networking and IP address requirements

- Grid federation connections can occur on the Grid Network, Admin Network, or Client Network.
- A grid federation connection connects one grid to another grid. The configuration for each grid specifies a grid federation endpoint on the other grid that consists of Admin Nodes, Gateway Nodes, or both.
- The best practice is to connect high availability (HA) groups of Gateway and Admin Nodes on each grid. Using HA groups helps ensure that grid federation connections will remain online if nodes become unavailable. If the active interface in either HA group fails, the connection can use a backup interface.
- Creating a grid federation connection that uses the IP address of a single Admin Node or Gateway Node is not recommended. If the node becomes unavailable, the grid federation connection will also become unavailable.

• Cross-grid replication of objects requires that the Storage Nodes on each grid be able to access the configured Admin and Gateway Nodes on the other grid. For each grid, confirm that all Storage Nodes have a high bandwidth route to as the Admin Nodes or Gateway Nodes used for the connection.

Use FQDNs to load balance the connection

For a production environment, use fully qualified domain names (FQDNs) to identify each grid in the connection. Then, create the appropriate DNS entries, as follows:

- The FQDN for Grid 1 mapped to one or more virtual IP (VIP) addresses for HA groups in Grid 1 or to the IP address of one or more Admin or Gateway Nodes in Grid 1.
- The FQDN for Grid 2 mapped to one or more VIP addresses for Grid 2 or to the IP address of one or more Admin or Gateway Nodes in Grid 2.

When you use multiple DNS entries, requests to use the connection are load balanced, as follows:

- DNS entries that map to the VIP addresses of multiple HA groups are load balanced between the active nodes in the HA groups.
- DNS entries that map to the IP addresses of multiple Admin Nodes or Gateway Nodes are load balanced between the mapped nodes.

Port requirements

When creating a grid federation connection, you can specify any unused port number from 23000 to 23999. Both grids in this connection will use the same port.

You must ensure that no node in either grid uses this port for other connections.

Certificate requirements

When you configure a grid federation connection, StorageGRID automatically generates four SSL certificates:

- Server and client certificates to authenticate and encrypt information sent from grid 1 to grid 2
- · Server and client certificates to authenticate and encrypt information sent from grid 2 to grid 1



By default, the certificates are valid for 730 days (2 years). When these certificates near their expiration date, the **Expiration of grid federation certificate** alert reminds you to rotate the certificates, which you can do using the Grid Manager.



If the certificates on either end of the connection expire, the connection will stop working. Data replication will be pending until the certificates are updated.

Learn more

- Create grid federation connections
- Manage grid federation connections
- Troubleshoot grid federation errors

What is account clone?

Account clone is the automatic replication of a tenant account, tenant groups, tenant users, and, optionally, S3 access keys between the StorageGRID systems in a grid federation connection.

Account clone is required for cross-grid replication. Cloning account information from a source StorageGRID system to a destination StorageGRID system ensures that tenant users and groups can access the corresponding buckets and objects on either grid.

Workflow for account clone

The workflow diagram shows the steps that grid administrators and permitted tenants will perform to set up account clone. These steps are performed after the grid federation connection is configured.



Grid admin workflow

The steps that grid admins perform depend on whether the StorageGRID systems in the grid federation connection use single sign-on (SSO) or identity federation.

Configure SSO for account clone (optional)

If either StorageGRID system in the grid federation connection uses SSO, both grids must use SSO. Before creating the tenant accounts for grid federation, the grid admins for the tenant's source and destination grids must perform these steps.

Steps

- 1. Configure the same identity source for both grids. See Use identity federation.
- 2. Configure the same SSO identity provider (IdP) for both grids. See Configure single sign-on.
- 3. Create the same admin group on both grids by importing the same federated group.

When you create the tenant, you will select this group to have the initial Root access permission for both the source and destination tenant accounts.



If this admin group doesn't exist on both grids before you create the tenant, the tenant isn't replicated to the destination.

Configure grid-level identity federation for account clone (optional)

If either StorageGRID system uses identity federation without SSO, both grids must use identity federation. Before creating the tenant accounts for grid federation, the grid admins for the tenant's source and destination grids must perform these steps.

Steps

- 1. Configure the same identity source for both grids. See Use identity federation.
- 2. Optionally, if a federated group will have initial Root access permission for both the source and destination tenant accounts, create the same admin group on both grids by importing the same federated group.



If you assign Root access permission to a federated group that doesn't exist on both grids, the tenant isn't replicated to the destination grid.

3. If you don't want a federated group to have initial Root access permission for both accounts, specify a password for the local root user.

Create permitted S3 tenant account

After optionally configuring SSO or identity federation, a grid admin performs these steps to determine which tenants can replicate bucket objects to other StorageGRID systems.

Steps

1. Determine which grid you want to be the tenant's source grid for account clone operations.

The grid where the tenant is originally created is known as the tenant's *source grid*. The grid where the tenant is replicated is known as the tenant's *destination grid*.

- 2. On that grid, create a new S3 tenant account or edit an existing account.
- 3. Assign the Use grid federation connection permission.
- 4. If the tenant account will manage its own federated users, assign the **Use own identity source** permission.

If this permission is assigned, both the source and destination tenant accounts must configure the same identity source before creating federated groups. Federated groups added to the source tenant can't be cloned to the destination tenant unless both grids use the same identity source.

- 5. Select a specific grid federation connection.
- 6. Save the new or modified tenant.

When a new tenant with the **Use grid federation connection** permission is saved, StorageGRID automatically creates a replica of that tenant on the other grid, as follows:

- Both tenant accounts have the same account ID, name, storage quota, and assigned permissions.
- If you selected a federated group to have Root access permission for the tenant, that group is cloned to the destination tenant.
- If you selected a local user to have Root access permission for the tenant, that user is cloned to the destination tenant. However, the password for that user is not cloned.

For details, see

Manage permitted tenants for grid federation.

Permitted tenant account workflow

After a tenant with the **Use grid federation connection** permission is replicated to the destination grid, permitted tenant accounts can perform these steps to clone tenant groups, users, and S3 access keys.

Steps

- 1. Sign in to the tenant account on the tenant's source grid.
- 2. If permitted, configure identify federation on both the source and destination tenant accounts.
- 3. Create groups and users on the source tenant.

When new groups or users are created on the source tenant, StorageGRID automatically clones them to the destination tenant, but no cloning occurs from the destination back to the source.

- 4. Create S3 access keys.
- 5. Optionally, clone S3 access keys from the source tenant to the destination tenant.

For details about the permitted tenant account workflow and to learn how groups, users, and S3 access keys are cloned, see Clone tenant groups and users and Clone S3 access keys using the API.

What is cross-grid replication?

Cross-grid replication is the automatic replication of objects between selected S3 buckets in two StorageGRID systems that are connected in a grid federation connection. Account clone is required for cross-grid replication.

Workflow for cross-grid replication

The workflow diagram summarize the steps for configuring cross-grid replication between buckets on two grids.



Requirements for cross-grid replication

If a tenant account has the **Use grid federation connection** permission to use one or more grid federation connections, a tenant user with Root access permission can create identical buckets in the corresponding tenant accounts on each grid. These buckets:

- Must have the same name but can have different regions
- · Must have versioning enabled
- Must have S3 Object Lock disabled
- · Must be empty

After both buckets have been created, cross-grid replication can be configured for either or both buckets.

Learn more

Manage cross-grid replication

How cross-grid replication works

Cross-grid replication can be configured to occur in one direction or in both directions.

Replication in one direction

If you enable cross-grid replication for a bucket on only one grid, objects added to that bucket (the source bucket) are replicated to the corresponding bucket on the other grid (the destination bucket). However, objects added to the destination bucket aren't replicated back to the source. In the figure, cross-grid replication is enabled for my-bucket from Grid 1 to Grid 2, but it is not enabled in the other direction.



Replication in both directions

If you enable cross-grid replication for the same bucket on both grids, objects added to either bucket are replicated to the other grid. In the figure, cross-grid replication is enabled for my-bucket in both directions.



What happens when objects are ingested?

When an S3 client adds an object to a bucket that has cross-grid replication enabled, the following happens:

1. StorageGRID automatically replicates the object from the source bucket to the destination bucket. The time to perform this background replication operation depends on several factors, including the number of other replication operations that are pending.

The S3 client can verify an object's replication status by issuing a GetObject or HeadObject request. The response includes a StorageGRID-specific x-ntap-sg-cgr-replication-status response header, which will have one of the following values:

The S3 client can verify an object's replication status by issuing a GetObject or HeadObject request. The response includes a StorageGRID-specific x-ntap-sg-cgr-replication-status response header, which will have one of the following values:

Grid	Replication status	
Source	 SUCCESS: The replication was successful for all grid connections. 	
	 PENDING: The object hasn't been replicated to at least one grid connection. 	
	• FAILURE : Replication is not pending for any grid connection and at least one failed with a permanent failure. A user must resolve the error.	
Destination	REPLICA : The object was replicated from the source grid.	



StorageGRID does not support the x-amz-replication-status header.

2. StorageGRID uses each grid's active ILM policies to manage the objects, just as it would any other object. For example, Object A on Grid 1 might be stored as two replicated copies and retained forever, while the copy of Object A that was replicated to Grid 2 might be stored using 2+1 erasure coding and deleted after three years.

What happens when objects are deleted?

As described in Delete data flow, StorageGRID can delete an object for any of these reasons:

- The S3 client issues a delete request.
- A Tenant Manager user selects the Delete objects in bucket option to remove all objects from a bucket.
- The bucket has a lifecycle configuration, which expires.
- The last time period in the ILM rule for the object ends, and there are no further placements specified.

When StorageGRID deletes an object because of a Delete objects in bucket operation, bucket lifecycle expiration, or ILM placement expiration, the replicated object is never deleted from the other grid in a grid federation connection. However, delete markers added to the source bucket by S3 client deletes can optionally be replicated to the destination bucket.

To understand what happens when an S3 client deletes objects from a bucket that has cross-grid replication enabled, review how S3 clients delete objects from buckets that have versioning enabled, as follows:

- If an S3 client issues a delete request that includes a version ID, that version of the object is permanently removed. No delete marker is added to the bucket.
- If an S3 client issues a delete request that does not include a version ID, StorageGRID does not delete any object versions. Instead, it adds a delete marker to the bucket. The delete marker causes StorageGRID to act as if the object was deleted:
 - A GetObject request without a version ID will fail with 404 No Object Found
 - A GetObject request with a valid version ID will succeed and return the requested object version.

When an S3 client deletes an object from a bucket that has cross-grid replication enabled, StorageGRID determines whether to replicate the delete request to the destination, as follows:

- If the delete request includes a version ID, that object version is permanently removed from the source grid. However, StorageGRID does not replicate delete requests that include a version ID, so the same object version is not deleted from the destination.
- If the delete request does not include a version ID, StorageGRID can optionally replicate the delete marker, based on how cross-grid replication is configured for the bucket:
 - If you choose to replicate delete markers (default), a delete marker is added to the source bucket and replicated to the destination bucket. In effect, the object appears to be deleted on both grids.
 - If you choose not to replicate delete markers, a delete marker is added to the source bucket but is not replicated to the destination bucket. In effect, objects that are deleted on the source grid aren't deleted on the destination grid.

In the figure, **Replicate delete markers** was set to **Yes** when cross-grid replication was enabled. Delete requests for the source bucket that include a version ID will not delete objects from the destination bucket. Delete requests for the source bucket that don't include a version ID will appear to delete objects in the destination bucket.



 (\mathbf{i})

If you want to keep object deletions synchronized between grids, create corresponding S3 lifecycle configurations for the buckets on both grids.

How encrypted objects are replicated

When you use cross-grid replication to replicate objects between grids, you can encrypt individual objects, use default bucket encryption, or configure grid-wide encryption. You can add, modify, or remove default bucket or

grid-wide encryption settings before or after you enable cross-grid replication for a bucket.

To encrypt individual objects, you can use SSE (server-side encryption with StorageGRID-managed keys) when adding the objects to the source bucket. Use the x-amz-server-side-encryption request header and specify AES256. See Use server-side encryption.



Using SSE-C (server-side encryption with customer-provided keys) is not supported for crossgrid replication. The ingest operation will fail.

To use default encryption for a bucket, use a PutBucketEncryption request and set the SSEAlgorithm parameter to AES256. Bucket-level encryption applies to any objects ingested without the x-amz-server-side-encryption request header. See Operations on buckets.

To use grid-level encryption, set the **Stored object encryption** option to **AES-256**. Grid-level encryption applies to any objects that aren't encrypted at the bucket level or that are ingested without the x-amz-server-side-encryption request header. See Configure network and object options.



SSE does not support AES-128. If the **Stored object encryption** option is enabled for the source grid using the **AES-128** option, the use of the AES-128 algorithm will not be propagated to the replicated object. Instead, the replicated object will use the destination's default bucket or grid-level encryption setting, if available.

When determining how to encrypt source objects, StorageGRID applies these rules:

- 1. Use the x-amz-server-side-encryption ingest header, if present.
- 2. If an ingest header is not present, use the bucket default encryption setting, if configured.
- 3. If a bucket setting is not configured, use the grid-wide encryption setting, if configured.
- 4. If a grid-wide setting is not present, don't encrypt the source object.

When determining how to encrypt replicated objects, StorageGRID applies these rules in this order:

- 1. Use the same encryption as the source object, unless that object uses AES-128 encryption.
- 2. If the source object is not encrypted or it uses AES-128, use the destination bucket's default encryption setting, if configured.
- 3. If the destination bucket does not have an encryption setting, use the destination's grid-wide encryption setting, if configured.
- 4. If a grid-wide setting is not present, don't encrypt the destination object.

PutObjectTagging and DeleteObjectTagging aren't supported

PutObjectTagging and DeleteObjectTagging requests aren't supported for objects in buckets that have crossgrid replication enabled.

If an S3 client issues a PutObjectTagging or DeleteObjectTagging request, 501 Not Implemented is returned. The message is Put(Delete) ObjectTagging is not available for buckets that have cross-grid replication configured.

How segmented objects are replicated

The source grid's maximum segment size applies to objects replicated to the destination grid. When objects

are replicated to another grid, the **Maximum Segment Size** setting (**CONFIGURATION** > **System** > **Storage options**) of the source grid will be used on both grids. For example, suppose the maximum segment size for the source grid is 1 GB, while the maximum segment size of the destination grid is 50 MB. If you ingest a 2-GB object on the source grid, that object is saved as two 1-GB segments. It will also be replicated to the destination grid as two 1-GB segments, even though that grid's maximum segment size is 50 MB.

Compare cross-grid replication and CloudMirror replication

As you begin using grid federation, review the similarities and differences between crossgrid replication and the StorageGRID CloudMirror replication service.

	Cross-grid replication	CloudMirror replication service
What is the primary purpose?	One StorageGRID system acts as a disaster recovery system. Objects in a bucket can be replicated between the grids in one or both directions.	Enables a tenant to automatically replicate objects from a bucket in StorageGRID (source) to an external S3 bucket (destination). CloudMirror replication creates an independent copy of an object in an independent S3 infrastructure. This independent copy is not used as a backup, but often further processed in the cloud.
How is it set up?	 Configure a grid federation connection between two grids. Add new tenant accounts, which are automatically cloned to the other grid. Add new tenant groups and users, which are also cloned. Create corresponding buckets on each grid and enable cross-grid replication to occur in one or both directions. 	 A tenant user configures CloudMirror replication by defining a CloudMirror endpoint (IP address, credentials, and so on) using the Tenant Manager or the S3 API. Any bucket owned by that tenant account can be configured to point to the CloudMirror endpoint.
Who is responsible for setting it up?	 A grid admin configures the connection and the tenants. Tenant users configure the groups, users, keys, and buckets. 	Typically, a tenant user.
What is the destination?	A corresponding and identical S3 bucket on the other StorageGRID system in the grid federation connection.	 Any compatible S3 infrastructure (including Amazon S3). Google Cloud Platform (GCP)
Is object versioning required?	Yes, both the source and destination buckets must have object versioning enabled.	No, CloudMirror replication supports any combination of unversioned and versioned buckets on both the source and destination.

	Cross-grid replication	CloudMirror replication service
What causes objects to be moved to the destination?	Objects are automatically replicated when they are added to a bucket that has cross-grid replication enabled.	Objects are automatically replicated when they are added to a bucket that has been configured with a CloudMirror endpoint. Objects that existed in the source bucket before the bucket was configured with the CloudMirror endpoint aren't replicated, unless they are modified.
How are objects replicated?	Cross-grid replication creates versioned objects, and it replicates the version ID from the source bucket to the destination bucket. This allows the version order to be maintained across both grids.	CloudMirror replication doesn't require versioning-enabled buckets, so CloudMirror can only maintain ordering for a key within a site. There are no guarantees that ordering will be maintained for requests to an object at different site.
What if an object can't be replicated?	The object is queued for replication, subject to metadata storage limits.	The object is queued for replication, subject to platform services limits (see Recommendations for using platform services).
Is the object's system metadata replicated?	Yes, when an object is replicated to the other grid, its system metadata is also replicated. The metadata will be identical on both grids.	No, when an object is replicated to the external bucket, its system metadata is updated. The metadata will differ between locations, depending on time of ingest and the behavior of the independent S3 infrastructure.
How are objects retrieved?	Applications can retrieve or read objects by making a request to the bucket on either grid.	Applications can retrieve or read objects by making a request either to StorageGRID or to the S3 destination. For example, suppose you use CloudMirror replication to mirror objects to a partner organization. The partner can use its own applications to read or update objects directly from the S3 destination. Using StorageGRID is not required.

	Cross-grid replication	CloudMirror replication service	
What happens if an object is deleted?	 Cross-grid replication Delete requests that include a version ID are never replicated to the destination grid. Delete requests that don't include a version ID add a delete marker to the source bucket, which can optionally be replicated to the destination grid. If cross-grid replication is configured for only one direction, objects in the destination bucket can be deleted without affecting the source. 	 Cloud Mirror replication service The results will vary based on the versioning state of the source and destination buckets (which don't need to be the same): If both buckets are versioned, a delete request will add a delete marker in both locations. If only the source bucket is versioned, a delete marker to the source bucket is versioned. 	
		 If neither bucket is versioned, a delete request will delete the object from the source but not from the destination. Similarly, objects in the destination bucket can be deleted without affecting the source. 	

Create grid federation connections

You can create a grid federation connection between two StorageGRID systems if you want to clone tenant details and replicate object data.

As shown in the figure, creating a grid federation connection includes steps on both grids. You add the connection on one grid and complete it on the other grid. You can start from either grid.



Before you begin

- You have reviewed the considerations and requirements for configuring grid federation connections.
- If you plan to use fully qualified domain names (FQDNs) for each grid instead of IP or VIP addresses, you know which names to use and you have confirmed that the DNS server for each grid has the appropriate entries.
- You are using a supported web browser.
- You have Root access permission and the provisioning passphrase for both grids.

Add connection

Perform these steps on either of the two StorageGRID systems.

Steps

- 1. Sign in to the Grid Manager from the primary Admin Node on either grid.
- 2. Select CONFIGURATION > System > Grid federation.
- 3. Select Add connection.
- 4. Enter details for the connection.

Field	Description
Connection name	A unique name to help you recognize this connection, for example, "Grid 1-Grid 2."
FQDN or IP for this grid	 One of the following: The FQDN of the grid you are currently signed into A VIP address of an HA group on this grid An IP address of an Admin Node or Gateway Node on this grid. The IP can be on any network that the destination grid can reach.
Port	The port you want to use for this connection. You can enter any unused port number from 23000 to 23999. Both grids in this connection will use the same port. You must ensure that no node in either grid uses this port for other connections.
Certificate valid days for this grid	The number of days you want the security certificates for this grid in the connection to be valid. The default value is 730 days (2 years), but you can enter any value from 1 to 762 days. StorageGRID automatically generates client and server certificates for each grid when you save the connection.
Provisioning passphrase for this grid	The provisioning passphrase for the grid you are signed in to.
FQDN or IP for the other grid	 One of the following: The FQDN of the grid you want to connect to A VIP address of an HA group on the other grid An IP address of an Admin Node or Gateway Node on the other grid. The IP can be on any network that the source grid can reach.

5. Select Save and continue.

6. For the Download verification file step, select **Download verification file**.

After the connection is completed on the other grid, you can no longer download the verification file from either grid.

7. Locate the downloaded file (connection-name.grid-federation), and save it to a safe location.



This file contains secrets (masked as *) and other sensitive details and must be securely stored and transmitted.

- 8. Select **Close** to return to the Grid federation page.
- 9. Confirm that the new connection is shown and that its Connection status is Waiting to connect.
- 10. Provide the *connection-name*.grid-federation file to the grid admin for the other grid.

Complete connection

Perform these steps on the StorageGRID system you are connecting to (the other grid).

Steps

- 1. Sign in to the Grid Manager from the primary Admin Node.
- 2. Select CONFIGURATION > System > Grid federation.
- 3. Select Upload verification file to access the Upload page.
- 4. Select **Upload verification file**. Then, browse to and select the file that was downloaded from the first grid (*connection-name*.grid-federation).

The details for the connection are shown.

 Optionally, enter a different number of valid days for the security certificates for this grid. The Certificate valid days entry defaults to the value you entered on the first grid, but each grid can use different expiration dates.

In general, use the same number of days for the certificates on both sides of the connection.



If the certificates on either end of the connection expire, the connection will stop working and replications will be pending until the certificates are updated.

- 6. Enter the provisioning passphrase for the grid you are currently signed in to.
- 7. Select Save and test.

The certificates are generated and the connection is tested. If the connection is valid, a success message appears and the new connection is listed on the Grid federation page. The **Connection status** will be **Connected**.

If an error message appears, address any issues. See Troubleshoot grid federation errors.

- 8. Go to the Grid federation page on the first grid and refresh the browser. Confirm that the **Connection status** is now **Connected**.
- 9. After the connection has been established, securely delete all copies of the verification file.

If you edit this connection, a new verification file will be created. The original file can't be reused.

After you finish

- Review the considerations for managing permitted tenants.
- Create one or more new tenant accounts, assign the **Use grid federation connection** permission, and select the new connection.
- Manage the connection as required. You can edit connection values, test a connection, rotate connection certificates, or remove a connection.
- Monitor the connection as part of your normal StorageGRID monitoring activities.
- Troubleshoot the connection, including resolving any alerts and errors related to account clone and crossgrid replication.

Manage grid federation connections

Managing grid federation connections between StorageGRID systems includes editing connection details, rotating the certificates, removing tenant permissions, and removing unused connections.

Before you begin

- You are signed in to the Grid Manager on either grid using a supported web browser.
- You have the Root access permission for the grid you are signed in to.

Edit a grid federation connection

You can edit a grid federation connection by signing in to the primary Admin Node on either grid in the connection. After you make changes to the first grid, you must download a new verification file and upload it to the other grid.



While the connection is being edited, account clone or cross-grid replication requests will continue to use the existing connection settings. Any edits you make to the first grid are saved locally but aren't used until they have been uploaded to the second grid, saved, and tested.

Start editing the connection

Steps

- 1. Sign in to the Grid Manager from the primary Admin Node on either grid.
- 2. Select NODES and confirm that all other Admin Nodes in your system are online.



When you edit a grid federation connection, StorageGRID attempts to save a "candidate configuration" file on all Admin Nodes on the first grid. If this file can't be saved to all Admin Nodes, a warning message appears when you select **Save and test**.

- 3. Select **CONFIGURATION > System > Grid federation**.
- 4. Edit the connection details using the **Actions** menu on the Grid federation page or the details page for a specific connection. See Create grid federation connections for what to enter.

Actions menu

- a. Select the radio button for the connection.
- b. Select Actions > Edit.
- c. Enter the new information.

Details page

- a. Select a connection name to display its details.
- b. Select Edit.
- c. Enter the new information.
- 5. Enter the provisioning passphrase for the grid you are signed in to.

6. Select Save and continue.

The new values are saved, but they will not be applied to the connection until you have uploaded the new verification file on the other grid.

7. Select Download verification file.

To download this file at a later time, go to the details page for the connection.

8. Locate the downloaded file (connection-name.grid-federation), and save it to a safe location.



The verification file contains secrets and must be securely stored and transmitted.

- 9. Select Close to return to the Grid federation page.
- 10. Confirm that the Connection status is Pending edit.



If the connection status was something other than **Connected** when you started editing the connection, it will not change to **Pending edit**.

11. Provide the *connection-name*.grid-federation file to the grid admin for the other grid.

Finish editing the connection

Finish editing the connection by uploading the verification file on the other grid.

Steps

- 1. Sign in to the Grid Manager from the primary Admin Node.
- 2. Select CONFIGURATION > System > Grid federation.
- 3. Select Upload verification file to access the upload page.
- 4. Select Upload verification file. Then, browse to and select the file that was downloaded from the first grid.
- 5. Enter the provisioning passphrase for the grid you are currently signed in to.
- 6. Select Save and test.

If the connection can be established using the edited values, a success message appears. Otherwise, an error message appears. Review the message and address any issues.

- 7. Close the wizard to return to the Grid federation page.
- 8. Confirm that the **Connection status** is **Connected**.
- 9. Go to the Grid federation page on the first grid and refresh the browser. Confirm that the **Connection status** is now **Connected**.
- 10. After the connection has been established, securely delete all copies of the verification file.

Test a grid federation connection

Steps

- 1. Sign in to the Grid Manager from the primary Admin Node.
- 2. Select CONFIGURATION > System > Grid federation.
- 3. Test the connection using the **Actions** menu on the Grid federation page or the details page for a specific connection.

Actions menu

a. Select the radio button for the connection.

b. Select Actions > Test.

Details page

- a. Select a connection name to display its details.
- b. Select Test connection.
- 4. Review the connection status:

Connection status	Description
Connected	Both grids are connected and communicating normally.
Error	The connection is in an error state. For example, a certificate has expired or a configuration value is no longer valid.
Pending edit	You have edited the connection on this grid, but the connection is still using the existing configuration. To complete the edit, upload the new verification file to the other grid.
Waiting to connect	You have configured the connection on this grid, but the connection hasn't been completed on the other grid. Download the verification file from this grid and upload it to the other grid.
Unknown	The connection is in an unknown state, possibly because of a networking issue or an offline node.

5. If the Connection status is **Error**, resolve any issues. Then, select **Test connection** again to confirm the issue has been fixed.

Rotate connection certificates

Each grid federation connection uses four automatically-generated SSL certificates to secure the connection. When the two certificates for each grid near their expiration date, the **Expiration of grid federation certificate** alert reminds you to rotate the certificates.



If the certificates on either end of the connection expire, the connection will stop working and replications will be pending until the certificates are updated.

Steps

1. Sign in to the Grid Manager from the primary Admin Node on either grid.

2. Select CONFIGURATION > System > Grid federation.

- 3. From either tab on the Grid federation page, select the connection name to display its details.
- 4. Select the Certificates tab.
- 5. Select Rotate certificates.
- 6. Specify how many days the new certificates should be valid.
- 7. Enter the provisioning passphrase for the grid you are signed in to.
- 8. Select Rotate certificates.
- 9. As required, repeat these steps on the other grid in the connection.

In general, use the same number of days for the certificates on both sides of the connection.

Remove a grid federation connection

You can remove a grid federation connection from either grid in the connection. As shown in the figure, you must perform prerequisite steps on both grids to confirm that the connection is not being used by any tenant on either grid.



Before removing a connection, note the following:

- Removing a connection does not delete any items that have already been copied between grids. For example, tenant users, groups, and objects that exist on both grids aren't deleted from either grid when the tenant's permission is removed. If you want to delete these items, you must manually delete them from both grids.
- When you remove a connection, any objects that are pending replication (ingested but not yet replicated to the other grid) will have their replication permanently failed.

Disable replication for all tenant buckets

Steps

- 1. Starting from either grid, sign in to the Grid Manager from the primary Admin Node.
- 2. Select CONFIGURATION > System > Grid federation.
- 3. Select the connection name to display its details.

- 4. On the **Permitted tenants** tab, determine if the connection is being used by any tenants.
- 5. If any tenants are listed, instruct all tenants to disable cross-grid replication for all of their buckets on both grids in the connection.



You can't remove the **Use grid federation connection** permission if any tenant buckets have cross-grid replication enabled. Each tenant account must disable cross-grid replication for their buckets on both grids.

Remove permission for each tenant

After cross-grid replication has been disabled for all tenant buckets, remove the **Use grid federation permission** from all tenants on both grids.

Steps

- 1. Select CONFIGURATION > System > Grid federation.
- 2. Select the connection name to display its details.
- 3. For each tenant on the **Permitted tenants** tab, remove the **Use grid federation connection** permission from each tenant. See Manage permitted tenants.
- 4. Repeat these steps for the permitted tenants on the other grid.

Remove connection

Steps

- 1. When no tenants on either grid are using the connection, select **Remove**.
- 2. Review the confirmation message, and select **Remove**.
 - If the connection can be removed, a success message is shown. The grid federation connection is now removed from both grids.
 - If the connection can't be removed (for example, it is still in use or there is a connection error), an error message is displayed. You can do either of the following:
 - Resolve the error (recommended). See Troubleshoot grid federation errors.
 - Remove the connection by force. See the next section.

Remove a grid federation connection by force

If necessary, you can force the removal of a connection that does not have **Connected** status.

Force removal only deletes the connection from the local grid. To completely remove the connection, perform the same steps on both grids.

Steps

1. From the confirmation dialog box, select **Force remove**.

A success message appears. This grid federation connection can no longer be used. However, tenant buckets might still have cross-grid replication enabled and some object copies might have already been replicated between the grids in the connection.

- 2. From the other grid in the connection, sign in to the Grid Manager from the primary Admin Node.
- 3. Select CONFIGURATION > System > Grid federation.

- 4. Select the connection name to display its details.
- 5. Select **Remove** and **Yes**.
- 6. Select Force remove to remove the connection from this grid.

Manage the permitted tenants for grid federation

You can allow S3 tenant accounts to use a grid federation connection between two StorageGRID systems. When tenants are allowed to use a connection, special steps are required to edit tenant details or to permanently remove a tenant's permission to use the connection.

Before you begin

- You are signed in to the Grid Manager on either grid using a supported web browser.
- · You have the Root access permission for the grid you are signed in to.
- You have created a grid federation connection between two grids.
- You have reviewed the workflows for account clone and cross-grid replication.
- As required, you have already configured single sign-on (SSO) or identify federation for both grids in the connection. See What is account clone.

Create a permitted tenant

If you want to allow a new or existing tenant account to use a grid federation connection for account clone and cross-grid replication, follow the general instructions to create a new S3 tenant or edit a tenant account and note the following:

- You can create the tenant from either grid in the connection. The grid where a tenant is created is the *tenant's source grid*.
- The status of the connection must be Connected.
- When the tenant is created or edited to enable the **Use grid federation connection** permission and then saved on the first grid, an identical tenant is automatically replicated to the other grid. The grid where the tenant is replicated is the *tenant's destination grid*.
- The tenants on both grids will have the same 20-digit account ID, name, description, quota, and permissions. Optionally, you can use the **Description** field to help identify which is the source tenant and which is the destination tenant. For example, this description for a tenant created on Grid 1 will also appear for the tenant replicated to Grid 2: "This tenant was created on Grid 1."
- For security reasons, the password for a local root user is not copied to the destination grid.



Before a local root user can sign in to the replicated tenant on the destination grid, a grid administrator for that grid must change the password for the local root user.

- After the new or edited tenant is available on both grids, tenant users can perform these operations:
 - From the tenant's source grid, create groups and local users, which are automatically cloned to the tenant's destination grid. See Clone tenant groups and users.
 - Create new S3 access keys, which can be optionally cloned to the tenant's destination grid. See Clone S3 access keys using the API.
 - Create identical buckets on both grids in the connection and enable cross-grid replication in one direction or in both directions. See Manage cross-grid replication.

View a permitted tenant

You can see details for a tenant that is permitted to use a grid federation connection.

Steps

- 1. Select TENANTS.
- 2. From the Tenants page, select the tenant name to view the tenant details page.

If this is the source grid for the tenant (that is, if the tenant was created on this grid), a banner appears to remind you that the tenant was cloned to another grid. If you edit or delete this tenant, your changes will not be synced to the other grid.

Tenants > tenant A for grid federation		
tenant A for grid federation		
Tenant ID: 0899 6970 1700 0930 0009	Quota — utilization:	
Object 0 count:	Logical 0 bytes space used:	
Description: this tenant was created on Grid 1	Quota: —	
Sign in Edit Actions ~		
This tenant has been cloned to another grid. If you edit o	or delete this tenant, your changes will not t	be synced to the other grid.
Space breakdown Allowed features	Grid federation	
Remove permission Clear error Search	Q	Displaying one result
Connection name 💠 Connection status	😧 💠 Remote grid hostname 🥥	Last error 2 \$
○ Grid 1 to Grid 2 Sconnected	10.96.106.230	Check for errors

3. Optionally select the Grid federation tab to monitor the grid federation connection.

Edit a permitted tenant

If you need to edit a tenant that has the **Use grid federation connection** permission, follow the general instructions for editing a tenant account and note the following:

• If a tenant has the **Use grid federation connection** permission, you can edit tenant details from either grid in the connection. However, any changes you make will not be copied to the other grid. If you want to keep the tenant details synchronized between grids, you must make the same edits on both grids.

- You can't clear the **Use grid federation connection** permission when you are editing a tenant.
- You can't select a different grid federation connection when you are editing a tenant.

Delete a permitted tenant

If you need to remove a tenant that has the **Use grid federation connection** permission, follow the general instructions for deleting a tenant account and note the following:

- Before you can remove the original tenant on the source grid, you must remove all buckets for the account on the source grid.
- Before you can remove the cloned tenant on the destination grid, you must remove all buckets for the account on the destination grid.
- If you remove either the original or the cloned tenant, the account can no longer be used for cross-grid replication.
- If you are removing the original tenant on the source grid, any tenant groups, users, or keys that were cloned to the destination grid will be unaffected. You can either delete the cloned tenant or allow it to manage its own groups, users, access keys, and buckets.
- If you are removing the cloned tenant on the destination grid, clone errors will occur if new groups or users are added to the original tenant.

To avoid these errors, remove the tenant's permission to use the grid federation connection before deleting the tenant from this grid.

Remove Use grid federation connection permission

To prevent a tenant from using a grid federation connection, you must remove the **Use grid federation** connection permission.



Before removing a tenant's permission to use a grid federation connection, note the following:

- You can't remove the **Use grid federation connection** permission if any of the tenant's buckets have cross-grid replication enabled. The tenant account must disable cross-grid replication for all of their buckets first.
- Removing the **Use grid federation connection** permission does not delete any items that have already been replicated between grids. For example, any tenant users, groups, and objects that exist on both grids aren't deleted from either grid when the tenant's permission is removed. If you want to delete these items, you must manually delete them from both grids.
- If you want to re-enable this permission with the same grid federation connection, delete this tenant on the destination grid first; otherwise, re-enabling this permission will result in an error.



Re-enabling the **Use grid federation connection** permission makes the local grid the source grid and triggers cloning to the remote grid specified by the selected grid federation connection. If the tenant account already exists on the remote grid, cloning will result in a conflict error.

Before you begin

- You are using a supported web browser.
- You have the Root access permission for both grids.

Disable replication for tenant buckets

As a first step, disable cross-grid replication for all tenant buckets.

Steps

- 1. Starting from either grid, sign in to the Grid Manager from the primary Admin Node.
- 2. Select CONFIGURATION > System > Grid federation.
- 3. Select the connection name to display its details.
- 4. On the **Permitted tenants** tab, determine if the tenant is using the connection.
- 5. If the tenant is listed, instruct them to disable cross-grid replication for all of their buckets on both grids in the connection.



You can't remove the **Use grid federation connection** permission if any tenant buckets have cross-grid replication enabled. The tenant must disable cross-grid replication for their buckets on both grids.

Remove permission for tenant

After cross-grid replication is disabled for tenant buckets, you can remove the tenant's permission to use the grid federation connection.

Steps

- 1. Sign in to the Grid Manager from the primary Admin Node.
- 2. Remove the permission from the Grid federation page or the Tenants page.

Grid federation page

- a. Select CONFIGURATION > System > Grid federation.
- b. Select the connection name to display its details page.
- c. On the Permitted tenants tab, select radio button for the tenant.
- d. Select Remove permission.

Tenants page

- a. Select TENANTS.
- b. Select the tenant's name to display the details page.
- c. On the Grid federation tab, select radio button for the connection.
- d. Select Remove permission.
- 3. Review the warnings in the confirmation dialog box, and select **Remove**.
 - If the permission can be removed, you are returned to the details page and a success message is shown. This tenant can no longer use the grid federation connection.

• If one or more tenant buckets still have cross-grid replication enabled, an error is displayed.



You can do either of the following:

- (Recommended.) Sign in to the Tenant Manager and disable replication for each of the tenant's buckets. See Manage cross-grid replication. Then, repeat the steps to remove the Use grid connection permission.
- Remove the permission by force. See the next section.

4. Go to the other grid and repeat these steps to remove the permission for the same tenant on the other grid.

Remove the permission by force

If necessary, you can force the removal of a tenant's permission to use a grid federation connection even if tenant buckets have cross-grid replication enabled.

Before removing a tenant's permission by force, note the general considerations for removing the permission as well as these additional considerations:

• If you remove the **Use grid federation connection** permission by force, any objects that are pending replication to the other grid (ingested but not yet replicated) will continue to be replicated. To prevent these

in-process objects from reaching the destination bucket, you must remove the tenant's permission on the other grid as well.

• Any objects ingested into the source bucket after you remove the **Use grid federation connection** permission will never be replicated to the destination bucket.

Steps

- 1. Sign in to the Grid Manager from the primary Admin Node.
- 2. Select CONFIGURATION > System > Grid federation.
- 3. Select the connection name to display its details page.
- 4. On the **Permitted tenants** tab, select radio button for the tenant.
- 5. Select Remove permission.
- 6. Review the warnings in the confirmation dialog box, and select **Force remove**.

A success message appears. This tenant can no longer use the grid federation connection.

7. As required, go to the other grid and repeat these steps to force-remove the permission for the same tenant account on the other grid. For example, you should repeat these steps on the other grid to prevent in-process objects from reaching the destination bucket.

Troubleshoot grid federation errors

You might need to troubleshoot alerts and errors related to grid federation connections, account clone, and cross-grid replication.

Grid federation connection alerts and errors

You might receive alerts or experience errors with your grid federation connections.

After making any changes to resolve a connection issue, test the connection to ensure that the connection status returns to **Connected**. For instructions, see Manage grid federation connections.

Grid federation connection failure alert

Issue

The Grid federation connection failure alert was triggered.

Details

This alert indicates that the grid federation connection between the grids is not working.

Recommended actions

- 1. Review the settings on the Grid Federation page for both grids. Confirm that all values are correct. See Manage grid federation connections.
- 2. Review the certificates used for the connection. Make sure there are no alerts for expired grid federation certificates and that the details for each certificate are valid. See the instructions for rotating connection certificates in Manage grid federation connections.
- 3. Confirm that all Admin and Gateway Nodes in both grids are online and available. Resolve any alerts that might be affecting these nodes and try again.
- 4. If you provided a fully qualified domain name (FQDN) for the local or remote grid, confirm the DNS server is online and available. See What is grid federation? for networking, IP address, and DNS requirements.

Expiration of grid federation certificate alert

Issue

The Expiration of grid federation certificate alert was triggered.

Details

This alert indicates that one or more grid federation certificates are about to expire.

Recommended actions

See the instructions for rotating connection certificates in Manage grid federation connections.

Error editing a grid federation connection

Issue

When editing a grid federation connection, you see the following warning message when you select **Save and test**: "Failed to create a candidate configuration file on one or more nodes."

Details

When you edit a grid federation connection, StorageGRID attempts to save a "candidate configuration" file on all Admin Nodes on the first grid. A warning message appears if this file can't be saved to all Admin Nodes, for example, because an Admin Node is offline.

Recommended actions

- 1. From the grid you are using to edit the connection, select **NODES**.
- 2. Confirm that all Admin Nodes for that grid are online.
- 3. If any nodes are offline, bring them back online and try editing the connection again.

Account clone errors

Can't sign in to a cloned tenant account

Issue

You can't sign in to a cloned tenant account. The error message on the Tenant Manager sign-in page is "Your credentials for this account were invalid. Please try again."

Details

For security reasons, when a tenant account is cloned from the tenant's source grid to the tenant's destination grid, the password you set for the tenant's local root user is not cloned. Similarly, when a tenant creates local users on its source grid, the local user passwords aren't cloned to the destination grid.

Recommended actions

Before the root user can sign in to the tenant's destination grid, a grid administrator must first change the password for the local root user on the destination grid.

Before a cloned local user can sign in to the tenant's destination grid, the root user for the cloned tenant must add a password for the user on the destination grid. For instructions, see Manage local users in the instructions for using the Tenant Manager.

Tenant created without a clone

Issue

You see the message "Tenant created without a clone" after creating a new tenant with the Use grid

federation connection permission.

Details

This issue can occur if updates to the Connection status are delayed, which might cause an unhealthy connection to be listed as **Connected**.

Recommended actions

- 1. Review the reason listed in the error message and resolve any networking or other issues that might be preventing the connection from working. See Grid federation connection alerts and errors.
- 2. Follow the instructions to test a grid federation connection in Manage grid federation connections to confirm the issue has been fixed.
- 3. From the tenant's source grid, select **TENANTS**.
- 4. Locate the tenant account that failed to be cloned.
- 5. Select the tenant name to display the details page.

6. Select Retry account clone.

Tenants > test			
test			
Tenant ID: Protocol:	0040 2213 8117 4859 6503 👘	Quota utilization:	-
Object count:	0	Logical space used:	0 bytes
		Quota:	-
Sign in E	idit Actions 🗸		
Tenant according to the second s	ount could not be cloned to the other grid. ernal server error. The server encountered an error a or count clone	nd could not complete your reques	t. Try again. If the problem persists, contact support. Internal

If the error has been resolved, the tenant account will now be cloned to the other grid.

Cross-grid replication alerts and errors

Last error shown for connection or tenant

Issue

When viewing a grid federation connection (or when managing the permitted tenants for a connection), you notice an error in the **Last error** column on the connection details page. For example:

Grid 1 - Grid 2			
Local hostname (this grid):	10.96.130.64		
Port:	23000		
Remote hostname (other grid):	10.96.130.76		
Connection status:	S Connected		
Edit Download file Test conn	ection Remove		
Permitted tenants Cer	tificates		
Remove permission Clear error	Search	Q	isplaying one result
Tenant name 🗘 Last error 🧿	¢		
O Tenant A Object Lock Check for en	6:19:20 MST plication has encountered an error. Failed to send cross bucket' to destination bucket 'my-bucket'. Error code: De tState. Confirm that the source and destination buckets disabled. (logID 13916508109026943924) ors	-grid replication request from estinationRequestError. Detail s have object versioning enabl	i source I: led and S3

Details

For each grid federation connection, the **Last error** column shows the most recent error to occur, if any, when a tenant's data was being replicated to the other grid. This column only shows the last cross-grid replication error to occur; previous errors that might have occurred will not be shown. An error in this column might occur for one of these reasons:

- The source object version was not found.
- The source bucket was not found.
- The destination bucket was deleted.
- The destination bucket was re-created by a different account.
- The destination bucket has versioning suspended.
- The destination bucket was re-created by the same account but is now unversioned.

Recommended actions

If an error message appears in the Last error column, follow these steps:

- 1. Review the message text.
- 2. Perform any recommended actions. For example, if versioning was suspended on the destination bucket for cross-grid replication, reenable versioning for that bucket.
- 3. Select the connection or tenant account from the table.
- 4. Select Clear error.
- 5. Select **Yes** to clear the message and update the system's status.
6. Wait 5-6 minutes and then ingest a new object into the bucket. Confirm that the error message does not reappear.



To ensure the error message is cleared, wait at least 5 minutes after the timestamp in the message before ingesting a new object.



After you clear the error, a new **Last error** might appear if objects are ingested in a different bucket that also has an error.

7. To determine if any objects failed to be replicated because of the bucket error, see Identify and retry failed replication operations.

Cross-grid replication permanent failure alert

Issue

The Cross-grid replication permanent failure alert was triggered.

Details

This alert indicates that tenant objects can't be replicated between the buckets on two grids for a reason that requires user intervention to resolve. This alert is typically caused by a change to either the source or the destination bucket.

Recommended actions

- 1. Sign in to the grid where the alert was triggered.
- 2. Go to CONFIGURATION > System > Grid federation, and locate the connection name listed in the alert.
- 3. On the Permitted tenants tab, look at the **Last error** column to determine which tenant accounts have errors.
- 4. To learn more about the failure, see the instructions in Monitor grid federation connections to review the cross-grid replication metrics.
- 5. For each affected tenant account:
 - a. See the instructions in Monitor tenant activity to confirm that the tenant has not exceeded its quota on the destination grid for cross-grid replication.
 - b. As required, increase the tenant's quota on the destination grid to allow new objects to be saved.
- 6. For each affected tenant, sign in to Tenant Manager on both grids, so you can compare the list of buckets.
- 7. For each bucket that has cross-grid replication enabled, confirm the following:
 - There is a corresponding bucket for the same tenant on the other grid (must use the exact name).
 - · Both buckets have object versioning enabled (versioning can't be suspended on either grid).
 - Both buckets have S3 Object Lock disabled.
 - Neither bucket is in the Deleting objects: read-only state.
- 8. To confirm that the issue was resolved, see the instructions in Monitor grid federation connections to review the cross-grid replication metrics, or perform these steps:
 - a. Go back to the Grid federation page.
 - b. Select the affected tenant, and select Clear Error in the Last error column.
 - c. Select Yes to clear the message and update the system's status.

d. Wait 5-6 minutes and then ingest a new object into the bucket. Confirm that the error message does not reappear.



To ensure the error message is cleared, wait at least 5 minutes after the timestamp in the message before ingesting a new object.

It might take up to a day for the alert to clear after it is resolved.

e. Go to Identify and retry failed replication operations to identify any objects or delete markers that failed to be replicated to the other grid and to retry replication as needed.

Cross-grid replication resource unavailable alert

Issue

The Cross-grid replication resource unavailable alert was triggered.

Details

This alert indicates that cross-grid replication requests are pending because a resource is unavailable. For example, there might be a network error.

Recommended actions

- 1. Monitor the alert to see if the issue resolves on its own.
- If the issue persists, determine if either grid has a Grid federation connection failure alert for the same connection or an Unable to communicate with node alert for a node. This alert might be resolved when you resolve those alerts.
- 3. To learn more about the failure, see the instructions in Monitor grid federation connections to review the cross-grid replication metrics.
- 4. If you can't resolve the alert, contact technical support.

Cross-grid replication will proceed as normal after the issue is resolved.

Identify and retry failed replication operations

After resolving the **Cross-grid replication permanent failure** alert, you should determine if any objects or delete markers failed to be replicated to the other grid. You can then reingest these objects or use the Grid Management API to retry replication.

The **Cross-grid replication permanent failure** alert indicates that tenant objects can't be replicated between the buckets on two grids for a reason that requires user intervention to resolve. This alert is typically caused by a change to either the source or the destination bucket. For details, see Troubleshoot grid federation errors.

Determine if any objects failed to be replicated

To determine if any objects or delete markers have not been replicated to the other grid, you can search the audit log for CGRR (Cross-Grid Replication Request) messages. This message is added to the log when StorageGRID fails to replicate an object, multipart object, or delete marker to the destination bucket.

You can use the audit-explain tool to translate the results into an easier-to-read format.

Before you begin

- You have Root access permission.
- You have the Passwords.txt file.
- You know the IP address of the primary Admin Node.

Steps

- 1. Log in to the primary Admin Node:
 - a. Enter the following command: ssh admin@primary_Admin_Node_IP
 - b. Enter the password listed in the Passwords.txt file.
 - c. Enter the following command to switch to root: su -
 - d. Enter the password listed in the <code>Passwords.txt</code> file.

When you are logged in as root, the prompt changes from \$ to #.

2. Search the audit.log for CGRR messages, and use the audit-explain tool to format the results.

For example, this command greps for all CGRR messages in the past 30 minutes and uses the auditexplain tool.

awk -vdate=\$(date -d "30 minutes ago" '+%Y-%m-%dT%H:%M:%S') '\$1\$2 >= date {
print }' audit.log | grep CGRR | audit-explain

The results of the command will look like this example, which has entries for six CGRR messages. In the example, all cross-grid replication requests returned a general error because the object could not be replicated. The first three errors are for "replicate object" operations, and the last three errors are for "replicate delete marker" operations.

```
CGRR Cross-Grid Replication Request tenant:50736445269627437748
connection:447896B6-6F9C-4FB2-95EA-AEBF93A774E9 operation:"replicate
object" bucket:bucket123 object:"audit-0"
version:QjRBNDIzODAtNjQ3My0xMUVELTq2QjEtODJBMjAwQkI3NEM4 error:qeneral
error
CGRR Cross-Grid Replication Request tenant: 50736445269627437748
connection:447896B6-6F9C-4FB2-95EA-AEBF93A774E9 operation:"replicate
object" bucket:bucket123 object:"audit-3"
version:QjRDOTRCOUMtNjQ3My0xMUVELTkzM0YtOTq1MTAwQkI3NEM4 error:general
error
CGRR Cross-Grid Replication Request tenant:50736445269627437748
connection:447896B6-6F9C-4FB2-95EA-AEBF93A774E9 operation:"replicate
delete marker" bucket:bucket123 object:"audit-1"
version:NUQ00EYxMDAtNjQ3NC0xMUVELTq2NjMtOTY5NzAwQkI3NEM4 error:general
error
CGRR Cross-Grid Replication Request tenant:50736445269627437748
connection:447896B6-6F9C-4FB2-95EA-AEBF93A774E9 operation:"replicate
delete marker" bucket:bucket123 object:"audit-5"
version:NUQ1ODUwQkUtNjQ3NC0xMUVELTg1NTItRDkwNzAwQkI3NEM4 error:general
error
```

Each entry contains the following information:

Field	Description
CGRR Cross-Grid Replication Request	The name of the request
tenant	The tenant's account ID
connection	The ID of the grid federation connection
operation	 The type of replication operation that was being attempted: replicate object replicate delete marker replicate multipart object
bucket	The bucket name
object	The object name
version	The version ID for the object

Field	Description
error	The type of error. If cross-grid replication failed, the error is "General error".

Retry failed replications

After generating a list of objects and delete markers that were not replicated to the destination bucket and resolving the underlying issues, you can retry replication in either of two ways:

- · Reingest each object into the source bucket.
- Use the Grid Management private API, as described.

Steps

- 1. From the top of the Grid Manager, select the help icon and select **API documentation**.
- 2. Select Go to private API documentation.



The StorageGRID API endpoints that are marked "Private" are subject to change without notice. StorageGRID private endpoints also ignore the API version of the request.

3. In the cross-grid-replication-advanced section, select the following endpoint:

POST /private/cross-grid-replication-retry-failed

- 4. Select Try it out.
- 5. In the **body** text box, replace the example entry for **versionID** with a version ID from the audit.log that corresponds to a failed cross-grid-replication request.

Be sure to retain the double quotes around the string.

- 6. Select Execute.
- 7. Confirm that the server response code is **204**, indicating that the object or delete marker has been marked as pending for cross-grid replication to the other grid.



Pending means the cross-grid replication request has been added to the internal queue for processing.

Monitor replication retries

You should monitor the replication retry operations to make sure they complete.



It might take several hours or longer for an object or delete marker to be replicated to the other grid.

You can monitor retry operations in either of two ways:

• Use an S3 HeadObject or GetObject request. The response includes the StorageGRID-specific x-ntapsg-cgr-replication-status response header, which will have one of the following values:

Grid	Replication status
Source	SUCCESS: The replication was successful.
	 PENDING: The object hasn't been replicated yet.
	 FAILURE: The replication failed with a permanent failure. A user must resolve the error.
Destination	REPLICA : The object was replicated from the source grid.

• Use the Grid Management private API, as described.

Steps

1. In the **cross-grid-replication-advanced** section of the private API documentation, select the following endpoint:

GET /private/cross-grid-replication-object-status/{id}

- 2. Select Try it out.
- 3. In the Parameter section, enter the version ID you used in the cross-grid-replication-retryfailed request.
- 4. Select **Execute**.
- 5. Confirm that the server response code is 200.
- 6. Review the replication status, which will be one of the following:
 - **PENDING**: The object hasn't been replicated yet.
 - COMPLETED: The replication was successful.
 - FAILED: The replication failed with a permanent failure. A user must resolve the error.

Manage security

Manage security: Overview

You can configure various security settings from the Grid Manager to help secure your StorageGRID system.

Manage encryption

StorageGRID provides several options for encrypting data. You should review the available encryption methods to determine which ones meet your data-protection requirements.

Manage certificates

You can configure and manage the server certificates used for HTTP connections or the client certificates used to authenticate a client or user identity to the server.

Configure key management servers

Using a key management server lets you protect StorageGRID data even if an appliance is removed from the data center. After the appliance volumes are encrypted, you can't access any data on the appliance unless the

node can communicate with the KMS.



To use encryption key management, you must enable the **Node Encryption** setting for each appliance during installation, before the appliance is added to the grid.

Manage proxy settings

If you are using S3 platform services or Cloud Storage Pools, you can configure a storage proxy server between Storage Nodes and the external S3 endpoints. If you send AutoSupport packages using HTTPS or HTTP, you can configure an admin proxy server between Admin Nodes and technical support.

Control firewalls

To enhance the security of your system, you can control access to StorageGRID Admin Nodes by opening or closing specific ports at the external firewall. You can also control network access to each node by configuring its internal firewall. You can prevent access on all ports except those needed for your deployment.

Review StorageGRID encryption methods

StorageGRID provides several options for encrypting data. You should review the available methods to determine which methods meet your data-protection requirements.

The table provides a high-level summary of the encryption methods available in StorageGRID.

Encryption option	How it works	Applies to
Key management server (KMS) in Grid Manager	You configure a key management server for the StorageGRID site and enable node encryption for the appliance. Then, an appliance node connects to the KMS to request a key encryption key (KEK). This key encrypts and decrypts the data encryption key (DEK) on each volume.	Appliance nodes that have Node Encryption enabled during installation. All data on the appliance is protected against physical loss or removal from the data center. Note : Managing encryption keys with a KMS is only supported for Storage Nodes and services appliances.
Drive Encryption page in StorageGRID Appliance Installer	If the appliance contains drives that support hardware encryption, you can set a drive passphrase during installation. When you set a drive passphrase, it's impossible for anyone to recover valid data from drives that have been removed from the system, unless they know the passphrase. Before starting installation, go to Configure Hardware > Drive Encryption to set a drive passphrase that applies to all StorageGRID-managed, self- encrypting drives in a node.	Appliances that contain self- encrypting drives. All data on the secured drives is protected against physical loss or removal from the data center. Drive encryption doesn't apply to SANtricity-managed drives. If you have a storage appliance with self- encrypting drives and SANtricity controllers, you can enable drive security in SANtricity.

Encryption option	How it works	Applies to
Drive security in SANtricity System Manager	If the Drive Security feature is enabled for an SG5700 or SG6000 storage appliance, you can use SANtricity System Manager to create and manage the security key. The key is required to access the data on the secured drives.	Storage appliances that have Full Disk Encryption (FDE) drives or self-encrypting drives. All data on the secured drives is protected against physical loss or removal from the data center. Can't be used with some storage appliances or with any services appliances.
Stored object encryption	You enable the Stored object encryption option in the Grid Manager. When enabled, any new objects that aren't encrypted at the bucket level or at the object level are encrypted during ingest.	Newly ingested S3 and Swift object data. Existing stored objects aren't encrypted. Object metadata and other sensitive data aren't encrypted.
S3 bucket encryption	You issue a PutBucketEncryption request to enable encryption for the bucket. Any new objects that aren't encrypted at the object level are encrypted during ingest.	Newly ingested S3 object data only. Encryption must be specified for the bucket. Existing bucket objects aren't encrypted. Object metadata and other sensitive data aren't encrypted. Operations on buckets
S3 object server-side encryption (SSE)	You issue an S3 request to store an object and include the x-amz- server-side-encryption request header.	Newly ingested S3 object data only. Encryption must be specified for the object. Object metadata and other sensitive data aren't encrypted. StorageGRID manages the keys. Use server-side encryption

Encryption option	How it works	Applies to
S3 object server-side encryption with customer-provided keys (SSE- C)	<pre>You issue an S3 request to store an object and include three request headers.</pre>	Newly ingested S3 object data only. Encryption must be specified for the object. Object metadata and other sensitive data aren't encrypted. Keys are managed outside of StorageGRID. Use server-side encryption
External volume or datastore encryption	You use an encryption method outside of StorageGRID to encrypt an entire volume or datastore, if your deployment platform supports it.	All object data, metadata, and system configuration data, assuming every volume or datastore is encrypted. An external encryption method provides tighter control over encryption algorithms and keys. Can be combined with the other methods listed.
Object encryption outside of StorageGRID	You use an encryption method outside of StorageGRID to encrypt object data and metadata before they are ingested into StorageGRID.	Object data and metadata only (system configuration data is not encrypted). An external encryption method provides tighter control over encryption algorithms and keys. Can be combined with the other methods listed. Amazon Simple Storage Service - Developer Guide: Protecting data using client-side encryption

Use multiple encryption methods

Depending on your requirements, you can use more than one encryption method at a time. For example:

- You can use a KMS to protect appliance nodes and also use the drive security feature in SANtricity System Manager to "double encrypt" data on the self-encrypting drives in the same appliances.
- You can use a KMS to secure data on appliance nodes and also use the Stored object encryption option to encrypt all objects when they are ingested.

If only a small portion of your objects require encryption, consider controlling encryption at the bucket or individual object level instead. Enabling multiple levels of encryption has an additional performance cost.

Manage certificates

Manage security certificates: Overview

Security certificates are small data files used to create secure, trusted connections between StorageGRID components and between StorageGRID components and external systems.

StorageGRID uses two types of security certificates:

- Server certificates are required when you use HTTPS connections. Server certificates are used to establish secure connections between clients and servers, authenticating the identity of a server to its clients and providing a secure communication path for data. The server and the client each have a copy of the certificate.
- **Client certificates** authenticate a client or user identity to the server, providing more secure authentication than passwords alone. Client certificates don't encrypt data.

When a client connects to the server using HTTPS, the server responds with the server certificate, which contains a public key. The client verifies this certificate by comparing the server signature to the signature on its copy of the certificate. If the signatures match, the client starts a session with the server using the same public key.

StorageGRID functions as the server for some connections (such as the load balancer endpoint) or as the client for other connections (such as the CloudMirror replication service).

Default Grid CA certificate

StorageGRID includes a built-in certificate authority (CA) that generates an internal Grid CA certificate during system installation. The Grid CA certificate is used, by default, to secure internal StorageGRID traffic. An external certificate authority (CA) can issue custom certificates that are fully compliant with your organization's information security policies. Although you can use the Grid CA certificate for a non-production environment, the best practice for a production environment is to use custom certificates signed by an external certificate authority. Unsecured connections with no certificate are also supported but aren't recommended.

- Custom CA certificates don't remove the internal certificates; however, the custom certificates should be the ones specified for verifying server connections.
- All custom certificates must meet the system hardening guidelines for server certificates.
- StorageGRID supports bundling of certificates from a CA into a single file (known as a CA certificate bundle).



StorageGRID also includes operating system CA certificates that are the same on all grids. In production environments, make sure that you specify a custom certificate signed by an external certificate authority in place of the operating system CA certificate.

Variants of the server and client certificate types are implemented in several ways. You should have all the certificates needed for your specific StorageGRID configuration ready before you configure the system.

Access security certificates

You can access information about all StorageGRID certificates in a single location, along with links to the configuration workflow for each certificate.

Steps

1. From Grid Manager, select **CONFIGURATION > Security > Certificates**.

~	cates that secure HT	TPS connections betwe	en StorageGRID and external clients, such	as S3 or Swift, and externa	al servers, such as a key management server (
Global	Grid CA	Client	Load balancer endpoints	Tenants	Other
e StorageGRID certificate	authority ("grid CA") generates and signs tw	o global certificates during installation. Th	he management interface (certificate on Admin Nodes secures the managed by an
ternal certificate authorit	y.	rage and Gateway Node	s secures client access, you should replac	e each deladit certificate w	nin your own custom certificate signed by an
Name		Description		Type 🔇	Expiration date 🧿 💠
		Secures the connecti	on between client web browsers and the	Grid	
Management interface co	ertificate	Manager, Tenant Mar Management API.	ager, Grid Management API, and Tenant	Custom	Jun 4th, 2022
		Secures the connecti	ons between S3 and Swift clients and Stor	rage	
S3 and Swift API certificate	te	Nodes or between clients and the deprecated CLB service on Gatewa		ateway Custom	Jun 4th, 2022
S3 and Swift API certifica	cc -	Nodac Vou can optio	0.0101111120 TRUE CONTINICATO TOY 0 10000 R010000		

- 2. Select a tab on the Certificates page for information about each certificate category and to access the certificate settings. You can access a tab if you have the appropriate permission.
 - Global: Secures StorageGRID access from web browsers and external API clients.
 - Grid CA: Secures internal StorageGRID traffic.
 - Client: Secures connections between external clients and the StorageGRID Prometheus database.
 - Load balancer endpoints: Secures connections between S3 and Swift clients and the StorageGRID Load Balancer.
 - **Tenants**: Secures connections to identity federation servers or from platform service endpoints to S3 storage resources.
 - Other: Secures StorageGRID connections requiring specific certificates.

Each tab is described below with links to additional certificate details.

Global

The global certificates secure StorageGRID access from web browsers and external S3 and Swift API clients. Two global certificates are initially generated by the StorageGRID certificate authority during installation. The best practice for a production environment is to use custom certificates signed by an external certificate authority.

- Management interface certificate: Secures client web-browser connections to StorageGRID management interfaces.
- S3 and Swift API certificate: Secures client API connections to Storage Nodes, Admin Nodes, and Gateway Nodes, which S3 and Swift client applications use to upload and download object data.

Information about the global certificates that are installed includes:

- Name: Certificate name with link to managing the certificate.
- Description
- **Type**: Custom or default. You should always use a custom certificate for improved grid security.
- Expiration date: If using the default certificate, no expiration date is shown.

You can:

- Replace the default certificates with custom certificates signed by an external certificate authority for improved grid security:
 - Replace the default StorageGRID-generated management interface certificate used for Grid Manager and Tenant Manager connections.
 - Replace the S3 and Swift API certificate used for Storage Node and load balancer endpoint (optional) connections.
- Restore the default management interface certificate.
- Restore the default S3 and Swift API certificate.
- Use a script to generate a new self-signed management interface certificate.
- Copy or download the management interface certificate or S3 and Swift API certificate.

Grid CA

The Grid CA certificate, generated by the StorageGRID certificate authority during StorageGRID installation, secures all internal StorageGRID traffic.

Certificate information includes the certificate expiration date and the certificate contents.

You can copy or download the Grid CA certificate, but you can't change it.

Client

Client certificates, generated by an external certificate authority, secure the connections between external monitoring tools and the StorageGRID Prometheus database.

The certificate table has a row for each configured client certificate and indicates whether the certificate can be used for Prometheus database access, along with the certificate expiration date.

You can:

- Upload or generate a new client certificate.
- Select a certificate name to display the certificate details where you can:
 - Change the client certificate name.
 - Set the Prometheus access permission.
 - Upload and replace the client certificate.
 - · Copy or download the client certificate.
 - Remove the client certificate.
- Select **Actions** to quickly edit, attach, or remove a client certificate. You can select up to 10 client certificates and remove them at one time using **Actions** > **Remove**.

Load balancer endpoints

Load balancer endpoint certificates secure the connections between S3 and Swift clients and the StorageGRID Load Balancer service on Gateway Nodes and Admin Nodes.

The load balancer endpoint table has a row for each configured load balancer endpoint and indicates whether the global S3 and Swift API certificate or a custom load balancer endpoint certificate is being used for the endpoint. The expiration date for each certificate is also displayed.



Changes to an endpoint certificate can take up to 15 minutes to be applied to all nodes.

You can:

- View a load balancer endpoint, including its certificate details.
- Specify a load balancer endpoint certificate for FabricPool.
- Use the global S3 and Swift API certificate instead of generating a new load balancer endpoint certificate.

Tenants

Tenants can use identity federation server certificates or platform service endpoint certificates to secure their connections with StorageGRID.

The tenant table has a row for each tenant and indicates if each tenant has permission to use its own identity source or platform services.

You can:

- · Select a tenant name to sign in to the Tenant Manager
- · Select a tenant name to view the tenant identity federation details
- · Select a tenant name to view tenant platform services details
- Specify a platform service endpoint certificate during endpoint creation

Other

StorageGRID uses other security certificates for specific purposes. These certificates are listed by their functional name. Other security certificates include:

- Cloud Storage Pool certificates
- Email alert notification certificates

- External syslog server certificates
- · Grid federation connection certificates
- Identity federation certificates
- Key management server (KMS) certificates
- Single sign-on certificates

Information indicates the type of certificate a function uses and its server and client certificate expiration dates, as applicable. Selecting a function name opens a browser tab where you can view and edit the certificate details.



You can only view and access information for other certificates if you have the appropriate permission.

You can:

- Specify a Cloud Storage Pool certificate for S3, C2S S3, or Azure
- Specify a certificate for alert email notifications
- Use a certificate for an external syslog server
- Rotate grid federation connection certificates
- · View and edit an identity federation certificate
- Upload key management server (KMS) server and client certificates
- · Manually specify an SSO certificate for a relying party trust

Security certificate details

Each type of security certificate is described below, with links to the implementation instructions.

Management interface certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates the connection between client web browsers and the StorageGRID management interface, allowing users to access the Grid Manager and Tenant Manager without security warnings. This certificate also authenticates Grid Management API and Tenant Management API connections. You can use the default certificate created during installation or upload a custom certificate.	CONFIGURATION > Security > Certificates, select the Global tab, and then select Management interface certificate	Configure management interface certificates

S3 and Swift API certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates secure S3 or Swift client connections to a Storage Node and to load balancer endpoints (optional).	CONFIGURATION > Security > Certificates, select the Global tab, and then select S3 and Swift API certificate	Configure S3 and Swift API certificates

Grid CA certificate

See the Default Grid CA certificate description.

Administrator client certificate

Certificate type	Description	Navigation location	Details
Client	 Installed on each client, allowing StorageGRID to authenticate external client access. Allows authorized external clients to access the StorageGRID Prometheus database. Allows secure monitoring of StorageGRID using external tools. 	CONFIGURATION > Security > Certificates and then select the Client tab	Configure client certificates

Load balancer endpoint certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates the connection between S3 or Swift clients and the StorageGRID Load Balancer service on Gateway Nodes and Admin Nodes. You can upload or generate a load balancer certificate when you configure a load balancer endpoint. Client applications use the load balancer certificate when connecting to StorageGRID to save and retrieve object data. You can also use a custom version of the global S3 and Swift API certificate certificate to authenticate connections to the Load Balancer service. If the global certificate is used to authenticate load balancer connections, you don't need to upload or generate a separate certificate for each load balancer endpoint. Note: The certificate used for load balancer authenticate during normal StorageGRID operation.	CONFIGURATION > Network > Load balancer endpoints	 Configure load balancer endpoints Create a load balancer endpoint for FabricPool

Cloud Storage Pool endpoint certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates the connection from a StorageGRID Cloud Storage Pool to an external storage location, such as S3 Glacier or Microsoft Azure Blob storage. A different certificate is required for each cloud provider type.	ILM > Storage pools	Create a Cloud Storage Pool

Email alert notification certificate

Certificate type	Description	Navigation location	Details
Server and client	 Authenticates the connection between an SMTP email server and StorageGRID that is used for alert notifications. If communications with the SMTP server requires Transport Layer Security (TLS), you must specify the email server CA certificate. Specify a client certificate only if the SMTP email server requires client certificates for authentication. 	ALERTS > Email setup	Set up email notifications for alerts

External syslog server certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates the TLS or RELP/TLS connection between an external syslog server that logs events in StorageGRID. Note: An external syslog server certificate is not required for TCP, RELP/TCP, and UDP connections to an external syslog server.	CONFIGURATION > Monitoring > Audit and syslog server	Use an external syslog server

Grid federation connection certificate

Certificate type	Description	Navigation location	Details
Server and client	Authenticate and encrypt information sent between the current StorageGRID system and another grid in a grid federation connection.	CONFIGURATION > System > Grid federation	 Create grid federation connections Rotate connection certificates

Identity federation certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates the connection between StorageGRID and an external identity provider, such as Active Directory, OpenLDAP, or Oracle Directory Server. Used for identity federation, which allows admin groups and users to be managed by an external system.	CONFIGURATION > Access Control > Identity federation	Use identity federation

Key management server (KMS) certificate

Certificate type	Description	Navigation location	Details
Server and client	Authenticates the connection between StorageGRID and an external key management server (KMS), which provides encryption keys to StorageGRID appliance nodes.	CONFIGURATION > Security > Key management server	Add key management server (KMS)

Platform services endpoint certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates the connection from the StorageGRID platform service to an S3 storage resource.	Tenant Manager > STORAGE (S3) > Platform services endpoints	Create platform services endpoint Edit platform services endpoint

Single sign-on (SSO) certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates the connection between identity federation services, such as Active Directory Federation Services (AD FS), and StorageGRID that are used for single sign-on (SSO) requests.	CONFIGURATION > Access control > Single sign-on	Configure single sign-on

Certificate examples

Example 1: Load Balancer service

In this example, StorageGRID acts as the server.

- 1. You configure a load balancer endpoint and upload or generate a server certificate in StorageGRID.
- 2. You configure an S3 or Swift client connection to the load balancer endpoint and upload the same certificate to the client.
- 3. When the client wants to save or retrieve data, it connects to the load balancer endpoint using HTTPS.
- 4. StorageGRID responds with the server certificate, which contains a public key, and with a signature based on the private key.
- 5. The client verifies this certificate by comparing the server signature to the signature on its copy of the certificate. If the signatures match, the client starts a session using the same public key.

6. The client sends object data to StorageGRID.

Example 2: External key management server (KMS)

In this example, StorageGRID acts as the client.

- 1. Using external Key Management Server software, you configure StorageGRID as a KMS client and obtain a CA-signed server certificate, a public client certificate, and the private key for the client certificate.
- 2. Using the Grid Manager, you configure a KMS server and upload the server and client certificates and the client private key.
- 3. When a StorageGRID node needs an encryption key, it makes a request to the KMS server that includes data from the certificate and a signature based on the private key.
- 4. The KMS server validates the certificate signature and decides that it can trust StorageGRID.
- 5. The KMS server responds using the validated connection.

Configure server certificates

Supported server certificate types

The StorageGRID system supports custom certificates encrypted with RSA or ECDSA (Elliptic Curve Digital Signature Algorithm).



The cipher type for the security policy must match the server certificate type. For example, RSA ciphers require RSA certificates, and ECDSA ciphers require ECDSA certificates. See Manage security certificates. If you configure a custom security policy that is not compatible with the server certificate, you can temporarily revert to the default security policy.

For more information about how StorageGRID secures client connections, see Security for S3 and Swift clients.

Configure management interface certificates

You can replace the default management interface certificate with a single custom certificate that allows users to access the Grid Manager and the Tenant Manager without encountering security warnings. You can also revert to the default management interface certificate or generate a new one.

About this task

By default, every Admin Node is issued a certificate signed by the grid CA. These CA signed certificates can be replaced by a single common custom management interface certificate and corresponding private key.

Because a single custom management interface certificate is used for all Admin Nodes, you must specify the certificate as a wildcard or multi-domain certificate if clients need to verify the hostname when connecting to the Grid Manager and Tenant Manager. Define the custom certificate such that it matches all Admin Nodes in the grid.

You need to complete configuration on the server, and depending on the root certificate authority (CA) you are using, users might also need to install the Grid CA certificate in the web browser they will use to access the Grid Manager and the Tenant Manager.

(i)

 (\mathbf{i})

To ensure that operations aren't disrupted by a failed server certificate, the **Expiration of server** certificate for Management Interface alert is triggered when this server certificate is about to expire. As required, you can view when the current certificate expires by selecting CONFIGURATION > Security > Certificates and looking at the Expiration date for the management interface certificate on the Global tab.

If you are accessing the Grid Manager or Tenant Manager using a domain name instead of an IP address, the browser shows a certificate error without an option to bypass if either of the following occurs:

- · Your custom management interface certificate expires.
- You revert from a custom management interface certificate to the default server certificate.

Add a custom management interface certificate

To add a custom management interface certificate, you can provide your own certificate or generate one using the Grid Manager.

Steps

- 1. Select CONFIGURATION > Security > Certificates.
- 2. On the Global tab, select Management interface certificate.
- 3. Select Use custom certificate.
- 4. Upload or generate the certificate.

Upload certificate

Upload the required server certificate files.

- a. Select Upload certificate.
- b. Upload the required server certificate files:
 - Server certificate: The custom server certificate file (PEM encoded).
 - Certificate private key: The custom server certificate private key file (.key).



EC private keys must be 224 bits or larger. RSA private keys must be 2048 bits or larger.

- **CA bundle**: A single optional file containing the certificates from each intermediate issuing certificate authority (CA). The file should contain each of the PEM-encoded CA certificate files, concatenated in certificate chain order.
- c. Expand **Certificate details** to see the metadata for each certificate you uploaded. If you uploaded an optional CA bundle, each certificate displays on its own tab.
 - Select Download certificate to save the certificate file or select Download CA bundle to save the certificate bundle.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid certificate.pem

- Select Copy certificate PEM or Copy CA bundle PEM to copy the certificate contents for pasting elsewhere.
- d. Select Save.

The custom management interface certificate is used for all subsequent new connections to the Grid Manager, Tenant Manager, Grid Manager API or Tenant Manager API.

Generate certificate

Generate the server certificate files.



The best practice for a production environment is to use a custom management interface certificate signed by an external certificate authority.

- a. Select Generate certificate.
- b. Specify the certificate information:

Field	Description
Domain name	One or more fully qualified domain names to include in the certificate. Use an * as a wildcard to represent multiple domain names.
IP	One or more IP addresses to include in the certificate.

Field	Description
Subject (optional)	X.509 subject or distinguished name (DN) of the certificate owner. If no value is entered in this field, the generated certificate uses the first domain name or IP address as the subject common name (CN).
Days valid	Number of days after creation that the certificate expires.
Add key usage extensions	If selected (default and recommended), key usage and extended key usage extensions are added to the generated certificate. These extensions define the purpose of the key contained in the certificate. Note : Leave this checkbox selected unless you experience connection problems with older clients when certificates include these extensions.

- c. Select Generate.
- d. Select Certificate details to see the metadata for the generated certificate.
 - Select Download certificate to save the certificate file.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid certificate.pem

- Select Copy certificate PEM to copy the certificate contents for pasting elsewhere.
- e. Select Save.

The custom management interface certificate is used for all subsequent new connections to the Grid Manager, Tenant Manager, Grid Manager API or Tenant Manager API.

5. Refresh the page to ensure the web browser is updated.



After uploading or generating a new certificate, allow up to one day for any related certificate expiration alerts to clear.

 After you add a custom management interface certificate, the Management interface certificate page displays detailed certificate information for the certificates that are in use. You can download or copy the certificate PEM as required.

Restore the default management interface certificate

You can revert to using the default management interface certificate for Grid Manager and Tenant Manager connections.

Steps

- 1. Select **CONFIGURATION > Security > Certificates**.
- 2. On the Global tab, select Management interface certificate.
- 3. Select Use default certificate.

When you restore the default management interface certificate, the custom server certificate files you configured are deleted and can't be recovered from the system. The default management interface certificate is used for all subsequent new client connections.

4. Refresh the page to ensure the web browser is updated.

Use a script to generate a new self-signed management interface certificate

If strict hostname validation is required, you can use a script to generate the management interface certificate.

Before you begin

- You have specific access permissions.
- You have the Passwords.txt file.

About this task

The best practice for a production environment is to use a certificate signed by an external certificate authority.

Steps

- 1. Obtain the fully qualified domain name (FQDN) of each Admin Node.
- 2. Log in to the primary Admin Node:
 - a. Enter the following command: ssh admin@primary_Admin_Node_IP
 - b. Enter the password listed in the Passwords.txt file.
 - c. Enter the following command to switch to root: su -
 - d. Enter the password listed in the Passwords.txt file.

When you are logged in as root, the prompt changes from \$ to #.

- 3. Configure StorageGRID with a new self-signed certificate.
 - \$ sudo make-certificate --domains wildcard-admin-node-fqdn --type management
 - For --domains, use wildcards to represent the fully qualified domain names of all Admin Nodes. For example, *.ui.storagegrid.example.com uses the * wildcard to represent admin1.ui.storagegrid.example.com and admin2.ui.storagegrid.example.com.
 - Set --type to management to configure the management interface certificate, which is used by Grid Manager and Tenant Manager.
 - By default, generated certificates are valid for one year (365 days) and must be recreated before they expire. You can use the --days argument to override the default validity period.



A certificate's validity period begins when make-certificate is run. You must ensure the management client is synchronized to the same time source as StorageGRID; otherwise, the client might reject the certificate.

```
$ sudo make-certificate --domains *.ui.storagegrid.example.com --type
management --days 720
```

The resulting output contains the public certificate needed by your management API client.

4. Select and copy the certificate.

Include the BEGIN and the END tags in your selection.

- 5. Log out of the command shell. \$ exit
- 6. Confirm the certificate was configured:
 - a. Access the Grid Manager.
 - b. Select **CONFIGURATION > Security > Certificates**
 - c. On the Global tab, select Management interface certificate.
- 7. Configure your management client to use the public certificate you copied. Include the BEGIN and END tags.

Download or copy the management interface certificate

You can save or copy the management interface certificate contents for use elsewhere.

Steps

- 1. Select CONFIGURATION > Security > Certificates.
- 2. On the Global tab, select Management interface certificate.
- 3. Select the **Server** or **CA bundle** tab and then download or copy the certificate.

Download certificate file or CA bundle

Download the certificate or CA bundle .pem file. If you are using an optional CA bundle, each certificate in the bundle displays on its own sub-tab.

a. Select Download certificate or Download CA bundle.

If you are downloading a CA bundle, all the certificates in the CA bundle secondary tabs download as a single file.

b. Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid_certificate.pem

Copy certificate or CA bundle PEM

Copy the certificate text to paste elsewhere. If you are using an optional CA bundle, each certificate in the bundle displays on its own sub-tab.

a. Select Copy certificate PEM or Copy CA bundle PEM.

If you are copying a CA bundle, all the certificates in the CA bundle secondary tabs copy together.

b. Paste the copied certificate into a text editor.

c. Save the text file with the extension .pem.

For example: storagegrid certificate.pem

Configure S3 and Swift API certificates

You can replace or restore the server certificate that is used for S3 or Swift client connections to Storage Nodes or to load balancer endpoints. The replacement custom server certificate is specific to your organization.

About this task

By default, every Storage Node is issued a X.509 server certificate signed by the grid CA. These CA signed certificates can be replaced by a single common custom server certificate and corresponding private key.

A single custom server certificate is used for all Storage Nodes, so you must specify the certificate as a wildcard or multi-domain certificate if clients need to verify the hostname when connecting to the storage endpoint. Define the custom certificate such that it matches all Storage Nodes in the grid.

After completing configuration on the server, you might also need to install the Grid CA certificate in the S3 or Swift API client you will use to access the system, depending on the root certificate authority (CA) you are using.



To ensure that operations aren't disrupted by a failed server certificate, the **Expiration of global** server certificate for S3 and Swift API alert is triggered when the root server certificate is about to expire. As required, you can view when the current certificate expires by selecting **CONFIGURATION > Security > Certificates** and looking at the Expiration date for the S3 and Swift API certificate on the Global tab. You can upload or generate a custom S3 and Swift API certificate.

Add a custom S3 and Swift API certificate

Steps

- 1. Select CONFIGURATION > Security > Certificates.
- 2. On the Global tab, select S3 and Swift API certificate.
- 3. Select Use custom certificate.
- 4. Upload or generate the certificate.

Upload certificate

Upload the required server certificate files.

- a. Select Upload certificate.
- b. Upload the required server certificate files:
 - Server certificate: The custom server certificate file (PEM encoded).
 - Certificate private key: The custom server certificate private key file (.key).



EC private keys must be 224 bits or larger. RSA private keys must be 2048 bits or larger.

- CA bundle: A single optional file containing the certificates from each intermediate issuing certificate authority. The file should contain each of the PEM-encoded CA certificate files, concatenated in certificate chain order.
- c. Select the certificate details to display the metadata and PEM for each custom S3 and Swift API certificate that was uploaded. If you uploaded an optional CA bundle, each certificate displays on its own tab.
 - Select Download certificate to save the certificate file or select Download CA bundle to save the certificate bundle.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid certificate.pem

- Select Copy certificate PEM or Copy CA bundle PEM to copy the certificate contents for pasting elsewhere.
- d. Select Save.

The custom server certificate is used for subsequent new S3 and Swift client connections.

Generate certificate

Generate the server certificate files.

- a. Select Generate certificate.
- b. Specify the certificate information:

Field	Description
Domain name	One or more fully qualified domain names to include in the certificate. Use an * as a wildcard to represent multiple domain names.
IP	One or more IP addresses to include in the certificate.
Subject (optional)	X.509 subject or distinguished name (DN) of the certificate owner. If no value is entered in this field, the generated certificate uses the first domain name or IP address as the subject common name (CN).

Field	Description
Days valid	Number of days after creation that the certificate expires.
Add key usage extensions	If selected (default and recommended), key usage and extended key usage extensions are added to the generated certificate. These extensions define the purpose of the key contained in the certificate.
	Note : Leave this checkbox selected unless you experience connection problems with older clients when certificates include these extensions.

- c. Select Generate.
- d. Select **Certificate Details** to display the metadata and PEM for the custom S3 and Swift API certificate that was generated.
 - Select **Download certificate** to save the certificate file.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid_certificate.pem

- Select Copy certificate PEM to copy the certificate contents for pasting elsewhere.
- e. Select Save.

The custom server certificate is used for subsequent new S3 and Swift client connections.

5. Select a tab to display metadata for the default StorageGRID server certificate, a CA signed certificate that was uploaded, or a custom certificate that was generated.



After uploading or generating a new certificate, allow up to one day for any related certificate expiration alerts to clear.

- 6. Refresh the page to ensure the web browser is updated.
- After you add a custom S3 and Swift API certificate the S3 and Swift API certificate page displays detailed certificate information for the custom S3 and Swift API certificate that is in use. You can download or copy the certificate PEM as required.

Restore the default S3 and Swift API certificate

You can revert to using the default S3 and Swift API certificate for S3 and Swift client connections to Storage Nodes. However, you can't use the default S3 and Swift API certificate for a load balancer endpoint.

Steps

- 1. Select CONFIGURATION > Security > Certificates.
- 2. On the Global tab, select S3 and Swift API certificate.
- 3. Select Use default certificate.

When you restore the default version of the global S3 and Swift API certificate, the custom server certificate files you configured are deleted and can't be recovered from the system. The default S3 and Swift API certificate will be used for subsequent new S3 and Swift client connections to Storage Nodes.

4. Select **OK** to confirm the warning and restore the default S3 and Swift API certificate.

If you have Root access permission and the custom S3 and Swift API certificate was used for load balancer endpoint connections, a list is displayed of load balancer endpoints that will no longer be accessible using the default S3 and Swift API certificate. Go to Configure load balancer endpoints to edit or remove the affected endpoints.

5. Refresh the page to ensure the web browser is updated.

Download or copy the S3 and Swift API certificate

You can save or copy the S3 and Swift API certificate contents for use elsewhere.

Steps

- 1. Select CONFIGURATION > Security > Certificates.
- 2. On the **Global** tab, select **S3 and Swift API certificate**.
- 3. Select the **Server** or **CA bundle** tab and then download or copy the certificate.

Download certificate file or CA bundle

Download the certificate or CA bundle .pem file. If you are using an optional CA bundle, each certificate in the bundle displays on its own sub-tab.

a. Select Download certificate or Download CA bundle.

If you are downloading a CA bundle, all the certificates in the CA bundle secondary tabs download as a single file.

b. Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid_certificate.pem

Copy certificate or CA bundle PEM

Copy the certificate text to paste elsewhere. If you are using an optional CA bundle, each certificate in the bundle displays on its own sub-tab.

a. Select Copy certificate PEM or Copy CA bundle PEM.

If you are copying a CA bundle, all the certificates in the CA bundle secondary tabs copy together.

- b. Paste the copied certificate into a text editor.
- c. Save the text file with the extension .pem.

For example: storagegrid_certificate.pem

Related information

- Use S3 REST API
- Use Swift REST API
- Configure S3 endpoint domain names

Copy the Grid CA certificate

StorageGRID uses an internal certificate authority (CA) to secure internal traffic. This certificate does not change if you upload your own certificates.

Before you begin

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

About this task

If a custom server certificate has been configured, client applications should verify the server using the custom server certificate. They should not copy the CA certificate from the StorageGRID system.

Steps

- 1. Select CONFIGURATION > Security > Certificates and then select the Grid CA tab.
- 2. In the Certificate PEM section, download or copy the certificate.

Download certificate file

Download the certificate .pem file.

- a. Select Download certificate.
- b. Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid_certificate.pem

Copy certificate PEM

Copy the certificate text to paste elsewhere.

- a. Select Copy certificate PEM.
- b. Paste the copied certificate into a text editor.
- c. Save the text file with the extension .pem.

For example: storagegrid_certificate.pem

Configure StorageGRID certificates for FabricPool

For S3 clients that perform strict hostname validation and don't support disabling strict hostname validation, such as ONTAP clients using FabricPool, you can generate or upload a server certificate when you configure the load balancer endpoint.

Before you begin

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

About this task

When you create a load balancer endpoint, you can generate a self-signed server certificate or upload a certificate that is signed by a known certificate authority (CA). In production environments, you should use a certificate that is signed by a known CA. Certificates signed by a CA can be rotated non-disruptively. They are also more secure because they provide better protection against man-in-the-middle attacks.

The following steps provide general guidelines for S3 clients that use FabricPool. For more detailed information and procedures, see Configure StorageGRID for FabricPool.

Steps

- 1. Optionally, configure a high availability (HA) group for FabricPool to use.
- 2. Create an S3 load balancer endpoint for FabricPool to use.

When you create an HTTPS load balancer endpoint, you are prompted to upload your server certificate, certificate private key, and optional CA bundle.

3. Attach StorageGRID as a cloud tier in ONTAP.

Specify the load balancer endpoint port and the fully qualified domain name used in the CA certificate you uploaded. Then, provide the CA certificate.



If an intermediate CA issued the StorageGRID certificate, you must provide the intermediate CA certificate. If the StorageGRID certificate was issued directly by the Root CA, you must provide the Root CA certificate.

Configure client certificates

Client certificates allow authorized external clients to access the StorageGRID Prometheus database, providing a secure way for external tools to monitor StorageGRID.

If you need to access StorageGRID using an external monitoring tool, you must upload or generate a client certificate using the Grid Manager and copy the certificate information to the external tool.

See Manage security certificates and Configure custom server certificates.



To ensure that operations aren't disrupted by a failed server certificate, the **Expiration of client** certificates configured on the Certificates page alert is triggered when this server certificate is about to expire. As required, you can view when the current certificate expires by selecting CONFIGURATION > Security > Certificates and looking at the Expiration date for the client certificate on the Client tab.



If you are using a key management server (KMS) to protect the data on specially configured appliance nodes, see the specific information about <u>uploading a KMS client certificate</u>.

Before you begin

- You have Root access permission.
- You are signed in to the Grid Manager using a supported web browser.

- To configure a client certificate:
 - You have the IP address or domain name of the Admin Node.
 - If you have configured the StorageGRID management interface certificate, you have the CA, client certificate, and private key used to configure the management interface certificate.
 - To upload your own certificate, the private key for the certificate is available on your local computer.
 - The private key must have been saved or recorded at the time it was created. If you don't have the original private key, you must create a new one.
- To edit a client certificate:
 - $\circ\,$ You have the IP address or domain name of the Admin Node.
 - To upload your own certificate or a new certificate, the private key, client certificate, and CA (if used) are available on your local computer.

Add client certificates

To add the client certificate, use one of these procedures:

- Management interface certificate already configured
- CA issued client certificate
- · Generated certificate from Grid Manager

Management interface certificate already configured

Use this procedure to add a client certificate if a management interface certificate is already configured using a customer-supplied CA, client certificate, and private key.

Steps

- 1. In the Grid Manager, select **CONFIGURATION > Security > Certificates** and then select the **Client** tab.
- 2. Select Add.
- 3. Enter a certificate name.
- 4. To access Prometheus metrics using your external monitoring tool, select Allow prometheus.
- 5. Select Continue.
- 6. For the **Attach certificates** step, upload the management interface certificate.
 - a. Select Upload certificate.
 - b. Select Browse and select the management interface certificate file (.pem).
 - Select Client certificate details to display the certificate metadata and certificate PEM.
 - Select Copy certificate PEM to copy the certificate contents for pasting elsewhere.
 - c. Select **Create** to save the certificate in the Grid Manager.

The new certificate appears on the Client tab.

7. Configure an external monitoring tool, such as Grafana.

CA issued client certificate

Use this procedure to add an administrator client certificate if a management interface certificate was not

configured and you plan to add a client certificate for Prometheus that uses a CA issued client certificate and private key.

Steps

- 1. Perform the steps to configure a management interface certificate.
- 2. In the Grid Manager, select CONFIGURATION > Security > Certificates and then select the Client tab.
- 3. Select Add.
- 4. Enter a certificate name.
- 5. To access Prometheus metrics using your external monitoring tool, select Allow prometheus.
- 6. Select Continue.
- 7. For the **Attach certificates** step, upload the client certificate, private key, and CA bundle files:
 - a. Select Upload certificate.
 - b. Select Browse and select the client certificate, private key, and CA bundle files (.pem).
 - Select Client certificate details to display the certificate metadata and certificate PEM.
 - Select Copy certificate PEM to copy the certificate contents for pasting elsewhere.
 - c. Select Create to save the certificate in the Grid Manager.

The new certificates appear on the Client tab.

8. Configure an external monitoring tool, such as Grafana.

Generated certificate from Grid Manager

Use this procedure to add an administrator client certificate if a management interface certificate was not configured and you plan to add a client certificate for Prometheus that uses the generate certificate function in Grid Manager.

Steps

- 1. In the Grid Manager, select **CONFIGURATION > Security > Certificates** and then select the **Client** tab.
- 2. Select Add.
- 3. Enter a certificate name.
- 4. To access Prometheus metrics using your external monitoring tool, select Allow prometheus.
- 5. Select Continue.
- 6. For the Attach certificates step, select Generate certificate.
- 7. Specify the certificate information:
 - Subject (optional): X.509 subject or distinguished name (DN) of the certificate owner.
 - **Days valid**: The number of days the generated certificate is valid, starting at the time it is generated.
 - Add key usage extensions: If selected (default and recommended), key usage and extended key usage extensions are added to the generated certificate.

These extensions define the purpose of the key contained in the certificate.



Leave this checkbox selected unless you experience connection problems with older clients when certificates include these extensions.

8. Select Generate.

9. Select **Client certificate details** to display the certificate metadata and certificate PEM.



You will not be able to view the certificate private key after you close the dialog. Copy or download the key to a safe location.

- Select Copy certificate PEM to copy the certificate contents for pasting elsewhere.
- Select **Download certificate** to save the certificate file.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid certificate.pem

- Select Copy private key to copy the certificate private key for pasting elsewhere.
- Select **Download private key** to save the private key as a file.

Specify the private key file name and download location.

10. Select Create to save the certificate in the Grid Manager.

The new certificate appears on the Client tab.

- 11. In the Grid Manager, select **CONFIGURATION > Security > Certificates** and then select the **Global** tab.
- 12. Select Management Interface certificate.
- 13. Select Use custom certificate.
- 14. Upload the certificate.pem and private_key.pem files from the client certificate details step. There is no need to upload CA bundle.
 - a. Select Upload certificate and then select Continue.
 - b. Upload each certificate file (.pem).
 - c. Select **Save** to save the certificate in the Grid Manager.

The new certificate appears on the Management Interface certificate page.

15. Configure an external monitoring tool, such as Grafana.

Configure an external monitoring tool

Steps

- 1. Configure the following settings on your external monitoring tool, such as Grafana.
 - a. **Name**: Enter a name for the connection.

StorageGRID does not require this information, but you must provide a name to test the connection.

b. URL: Enter the domain name or IP address for the Admin Node. Specify HTTPS and port 9091.

For example: https://admin-node.example.com:9091

- c. Enable TLS Client Auth and With CA Cert.
- d. Under TLS/SSL Auth Details, copy and paste: +
- The management interface CA certificate to CA Cert
- The client certificate to Client Cert
- The private key to **Client Key**
- e. ServerName: Enter the domain name of the Admin Node.

ServerName must match the domain name as it appears in the management interface certificate.

2. Save and test the certificate and private key that you copied from StorageGRID or a local file.

You can now access the Prometheus metrics from StorageGRID with your external monitoring tool.

For information about the metrics, see the instructions for monitoring StorageGRID.

Edit client certificates

You can edit an administrator client certificate to change its name, enable or disable Prometheus access, or upload a new certificate when the current one has expired.

Steps

1. Select **CONFIGURATION > Security > Certificates** and then select the **Client** tab.

Certificate expiration dates and Prometheus access permissions are listed in the table. If a certificate will expire soon or is already expired, a message appears in the table and an alert is triggered.

- 2. Select the certificate you want to edit.
- 3. Select Edit and then select Edit name and permission
- 4. Enter a certificate name.
- 5. To access Prometheus metrics using your external monitoring tool, select Allow prometheus.
- 6. Select **Continue** to save the certificate in the Grid Manager.

The updated certificate displays on the Client tab.

Attach new client certificate

You can upload a new certificate when the current one has expired.

Steps

1. Select CONFIGURATION > Security > Certificates and then select the Client tab.

Certificate expiration dates and Prometheus access permissions are listed in the table. If a certificate will expire soon or is already expired, a message appears in the table and an alert is triggered.

- 2. Select the certificate you want to edit.
- 3. Select **Edit** and then select an edit option.

Upload certificate

Copy the certificate text to paste elsewhere.

- a. Select **Upload certificate** and then select **Continue**.
- b. Upload the client certificate name (.pem).

Select Client certificate details to display the certificate metadata and certificate PEM.

• Select **Download certificate** to save the certificate file.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid certificate.pem

- Select Copy certificate PEM to copy the certificate contents for pasting elsewhere.
- c. Select Create to save the certificate in the Grid Manager.

The updated certificate displays on the Client tab.

Generate certificate

Generate the certificate text to paste elsewhere.

- a. Select Generate certificate.
- b. Specify the certificate information:
 - Subject (optional): X.509 subject or distinguished name (DN) of the certificate owner.
 - **Days valid**: The number of days the generated certificate is valid, starting at the time it is generated.
 - Add key usage extensions: If selected (default and recommended), key usage and extended key usage extensions are added to the generated certificate.

These extensions define the purpose of the key contained in the certificate.



Leave this checkbox selected unless you experience connection problems with older clients when certificates include these extensions.

- c. Select Generate.
- d. Select Client certificate details to display the certificate metadata and certificate PEM.



You will not be able to view the certificate private key after you close the dialog. Copy or download the key to a safe location.

- Select Copy certificate PEM to copy the certificate contents for pasting elsewhere.
- Select **Download certificate** to save the certificate file.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid_certificate.pem

- Select Copy private key to copy the certificate private key for pasting elsewhere.
- Select Download private key to save the private key as a file.

Specify the private key file name and download location.

e. Select Create to save the certificate in the Grid Manager.

The new certificate appears on the Client tab.

Download or copy client certificates

You can download or copy a client certificate for use elsewhere.

Steps

- 1. Select CONFIGURATION > Security > Certificates and then select the Client tab.
- 2. Select the certificate you want to copy or download.
- 3. Download or copy the certificate.

Download certificate file

Download the certificate .pem file.

- a. Select Download certificate.
- b. Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid certificate.pem

Copy certificate

Copy the certificate text to paste elsewhere.

- a. Select Copy certificate PEM.
- b. Paste the copied certificate into a text editor.
- c. Save the text file with the extension .pem.

For example: storagegrid certificate.pem

Remove client certificates

If you no longer need an administrator client certificate, you can remove it.

Steps

- 1. Select **CONFIGURATION > Security > Certificates** and then select the **Client** tab.
- 2. Select the certificate you want to remove.
- 3. Select **Delete** and then confirm.



To remove up to 10 certificates, select each certificate to remove on the Client tab and then select **Actions > Delete**.

After a certificate is removed, clients that used the certificate must specify a new client certificate to access the StorageGRID Prometheus database.

Configure security settings

Manage the TLS and SSH policy

The TLS and SSH policy determines which protocols and ciphers are used to establish secure TLS connections with client applications and secure SSH connections to internal StorageGRID services.

The security policy controls how TLS and SSH encrypt data in motion. In general, use the Modern compatibility (default) policy, unless your system needs to be Common Criteria-compliant or you need to use other ciphers.



Some StorageGRID services have not been updated to use the ciphers in these policies.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.

Select a security policy

Steps

1. Select CONFIGURATION > Security > Security settings.

The **TLS and SSH policies** tab shows the available policies. The currently active policy is noted by a green check mark on the policy tile.



2. Review the tiles to learn about the available policies.

Policy	Description
Modern compatibility (default)	Use the default policy if you need strong encryption and unless you have special requirements. This policy is compatible with most TLS and SSH clients.

Policy	Description
Legacy compatibility	Use this policy if you need additional compatibility options for older clients. The additional options in this policy might make it less secure than the Modern compatibility policy.
Common Criteria	Use this policy if you require Common Criteria certification.
FIPS strict	Use this policy if you require Common Criteria certification and need to use the NetApp Cryptographic Security Module 3.0.8 for external client connections to load balancer endpoints, Tenant Manager, and Grid Manager. Using this policy might reduce performance. Note : After you select this policy, all nodes must be rebooted in a rolling fashion to activate the NetApp Cryptographic Security Module. Use Maintenance > Rolling reboot to initiate and monitor reboots.
Custom	Create a custom policy if you need to apply your own ciphers.

- 3. To see details about each policy's ciphers, protocols, and algorithms, select View details.
- 4. To change the current policy, select **Use policy**.

A green check mark appears next to Current policy on the policy tile.

Create a custom security policy

You can create a custom policy if you need to apply your own ciphers.

Steps

- 1. From the tile of the policy that is the most similar to the custom policy you want to create, select **View** details.
- 2. Select Copy to clipboard, and then select Cancel.



3. From the Custom policy tile, select Configure and use.

- 4. Paste the JSON you copied and make any changes required.
- 5. Select Use policy.

A green check mark appears next to Current policy on the Custom policy tile.

6. Optionally, select Edit configuration to make more changes to the new custom policy.

Temporarily revert to the default security policy

If you configured a custom security policy, you might not be able to sign in to the Grid Manager if the configured TLS policy is incompatible with the configured server certificate.

You can temporarily revert to the default security policy.

Steps

- 1. Log in to an Admin Node:
 - a. Enter the following command: ssh admin@Admin_Node_IP
 - b. Enter the password listed in the Passwords.txt file.
 - c. Enter the following command to switch to root: su -
 - d. Enter the password listed in the <code>Passwords.txt</code> file.

When you are logged in as root, the prompt changes from \$ to #.

2. Run the following command:

restore-default-cipher-configurations

- 3. From a web browser, access the Grid Manager on the same Admin Node.
- 4. Follow the steps in Select a security policy to configure the policy again.

Configure network and object security

You can configure network and object security to encrypt stored objects, to prevent certain S3 and Swift requests, or to allow client connections to Storage Nodes to use HTTP instead of HTTPS.

Stored object encryption

Stored object encryption enables the encryption of all object data as it is ingested through S3. By default, stored objects aren't encrypted but you can choose to encrypt objects using the AES-128 or AES-256 encryption algorithm. When you enable the setting, all newly ingested objects are encrypted but no change is made to existing stored objects. If you disable encryption, currently encrypted objects remain encrypted but newly ingested objects aren't encrypted.

The Stored object encryption setting applies only to S3 objects that have not been encrypted by bucket-level or object-level encryption.

For more details on StorageGRID encryption methods, see Review StorageGRID encryption methods.

Prevent client modification

Prevent client modification is a system wide setting. When the **Prevent client modification** option is selected, the following requests are denied.

S3 REST API

- DeleteBucket requests
- · Any requests to modify an existing object's data, user-defined metadata, or S3 object tagging

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.

Enable HTTP for Storage Node connections

By default, client applications use the HTTPS network protocol for any direct connections to Storage Nodes. You can optionally enable HTTP for these connections, for example, when testing a non-production grid.

Use HTTP for Storage Node connections only if S3 and Swift clients need to make HTTP connections directly to Storage Nodes. You don't need to use this option for clients that only use HTTPS connections or for clients that connect to the Load Balancer service (because you can configure each load balancer endpoint to use either HTTP or HTTPS).

See Summary: IP addresses and ports for client connections to learn which ports S3 and Swift clients use when connecting to Storage Nodes using HTTP or HTTPS.

Select options

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have Root access permission.

Steps

- 1. Select CONFIGURATION > Security > Security settings.
- 2. Select the Network and objects tab.
- 3. For Stored object encryption, use the **None** (default) setting if you don't want stored objects to be encrypted, or select **AES-128** or **AES-256** to encrypt stored objects.
- Optionally select Prevent client modification if you want to prevent S3 and Swift clients from making specific requests.



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

5. Optionally select **Enable HTTP for Storage Node connections** if clients connect directly to Storage Nodes and you want to use HTTP connections.



Be careful when enabling HTTP for a production grid because requests will be sent unencrypted.

6. Select Save.

Change interface security settings

The interface security settings let you control whether users are signed out if they are inactive for more than the specified amount of time and whether a stack trace is included in API error responses.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have Root access permission.

About this task

The **Security settings** page includes the **Browser inactivity timeout** and **Management API stack trace** settings.

Browser inactivity timeout

Indicates how long a user's browser can be inactive before the user is signed out. The default is 15 minutes.

Browser inactivity timeout is also controlled by the following:

- A separate, non-configurable StorageGRID timer, which is included for system security. Each user's authentication token expires 16 hours after the user signs in. When a user's authentication expires, that user is automatically signed out, even if browser inactivity timeout is disabled or the value for the browser timeout has not been reached. To renew the token, the user must sign back in.
- Timeout settings for the identity provider, assuming single sign-on (SSO) is enabled for StorageGRID.

If SSO is enabled and a user's browser times out, the user must reenter their SSO credentials to access StorageGRID again. See Configure single sign-on.

Management API stack trace

Controls whether a stack trace is returned in Grid Manager and Tenant Manager API error responses.

This option is disabled by default, but you might want to enable this functionality for a test environment. In general, you should leave stack trace disabled in production environments to avoid revealing internal software details when API errors occur.

Steps

- 1. Select CONFIGURATION > Security > Security settings.
- 2. Select the Interface tab.
- 3. To change the setting for browser inactivity timeout:
 - a. Expand the accordion.
 - b. To change the timeout period, specify a value between 60 seconds and 7 days. The default timeout is 15 minutes.
 - c. To disable this feature, unselect the checkbox.

d. Select Save.

The new setting doesn't affect users who are currently signed in. Users must sign in again or refresh their browsers for the new timeout setting to take effect.

- 4. To change the setting for Management API stack trace:
 - a. Expand the accordion.
 - b. Select the checkbox to return a stack trace in Grid Manager and Tenant Manager API error responses.



Leave stack trace disabled in production environments to avoid revealing internal software details when API errors occur.

c. Select Save.

Configure key management servers

Configure key management servers: Overview

You can configure one or more external key management servers (KMS) to protect the data on specially configured appliance nodes.



StorageGRID supports only certain key management servers. For a list of supported products and versions, use the NetApp Interoperability Matrix Tool (IMT).

What is a key management server (KMS)?

A key management server (KMS) is an external, third-party system that provides encryption keys to StorageGRID appliance nodes at the associated StorageGRID site using the Key Management Interoperability Protocol (KMIP).

You can use one or more key management servers to manage the node encryption keys for any StorageGRID appliance nodes that have the **Node Encryption** setting enabled during installation. Using key management servers with these appliance nodes lets you protect your data even if an appliance is removed from the data center. After the appliance volumes are encrypted, you can't access any data on the appliance unless the node can communicate with the KMS.



StorageGRID does not create or manage the external keys used to encrypt and decrypt appliance nodes. If you plan to use an external key management server to protect StorageGRID data, you must understand how to set up that server, and you must understand how to manage the encryption keys. Performing key management tasks is beyond the scope of these instructions. If you need help, see the documentation for your key management server or contact technical support.

Overview of KMS and appliance configuration

Before you can use a key management server (KMS) to secure StorageGRID data on appliance nodes, you must complete two configuration tasks: setting up one or more KMS servers and enabling node encryption for the appliance nodes. When these two configuration tasks are complete, the key management process occurs automatically.

The flowchart shows the high-level steps for using a KMS to secure StorageGRID data on appliance nodes.



The flowchart shows KMS setup and appliance setup occurring in parallel; however, you can set up the key

management servers before or after you enable node encryption for new appliance nodes, based on your requirements.

Set up the key management server (KMS)

Setting up a key management server includes the following high-level steps.

Step	Refer to
Access the KMS software and add a client for StorageGRID to each KMS or KMS cluster.	Configure StorageGRID as a client in the KMS
Obtain the required information for the StorageGRID client on the KMS.	Configure StorageGRID as a client in the KMS
Add the KMS to the Grid Manager, assign it to a single site or to a default group of sites, upload the required certificates, and save the KMS configuration.	Add a key management server (KMS)

Set up the appliance

Setting up an appliance node for KMS use includes the following high-level steps.

1. During the hardware configuration stage of appliance installation, use the StorageGRID Appliance Installer to enable the **Node Encryption** setting for the appliance.



You can't enable the **Node Encryption** setting after an appliance is added to the grid, and you can't use external key management for appliances that don't have node encryption enabled.

- Run the StorageGRID Appliance Installer. During installation, a random data encryption key (DEK) is assigned to each appliance volume, as follows:
 - The DEKs are used to encrypt the data on each volume. These keys are generated using Linux Unified Key Setup (LUKS) disk encryption in the appliance OS and can't be changed.
 - Each individual DEK is encrypted by a master key encryption key (KEK). The initial KEK is a temporary key that encrypts the DEKs until the appliance can connect to the KMS.
- 3. Add the appliance node to StorageGRID.

See Enable node encryption for details.

Key management encryption process (occurs automatically)

Key management encryption includes the following high-level steps that are performed automatically.

- 1. When you install an appliance that has node encryption enabled into the grid, StorageGRID determines if a KMS configuration exists for the site that contains the new node.
 - If a KMS has already been configured for the site, the appliance receives the KMS configuration.
 - If a KMS has not yet been configured for the site, data on the appliance continues to be encrypted by the temporary KEK until you configure a KMS for the site and the appliance receives the KMS configuration.

- 2. The appliance uses the KMS configuration to connect to the KMS and request an encryption key.
- 3. The KMS sends an encryption key to the appliance. The new key from the KMS replaces the temporary KEK and is now used to encrypt and decrypt the DEKs for the appliance volumes.



Any data that exists before the encrypted appliance node connects to the configured KMS is encrypted with a temporary key. However, the appliance volumes should not be considered protected from removal from the data center until the temporary key is replaced by the KMS encryption key.

4. If the appliance is powered on or rebooted, it reconnects to the KMS to request the key. The key, which is saved in volatile memory, can't survive a loss of power or a reboot.

Considerations and requirements for using a key management server

Before configuring an external key management server (KMS), you must understand the considerations and requirements.

Which version of KMIP is supported?

StorageGRID supports KMIP version 1.4.

Key Management Interoperability Protocol Specification Version 1.4

What are the network considerations?

The network firewall settings must allow each appliance node to communicate through the port used for Key Management Interoperability Protocol (KMIP) communications. The default KMIP port is 5696.

You must ensure that each appliance node that uses node encryption has network access to the KMS or KMS cluster you configured for the site.

Which versions of TLS are supported?

Communications between the appliance nodes and the configured KMS use secure TLS connections. StorageGRID can support either the TLS 1.2 or TLS 1.3 protocol when it makes KMIP connections to a KMS or KMS cluster, based on what the KMS supports and which TLS and SSH policy you are using.

StorageGRID negotiates the protocol and cipher (TLS 1.2) or cipher suite (TLS 1.3) with the KMS when it makes the connection. To see which protocol versions and ciphers/cipher suites are available, review the tlsOutbound section of the grid's active TLS and SSH policy (**CONFIGURATION** > **Security Security settings**).

Which appliances are supported?

You can use a key management server (KMS) to manage encryption keys for any StorageGRID appliance in your grid that has the **Node Encryption** setting enabled. This setting can only be enabled during the hardware configuration stage of appliance installation using the StorageGRID Appliance Installer.



You can't enable node encryption after an appliance is added to the grid, and you can't use external key management for appliances that don't have node encryption enabled.

You can use the configured KMS for StorageGRID appliances and appliance nodes.

You can't use the configured KMS for software-based (non-appliance) nodes, including the following:

- Nodes deployed as virtual machines (VMs)
- · Nodes deployed within container engines on Linux hosts

Nodes deployed on these other platforms can use encryption outside of StorageGRID at the datastore or disk level.

When should I configure key management servers?

For a new installation, you should typically set up one or more key management servers in the Grid Manager before creating tenants. This order ensures that the nodes are protected before any object data is stored on them.

You can configure the key management servers in the Grid Manager before or after you install the appliance nodes.

How many key management servers do I need?

You can configure one or more external key management servers to provide encryption keys to the appliance nodes in your StorageGRID system. Each KMS provides a single encryption key to the StorageGRID appliance nodes at a single site or at a group of sites.

StorageGRID supports the use of KMS clusters. Each KMS cluster contains multiple, replicated key management servers that share configuration settings and encryption keys. Using KMS clusters for key management is recommended because it improves the failover capabilities of a high availability configuration.

For example, suppose your StorageGRID system has three data center sites. You might configure one KMS cluster to provide a key to all appliance nodes at Data Center 1 and a second KMS cluster to provide a key to all appliance nodes at all other sites. When you add the second KMS cluster, you can configure a default KMS for Data Center 2 and Data Center 3.

Note that you can't use a KMS for non-appliance nodes or for any appliance nodes that did not have the **Node Encryption** setting enabled during installation.



X

Appliance node with node encryption enabled

Appliance node without node encryption enabled

Non-appliance node (not encrypted)

What happens when a key is rotated?

As a security best practice, you should periodically rotate the encryption key used by each configured KMS.

When the new key version is available:

- It is automatically distributed to the encrypted appliance nodes at the site or sites associated with the KMS. The distribution should occur within an hour of when the key is rotated.
- If the encrypted appliance node is offline when the new key version is distributed, the node will receive the new key as soon as it reboots.
- If the new key version can't be used to encrypt appliance volumes for any reason, the **KMS encryption key rotation failed** alert is triggered for the appliance node. You might need to contact technical support for help in resolving this alert.

Can I reuse an appliance node after it has been encrypted?

If you need to install an encrypted appliance into another StorageGRID system, you must first decommission the grid node to move object data to another node. Then, you can use the StorageGRID Appliance Installer to clear the KMS configuration. Clearing the KMS configuration disables the **Node Encryption** setting and removes the association between the appliance node and the KMS configuration for the StorageGRID site.



With no access to the KMS encryption key, any data that remains on the appliance can no longer be accessed and is permanently locked.

Considerations for changing the KMS for a site

Each key management server (KMS) or KMS cluster provides an encryption key to all appliance nodes at a single site or at a group of sites. If you need to change which KMS is used for a site, you might need to copy the encryption key from one KMS to another.

If you change the KMS used for a site, you must ensure that the previously encrypted appliance nodes at that site can be decrypted using the key stored on the new KMS. In some cases, you might need to copy the current version of the encryption key from the original KMS to the new KMS. You must ensure that the KMS has the correct key to decrypt the encrypted appliance nodes at the site.

For example:

- 1. You initially configure a default KMS that applies to all sites that don't have a dedicated KMS.
- 2. When the KMS is saved, all appliance nodes that have the **Node Encryption** setting enabled connect to the KMS and request the encryption key. This key is used to encrypt the appliance nodes at all sites. This same key must also be used to decrypt those appliances.



3. You decide to add a site-specific KMS for one site (Data Center 3 in the figure). However, because the appliance nodes are already encrypted, a validation error occurs when you attempt to save the configuration for the site-specific KMS. The error occurs because the site-specific KMS does not have the correct key to decrypt the nodes at that site.



4. To address the issue, you copy the current version of the encryption key from the default KMS to the new KMS. (Technically, you copy the original key to a new key with the same alias. The original key becomes a prior version of the new key.) The site-specific KMS now has the correct key to decrypt the appliance nodes at Data Center 3, so it can be saved in StorageGRID.



Use cases for changing which KMS is used for a site

The table summarizes the required steps for the most common cases for changing the KMS for a site.

Use case for changing a site's KMS	Required steps
You have one or more site-specific KMS entries, and you want to use one of them as the default KMS.	Edit the site-specific KMS. In the Manages keys for field, select Sites not managed by another KMS (default KMS) . The site-specific KMS will now be used as the default KMS. It will apply to any sites that don't have a dedicated KMS. Edit a key management server (KMS)

Use case for changing a site's KMS	Required steps
You have a default KMS and you add a new site in an expansion. You don't want to use the default KMS for the new site.	 If the appliance nodes at the new site have already been encrypted by the default KMS, use the KMS software to copy the current version of the encryption key from the default KMS to a new KMS.
	2. Using the Grid Manager, add the new KMS and select the site. Add a key management server (KMS)
You want the KMS for a site to use a different server.	 If the appliance nodes at the site have already been encrypted by the existing KMS, use the KMS software to copy the current version of the encryption key from the existing KMS to the new KMS.
	2. Using the Grid Manager, edit the existing KMS configuration and enter the new host name or IP address.
	Add a key management server (KMS)

Configure StorageGRID as a client in the KMS

You must configure StorageGRID as a client for each external key management server or KMS cluster before you can add the KMS to StorageGRID.



These instructions apply to Thales CipherTrust Manager and Hashicorp Vault. For a list of supported products and versions, use the NetApp Interoperability Matrix Tool (IMT).

Steps

1. From the KMS software, create a StorageGRID client for each KMS or KMS cluster you plan to use.

Each KMS manages a single encryption key for the StorageGRID appliances nodes at a single site or at a group of sites.

- 2. Create a key using one of the following two methods:
 - Use the key management page of your KMS product. Create an AES encryption key for each KMS or KMS cluster.

The encryption key must be 2,048 bits or more, and it must be exportable.

- Have StorageGRID create the key. You will be prompted when you test and save after uploading client certificates.
- 3. Record the following information for each KMS or KMS cluster.

You need this information when you add the KMS to StorageGRID:

- Host name or IP address for each server.
- KMIP port used by the KMS.
- Key alias for the encryption key in the KMS.
- 4. For each KMS or KMS cluster, obtain a server certificate signed by a certificate authority (CA) or a certificate bundle that contains each of the PEM-encoded CA certificate files, concatenated in certificate

chain order.

The server certificate allows the external KMS to authenticate itself to StorageGRID.

- The certificate must use the Privacy Enhanced Mail (PEM) Base-64 encoded X.509 format.
- The Subject Alternative Name (SAN) field in each server certificate must include the fully qualified domain name (FQDN) or IP address that StorageGRID will connect to.



When you configure the KMS in StorageGRID, you must enter the same FQDNs or IP addresses in the **Hostname** field.

- The server certificate must match the certificate used by the KMIP interface of the KMS, which typically uses port 5696.
- 5. Obtain the public client certificate issued to StorageGRID by the external KMS and the private key for the client certificate.

The client certificate allows StorageGRID to authenticate itself to the KMS.

Add a key management server (KMS)

You use the StorageGRID Key Management Server wizard to add each KMS or KMS cluster.

Before you begin

- You have reviewed the considerations and requirements for using a key management server.
- You have configured StorageGRID as a client in the KMS, and you have the required information for each KMS or KMS cluster.
- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.

About this task

If possible, configure any site-specific key management servers before configuring a default KMS that applies to all sites not managed by another KMS. If you create the default KMS first, all node-encrypted appliances in the grid will be encrypted by the default KMS. If you want to create a site-specific KMS later, you must first copy the current version of the encryption key from the default KMS to the new KMS. See Considerations for changing the KMS for a site for details.

Step 1: KMS details

In Step 1 (KMS details) of the Add a Key Management Server wizard, you provide details about the KMS or KMS cluster.

Steps

1. Select CONFIGURATION > Security > Key management server.

The Key management server page appears with the Configuration details tab selected.

2. Select Create.

Step 1 (KMS details) of the Add a Key Management Server wizard appears.

3. Enter the following information for the KMS and the StorageGRID client you configured in that KMS.

Field	Description	
KMS name	A descriptive name to help you identify this KMS. Must be between 1 and 64 characters.	
Key name	The exact key alias for the StorageGRID client in the KMS. Must be between 1 and 255 characters. Note : If you haven't created a key using your KMS product, you'll be prompted to have StorageGRID create the key.	
Manages keys for	 The StorageGRID site that will be associated with this KMS. If possible, you should configure any site-specific key management servers before configuring a default KMS that applies to all sites not managed by another KMS. Select a site if this KMS will manage encryption keys for the 	
	 appliance nodes at a specific site. Select Sites not managed by another KMS (default KMS) to configure a default KMS that will apply to any sites that don't have a dedicated KMS and to any sites you add in subsequent expansions. 	
	Note: A validation error will occur when you save the KMS configuration if you select a site that was previously encrypted by the default KMS but you did not provide the current version of original encryption key to the new KMS.	
Port	The port the KMS server uses for Key Management Interoperability Protocol (KMIP) communications. Defaults to 5696, which is the KMIP standard port.	
Hostname	The fully qualified domain name or IP address for the KMS. Note: The Subject Alternative Name (SAN) field of the server certificate must include the FQDN or IP address you enter here. Otherwise, StorageGRID will not be able to connect to the KMS or to all servers in a KMS cluster.	

- 4. If you are configuring a KMS cluster, select **Add another hostname** to add a hostname for each server in the cluster.
- 5. Select Continue.

Step 2: Upload server certificate

In Step 2 (Upload server certificate) of the Add a Key Management Server wizard, you upload the server certificate (or certificate bundle) for the KMS. The server certificate allows the external KMS to authenticate itself to StorageGRID.

Steps

- 1. From **Step 2 (Upload server certificate)**, browse to the location of the saved server certificate or certificate bundle.
- 2. Upload the certificate file.

The server certificate metadata appears.



If you uploaded a certificate bundle, the metadata for each certificate appears on its own tab.

3. Select Continue.

Step 3: Upload client certificates

In Step 3 (Upload client certificates) of the Add a Key Management Server wizard, you upload the client certificate and the client certificate private key. The client certificate allows StorageGRID to authenticate itself to the KMS.

Steps

- 1. From Step 3 (Upload client certificates), browse to the location of the client certificate.
- 2. Upload the client certificate file.

The client certificate metadata appears.

- 3. Browse to the location of the private key for the client certificate.
- 4. Upload the private key file.
- 5. Select Test and save.

If a key doesn't exist, you are prompted to have StorageGRID create one.

The connections between the key management server and the appliance nodes are tested. If all connections are valid and the correct key is found on the KMS, the new key management server is added to the table on the Key Management Server page.



Immediately after you add a KMS, the certificate status on the Key Management Server page appears as Unknown. It might take StorageGRID as long as 30 minutes to get the actual status of each certificate. You must refresh your web browser to see the current status.

6. If an error message appears when you select **Test and save**, review the message details and then select **OK**.

For example, you might receive a 422: Unprocessable Entity error if a connection test failed.

7. If you need to save the current configuration without testing the external connection, select Force save.



Selecting **Force save** saves the KMS configuration, but it does not test the external connection from each appliance to that KMS. If there is an issue with the configuration, you might not be able to reboot appliance nodes that have node encryption enabled at the affected site. You might lose access to your data until the issues are resolved.

8. Review the confirmation warning, and select **OK** if you are sure you want to force save the configuration.

The KMS configuration is saved but the connection to the KMS is not tested.

Manage a KMS

Managing a key management server (KMS) involves viewing or editing details, managing certificates, viewing encrypted nodes, and removing a KMS when it is no longer needed.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- · You have the required access permission.

View KMS details

You can view information about each key management server (KMS) in your StorageGRID system, including key details and the current status of the server and client certificates.

Steps

1. Select CONFIGURATION > Security > Key management server.

The Key management server page appears and shows the following information:

- The Configuration details tab lists any key management servers that are configured.
- The Encrypted nodes tab lists any nodes that have node encryption enabled.
- 2. To view the details for a specific KMS and perform operations on that KMS, select the name of the KMS. The details page for the KMS lists the following information:

Field	Description
Manages keys for	The StorageGRID site associated with the KMS. This field displays the name of a specific StorageGRID site or Sites not managed by another KMS (default KMS).
Hostname	The fully qualified domain name or IP address of the KMS. If there is a cluster of two key management servers, the fully qualified domain name or IP address of both servers are listed. If there are more than two key management servers in a cluster, the fully qualified domain name or IP address of the first KMS is listed along with the number of additional key management servers in the cluster. For example: 10.10.10.10 and 10.10.10.11 or 10.10.10.10 and 2 others. To view all hostnames in a cluster, select a KMS and select Edit or Actions > Edit.

3. Select a tab on the KMS details page to view the following information:

Tab	Field	Description
Key details	Key name	The key alias for the StorageGRID client in the KMS.
	Key UID	The unique identifier of the latest version of the key.
	Last modified	The date and time of the latest version of the key.
Server certificate	Metadata	The metadata for the certificate, such as serial number, expiration date and time, and the certificate PEM.
	Certificate PEM	The contents of the PEM (privacy enhanced mail) file for the certificate.
Client certificate	Metadata	The metadata for the certificate, such as serial number, expiration date and time, and the certificate PEM.
	Certificate PEM	The contents of the PEM (privacy enhanced mail) file for the certificate.

4. As often as required by your organization's security practices, select **Rotate key**, or use the KMS software, to create a new version of the key.

When key rotation is successful, the Key UID and Last modified fields are updated.

If you rotate the encryption key using the KMS software, rotate it from the last used version of the key to a new version of the same key. Don't rotate to an entirely different key.



Never attempt to rotate a key by changing the key name (alias) for the KMS. StorageGRID requires all previously used key versions (as well as any future ones) to be accessible from the KMS with the same key alias. If you change the key alias for a configured KMS, StorageGRID might not be able to decrypt your data.

Manage certificates

Promptly address any server or client certificate issues. If possible, replace certificates before they expire.



You must address any certificate issues as soon as possible to maintain data access.

Steps

1. Select **CONFIGURATION > Security > Key management server**.

- 2. In the table, look at the value for Certificate expiration for each KMS.
- 3. If Certificate expiration for any KMS is Unknown, wait up to 30 minutes and then refresh your web browser.
- If the Certificate expiration column indicates that a certificate has expired or is nearing expiration, select the KMS to go to the KMS details page.
 - a. Select Server certificate and verify the value for the "Expires on" field.
 - b. To replace the certificate, select **Edit certificate** to upload a new certificate.

- c. Repeat these sub-steps and select **Client certificate** instead of Server certificate.
- 5. When the **KMS CA certificate expiration**, **KMS client certificate expiration**, and **KMS server certificate expiration** alerts are triggered, note the description of each alert and perform the recommended actions.



It might take StorageGRID as long as 30 minutes to get updates to the certificate expiration. Refresh your web browser to see the current values.

View encrypted nodes

You can view information about the appliance nodes in your StorageGRID system that have the **Node Encryption** setting enabled.

Steps

1. Select CONFIGURATION > Security > Key management server.

The Key Management Server page appears. The Configuration Details tab shows any key management servers that have been configured.

2. From the top of the page, select the **Encrypted nodes** tab.

The Encrypted nodes tab lists the appliance nodes in your StorageGRID system that have the **Node Encryption** setting enabled.

3. Review the information in the table for each appliance node.

Column	Description
Node name	The name of the appliance node.
Node type	The type of node: Storage, Admin, or Gateway.
Site	The name of the StorageGRID site where the node is installed.
KMS name	The descriptive name of the KMS used for the node. If no KMS is listed, select the Configuration details tab to add a KMS. Add a key management server (KMS)
Key UID	The unique ID of the encryption key used to encrypt and decrypt data on the appliance node. To view an entire key UID, select the text. A dash () indicates the key UID is unknown, possibly because of a connection issue between the appliance node and the KMS.

Column	Description
Status	The status of the connection between the KMS and the appliance node. If the node is connected, the timestamp updates every 30 minutes. It can take several minutes for the connection status to update after the KMS configuration changes. Note: Refresh your web browser to see the new values.

4. If the Status column indicates a KMS issue, address the issue immediately.

During normal KMS operations, the status will be **Connected to KMS**. If a node is disconnected from the grid, the node connection state is shown (Administratively Down or Unknown).

Other status messages correspond to StorageGRID alerts with the same names:

- KMS configuration failed to load
- · KMS connectivity error
- · KMS encryption key name not found
- KMS encryption key rotation failed
- KMS key failed to decrypt an appliance volume
- KMS is not configured

Perform the recommended actions for these alerts.



You must address any issues immediately to ensure that your data is fully protected.

Edit a KMS

You might need to edit the configuration of a key management server, for example, if a certificate is about to expire.

Before you begin

- If you plan to update the site selected for a KMS, you have reviewed the considerations for changing the KMS for a site.
- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.

Steps

1. Select CONFIGURATION > Security > Key management server.

The Key management server page appears and shows all key management servers that have been configured.

2. Select the KMS you want to edit, and select Actions > Edit.

You can also edit a KMS by selecting the KMS name in the table and selecting **Edit** on the KMS details page.

3. Optionally, update the details in Step 1 (KMS details) of the Edit a Key Management Server wizard.

Field	Description
KMS name	A descriptive name to help you identify this KMS. Must be between 1 and 64 characters.
Key name	The exact key alias for the StorageGRID client in the KMS. Must be between 1 and 255 characters. You only need to edit the key name in rare cases. For example, you must edit the key name if the alias is renamed in the KMS or if all versions of the previous key have been copied to the version history of the new alias.
Manages keys for	If you are editing a site-specific KMS and you don't already have a default KMS, optionally select Sites not managed by another KMS (default KMS) . This selection converts a site-specific KMS to the default KMS, which will apply to all sites that don't have a dedicated KMS and to any sites added in an expansion. Note: If you are editing a site-specific KMS, you can't select another site. If you are editing the default KMS, you can't select a specific site.
Port	The port the KMS server uses for Key Management Interoperability Protocol (KMIP) communications. Defaults to 5696, which is the KMIP standard port.
Hostname	The fully qualified domain name or IP address for the KMS. Note: The Subject Alternative Name (SAN) field of the server certificate must include the FQDN or IP address you enter here. Otherwise, StorageGRID will not be able to connect to the KMS or to all servers in a KMS cluster.

- 4. If you are configuring a KMS cluster, select **Add another hostname** to add a hostname for each server in the cluster.
- 5. Select Continue.

Step 2 (Upload server certificate) of the Edit a Key Management Server wizard appears.

- 6. If you need to replace the server certificate, select **Browse** and upload the new file.
- 7. Select Continue.

Step 3 (Upload client certificates) of the Edit a Key Management Server wizard appears.

- 8. If you need to replace the client certificate and the client certificate private key, select **Browse** and upload the new files.
- 9. Select Test and save.

The connections between the key management server and all node-encrypted appliance nodes at the affected sites are tested. If all node connections are valid and the correct key is found on the KMS, the key management server is added to the table on the Key Management Server page.

10. If an error message appears, review the message details, and select **OK**.

For example, you might receive a 422: Unprocessable Entity error if the site you selected for this KMS is already managed by another KMS, or if a connection test failed.

11. If you need to save the current configuration before resolving the connection errors, select **Force save**.



Selecting **Force save** saves the KMS configuration, but it does not test the external connection from each appliance to that KMS. If there is an issue with the configuration, you might not be able to reboot appliance nodes that have node encryption enabled at the affected site. You might lose access to your data until the issues are resolved.

The KMS configuration is saved.

12. Review the confirmation warning, and select **OK** if you are sure you want to force save the configuration.

The KMS configuration is saved, but the connection to the KMS is not tested.

Remove a key management server (KMS)

You might want to remove a key management server in some cases. For example, you might want to remove a site-specific KMS if you have decommissioned the site.

Before you begin

- You have reviewed the considerations and requirements for using a key management server.
- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.

About this task

You can remove a KMS in these cases:

- You can remove a site-specific KMS if the site has been decommissioned or if the site includes no appliance nodes with node encryption enabled.
- You can remove the default KMS if a site-specific KMS already exists for each site that has appliance nodes with node encryption enabled.

Steps

1. Select CONFIGURATION > Security > Key management server.

The Key management server page appears and shows all key management servers that have been configured.

2. Select the KMS you want to remove, and select **Actions > Remove**.

You can also remove a KMS by selecting the KMS name in the table and selecting **Remove** from the KMS details page.

- 3. Confirm the following is true:
 - You are removing a site-specific KMS for a site that has no appliance node with node encryption enabled.
 - You are removing the default KMS, but a site-specific KMS already exists for each site with node encryption.
- 4. Select Yes.

The KMS configuration is removed.

Manage proxy settings

Configure storage proxy

If you are using platform services or Cloud Storage Pools, you can configure a nontransparent proxy between Storage Nodes and the external S3 endpoints. For example, you might need a non-transparent proxy to allow platform services messages to be sent to external endpoints, such as an endpoint on the internet.



Configured storage proxy settings do not apply to Kafka platform services endpoints.

Before you begin

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

About this task

You can configure the settings for a single storage proxy.

Steps

- 1. Select CONFIGURATION > Security > Proxy settings.
- 2. On the Storage tab, select the Enable storage proxy checkbox.
- 3. Select the protocol for the storage proxy.
- 4. Enter the hostname or IP address of the proxy server.
- 5. Optionally, enter the port used to connect to the proxy server.

Leave this field blank to use the default port for the protocol: 80 for HTTP or 1080 for SOCKS5.

6. Select Save.

After the storage proxy is saved, new endpoints for platform services or Cloud Storage Pools can be configured and tested.



Proxy changes can take up to 10 minutes to take effect.

- 7. Check the settings of your proxy server to ensure that platform service-related messages from StorageGRID will not be blocked.
- 8. If you need to disable a storage proxy, clear the checkbox, and select **Save**.

Configure admin proxy settings

If you send AutoSupport packages using HTTP or HTTPS, you can configure a nontransparent proxy server between Admin Nodes and technical support (AutoSupport).

For more information about AutoSupport, see Configure AutoSupport.

Before you begin

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

About this task

You can configure the settings for a single admin proxy.

Steps

1. Select CONFIGURATION > Security > Proxy settings.

The Proxy Settings page appears. By default, Storage is selected in the tab menu.

- 2. Select the **Admin** tab.
- 3. Select the Enable Admin Proxy checkbox.
- 4. Enter the hostname or IP address of the proxy server.
- 5. Enter the port used to connect to the proxy server.
- 6. Optionally, enter a username and password for the proxy server.

Leave these fields blank if your proxy server does not require a username or a password.

- 7. Select one of the following:
 - If you want to secure the connection to the admin proxy, select **Verify certificate**. Upload a CA bundle to verify the authenticity of SSL certificates presented by the admin proxy server.



AutoSupport on Demand, E-Series AutoSupport through StorageGRID, and Update Path determination on the StorageGRID Upgrade page will not work if a proxy certificate is verified.

After you upload the CA bundle, its metadata appears.

- If you don't want to validate certificates when communicating with the admin proxy server, select **Do** not verify certificate.
- 8. Select Save.

After the admin proxy is saved, the proxy server between Admin Nodes and technical support is configured.



Proxy changes can take up to 10 minutes to take effect.

9. If you need to disable the admin proxy, clear the Enable Admin Proxy checkbox, and then select Save.

Control firewalls

Control access at external firewall

You can open or close specific ports at the external firewall.

You can control access to the user interfaces and APIs on StorageGRID Admin Nodes by opening or closing specific ports at the external firewall. For example, you might want to prevent tenants from being able to connect to the Grid Manager at the firewall, in addition to using other methods to control system access.

Port	Description	If port is open
443	Default HTTPS port for Admin Nodes	Web browsers and management API clients can access the Grid Manager, the Grid Management API, the Tenant Manager, and the Tenant Management API. Note: Port 443 is also used for some internal traffic.
8443	Restricted Grid Manager port on Admin Nodes	 Web browsers and management API clients can access the Grid Manager and the Grid Management API using HTTPS. Web browsers and management API clients can't access the Tenant Manager or the Tenant Management API. Requests for internal content will be rejected.
9443	Restricted Tenant Manager port on Admin Nodes	 Web browsers and management API clients can access the Tenant Manager and the Tenant Management API using HTTPS. Web browsers and management API clients can't access the Grid Manager or the Grid Management API. Requests for internal content will be rejected.



Single sign-on (SSO) is not available on the restricted Grid Manager or Tenant Manager ports. You must use the default HTTPS port (443) if you want users to authenticate with single sign-on.

Related information

- Sign in to the Grid Manager
- Create tenant account
- External communications

Manage internal firewall controls

StorageGRID includes an internal firewall on each node that enhances the security of your grid by enabling you to control network access to the node. Use the firewall to prevent network access on all ports except those necessary for your specific grid deployment. The configuration changes you make on the Firewall control page are deployed to each node.

Use the three tabs on the Firewall control page to customize the access you need for your grid.

- **Privileged address list**: Use this tab to allow selected access to closed ports. You can add IP addresses or subnets in CIDR notation that can access ports closed using the Manage external access tab.
- Manage external access: Use this tab to close ports that are open by default, or reopen ports previously

closed.

• **Untrusted Client Network**: Use this tab to specify whether a node trusts inbound traffic from the Client Network.

The settings on this tab override the settings in the Manage external access tab.

- A node with an untrusted Client Network will accept only connections on load balancer endpoint ports configured on that node (global, node interface and node type bound endpoints).
- Load balancer endpoint ports *are the only open ports* on untrusted Client Networks, regardless of the settings on the Manage external networks tab.
- When trusted, all ports opened under the Manage external access tab are accessible, as well as any load balancer endpoints opened on the Client Network.



The settings you make on one tab can affect the access changes you make on another tab. Be sure to check the settings on all tabs to ensure your network behaves in the way you expect.

To configure internal firewall controls, see Configure firewall controls.

For more information about external firewalls and network security, see Control access at external firewall.

Privileged address list and Manage external access tabs

The Privileged address list tab enables you to register one or more IP addresses that are granted access to grid ports that are closed. The Manage external access tab enables you to close external access to selected external ports or all open external ports (external ports are ports that are accessible by non-grid nodes by default). These two tabs often can be used together to customize the exact network access you need to allow for your grid.



Privileged IP addresses don't have internal grid port access by default.

Example 1: Use a jump host for maintenance tasks

Suppose you want to use a jump host (a security hardened host) for network administration. You could use these general steps:

- 1. Use the Privileged address list tab to add the IP address of the jump host.
- 2. Use the Manage external access tab to block all ports.



Add the privileged IP address before you block ports 443 and 8443. Any users currently connected on a blocked port, including you, will lose access to Grid Manager unless their IP address has been added to the Privileged address list.

After you save your configuration, all external ports on the Admin Node in your grid will be blocked for all hosts except the jump host. You can then use the jump host to perform maintenance tasks on your grid more securely.

Example 2: Limit access to the Grid Manager and Tenant Manager

Suppose you want to limit access to the Grid Manager and Tenant manager (preset ports) for security reasons. You could use these general steps:

- 1. Use the toggle on the Manage external access tab to block port 443.
- 2. Use the toggle on the Manage external access tab to allow access to port 8443.
- 3. Use the toggle on the Manage external access tab to allow access to port 9443.

After you save your configuration, hosts will not be able to access port 443, but they can still access the Grid Manager through port 8443 and the Tenant Manager through port 9443.



Ports 443, 8443, and 9443 are the preset ports for Grid Manager and Tenant Manager. You can toggle any port to limit access to a specific Grid Manager or Tenant manager.

Example 3: Lock down sensitive ports

Suppose you want to lock down sensitive ports and the service on that port (for example, SSH on port 22). You could use the following general steps:

- 1. Use the Privileged address list tab to grant access only to the hosts that need access to the service.
- 2. Use the Manage external access tab to block all ports.



Add the privileged IP address before you block access to any ports assigned to access Grid Manager and Tenant manager (preset ports are 443 and 8443). Any users currently connected on a blocked port, including you, will lose access to Grid Manager unless their IP address has been added to the Privileged address list.

After you save your configuration, port 22 and SSH service will be available to hosts on the privileged address list. All other hosts will be denied access to the service no matter what interface the request comes from.

Example 4: Disable access to unused services

At a network level, you could disable some services that you don't intend to use. For example if you will not provide Swift access, you would perform the following general steps:

- 1. Use the toggle on the Manage external access tab to block port 18083.
- 2. Use the toggle on the Manage external access tab to block port 18085.

After you save your configuration, the Storage Node no longer allows Swift connectivity, but continues to allow access to other services on unblocked ports.

Untrusted Client Networks tab

If you are using a Client Network, you can help secure StorageGRID from hostile attacks by accepting inbound client traffic only on explicitly configured endpoints.

By default, the Client Network on each grid node is *trusted*. That is, by default, StorageGRID trusts inbound connections to each grid node on all available external ports.

You can reduce the threat of hostile attacks on your StorageGRID system by specifying that the Client Network on each node be *untrusted*. If a node's Client Network is untrusted, the node only accepts inbound connections on ports explicitly configured as load balancer endpoints. See Configure load balancer endpoints and Configure firewall controls.

Example 1: Gateway Node only accepts HTTPS S3 requests

Suppose you want a Gateway Node to refuse all inbound traffic on the Client Network except for HTTPS S3 requests. You would perform these general steps:

- 1. From the Load balancer endpoints page, configure a load balancer endpoint for S3 over HTTPS on port 443.
- 2. From the Firewall control page, select Untrusted to specify that the Client Network on the Gateway Node is untrusted.

After you save your configuration, all inbound traffic on the Gateway Node's Client Network is dropped except for HTTPS S3 requests on port 443 and ICMP echo (ping) requests.

Example 2: Storage Node sends S3 platform services requests

Suppose you want to enable outbound S3 platform services traffic from a Storage Node, but you want to prevent any inbound connections to that Storage Node on the Client Network. You would perform this general step:

• From the Untrusted Client Networks tab of the Firewall control page, indicate that the Client Network on the Storage Node is untrusted.

After you save your configuration, the Storage Node no longer accepts any incoming traffic on the Client Network, but it continues to allow outbound requests to configured platform services destinations.

Example 3: Limiting access to Grid Manager to a subnet

Suppose you want to allow Grid Manager access only on a specific subnet. You would perform the following steps:

- 1. Attach the Client Network of your Admin Nodes to the subnet.
- 2. Use the Untrusted Client Network tab to configure the Client Network as untrusted.
- 3. When you create a management interface load balancer endpoint, enter port and select the management interface that the port will access.
- 4. Select Yes for Untrusted Client Network.
- 5. Use the Manage external access tab to block all external ports (with or without privileged IP addresses set for hosts outside that subnet).

After you save your configuration, only hosts on the subnet you specified can access the Grid Manager. All other hosts are are blocked.

Configure internal firewall

You can configure the StorageGRID firewall to control network access to specific ports on your StorageGRID nodes.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.
- You have reviewed the information in Manage firewall controls and Networking guidelines.

• If you want an Admin Node or Gateway Node to accept inbound traffic only on explicitly configured endpoints, you have defined the load balancer endpoints.



When changing the configuration of the Client Network, existing client connections might fail if load balancer endpoints have not been configured.

About this task

StorageGRID includes an internal firewall on each node that enables you to open or close some of the ports on the nodes of your grid. You can use the Firewall control tabs to open or close ports that are open by default on the Grid Network, Admin Network, and Client Network. You can also create a list of privileged IP addresses that can access grid ports that are closed. If you are using a Client Network, you can specify whether a node trusts inbound traffic from the Client Network, and you can configure the access of specific ports on the Client Network.

Limiting the number of ports open to IP addresses outside of your grid to only those that are absolutely necessary enhances the security of your grid. You use the settings on each of the three Firewall control tabs to ensure only the needed ports are open.

For more information about using firewall controls, including examples, see Manage firewall controls.

For more information about external firewalls and network security, see Control access at external firewall.

Access firewall controls

Steps

1. Select CONFIGURATION > Security > Firewall control.

The three tabs on this page are described in Manage firewall controls.

2. Select any tab to configure the firewall controls.

You can use these tabs in any order. The configurations you set on one tab don't limit what you can do on the other tabs; however, configuration changes you make on one tab might change the behavior of ports configured on other tabs.

Privileged address list

You use the Privileged address list tab to grant hosts access to ports that are closed by default or closed by settings on the Manage external access tab.

Privileged IP addresses and subnets don't have internal grid access by default. Also, load balancer endpoints and additional ports opened in the Privileged address list tab are accessible even if blocked in the Manage external access tab.



Settings on the Privileged address list tab can't override settings on the Untrusted Client Network tab.

Steps

- 1. On the Privileged address list tab, enter the address or IP subnet you want to grant access to closed ports.
- 2. Optionally, select **Add another IP address or subnet in CIDR notation** to add additional privileged clients.



Add as few addresses as possible to the privileged list.

3. Optionally, select Allow privileged IP addresses to access StorageGRID internal ports. See StorageGRID internal ports.



This option removes some protections for internal services. Leave it disabled if possible.

4. Select Save.

Manage external access

When a port is closed in the Manage external access tab, the port can't be accessed by any non-grid IP address unless you add the IP address to the privileged address list. You can only close ports that are open by default, and you can only open ports that you have closed.



Settings on the Manage external access tab can't override settings on the Untrusted Client Network tab. For example, if a node is untrusted, port SSH/22 is blocked on the Client Network even if it is open on the Manage external access tab. Settings on the Untrusted Client Network tab override closed ports (such as 443, 8443, 9443) on the Client Network.

Steps

1. Select Manage external access.

The tab displays a table with all of the external ports (ports that are accessible by non-grid nodes by default) for the nodes in your grid.

- 2. Configure the ports you want open and closed using the following options:
 - Use the toggle beside each port to open or close the selected port.
 - Select Open all displayed ports to open all ports listed in the table.
 - Select Close all displayed ports to close all ports listed in the table.



If you close Grid Manager ports 443 or 8443, any users currently connected on a blocked port, including you, will lose access to Grid Manager unless their IP address has been added to the Privileged address list.



Use the scroll bar on the right side of the table to be sure you have viewed all available ports. Use the search field to find the settings for any external port by entering a port number. You can enter a partial port number. For example, if you enter a **2**, all ports that have the string "2" as part of their name are displayed.

3. Select Save

Untrusted Client Network

If the Client Network for a node is untrusted, the node only accepts inbound traffic on ports configured as load balancer endpoints and, optionally, additional ports you select on this tab. You can also use this tab to specify the default setting for new nodes added in an expansion.



Existing client connections might fail if load balancer endpoints have not been configured.

The configuration changes you make on the Untrusted Client Network tab override the settings on the

Manage external access tab.

Steps

- 1. Select Untrusted Client Network.
- 2. In the Set New Node Default section, specify what the default setting should be when new nodes are added to the grid in an expansion procedure.
 - Trusted (default): When a node is added in an expansion, its Client Network is trusted.
 - Untrusted: When a node is added in an expansion, its Client Network is untrusted.

As required, you can return to this tab to change the setting for a specific new node.



This setting does not affect the existing nodes in your StorageGRID system.

- 3. Use the following options to select the nodes that should allow client connections only on explicitly configured load balancer endpoints or additional selected ports:
 - Select Untrust on displayed nodes to add all nodes displayed in the table to the Untrusted Client Network list.
 - Select Trust on displayed nodes to remove all nodes displayed in the table from the Untrusted Client Network list.
 - Use the toggle beside each node to set the Client Network as Trusted or Untrusted for the selected node.

For example, you could select **Untrust on displayed nodes** to add all nodes to the Untrusted Client Network list and then use the toggle besides an individual node to add that single node to the Trusted Client Network list.



Use the scroll bar on the right side of the table to be sure you have viewed all available nodes. Use the search field to find the settings for any node by entering the node name. You can enter a partial name. For example, if you enter a **GW**, all nodes that have the string "GW" as part of their name are displayed.

4. Select Save.

The new firewall settings are immediately applied and enforced. Existing client connections might fail if load balancer endpoints have not been configured.

Manage tenants

Manage tenants: Overview

As a grid administrator, you create and manage the tenant accounts that S3 and Swift clients use to store and retrieve objects.



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

What are tenant accounts?

A tenant account allows you to use either the Simple Storage Service (S3) REST API or the Swift REST API to

store and retrieve objects in a StorageGRID system.

Each tenant account has federated or local groups, users, S3 buckets or Swift containers, and objects.

Tenant accounts can be used to segregate stored objects by different entities. For example, multiple tenant accounts can be used for either of these use cases:

• Enterprise use case: If you are administering a StorageGRID system in an enterprise application, you might want to segregate the grid's object storage by the different departments in your organization. In this case, you could create tenant accounts for the Marketing department, the Customer Support department, the Human Resources department, and so on.



If you use the S3 client protocol, you can use S3 buckets and bucket policies to segregate objects between the departments in an enterprise. You don't need to use tenant accounts. See instructions for implementing S3 buckets and bucket policies for more information.

• Service provider use case: If you are administering a StorageGRID system as a service provider, you can segregate the grid's object storage by the different entities that will lease the storage on your grid. In this case, you would create tenant accounts for Company A, Company B, Company C, and so on.

For more information, see Use a tenant account.

How do I create a tenant account?

When you create a tenant account, you specify the following information:

- Basic information including the tenant name, client type (S3 or Swift) and optional storage quota.
- Permissions for the tenant account, such as whether the tenant account can use S3 platform services, configure its own identity source, use S3 Select, or use a grid federation connection.
- The initial root access for the tenant, based on whether the StorageGRID system uses local groups and users, identity federation, or single sign-on (SSO).

In addition, you can enable the S3 Object Lock setting for the StorageGRID system if S3 tenant accounts need to comply with regulatory requirements. When S3 Object Lock is enabled, all S3 tenant accounts can create and manage compliant buckets.

What is Tenant Manager used for?

After you create the tenant account, tenant users can sign in to the Tenant Manager to perform tasks such as the following:

- Set up identity federation (unless the identity source is shared with the grid)
- · Manage groups and users
- Use grid federation for account clone and cross-grid replication
- Manage S3 access keys
- Create and manage S3 buckets
- Use S3 platform services
- Use S3 Select
- Monitor storage usage


While S3 tenant users can create and manage S3 access key and buckets with the Tenant Manager, they must use an S3 client application to ingest and manage objects. See Use S3 REST API for details.



Swift users must have the Root access permission to access the Tenant Manager. However, the Root access permission does not allow users to authenticate into the Swift REST API to create containers and ingest objects. Users must have the Swift Administrator permission to authenticate into the Swift REST API.

Create a tenant account

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

The steps for creating a tenant account vary based on whether identity federation and single sign-on are configured and whether the Grid Manager account you use to create the tenant account belongs to an admin group with the Root access permission.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- · You have the Root access or Tenant accounts permission.
- If the tenant account will use the identity source that was configured for the Grid Manager, and you want to grant Root access permission for the tenant account to a federated group, you have imported that federated group into the Grid Manager. You don't need to assign any Grid Manager permissions to this admin group. See Manage admin groups.
- If you want to allow an S3 tenant to clone account data and replicate bucket objects to another grid using a grid federation connection:
 - You have configured the grid federation connection.
 - The status of the connection is **Connected**.
 - · You have Root access permission.
 - You have reviewed the considerations for managing the permitted tenants for grid federation.
 - If the tenant account will use the identity source that was configured for Grid Manager, you have imported the same federated group into Grid Manager on both grids.

When you create the tenant, you will select this group to have the initial Root access permission for both the source and destination tenant accounts.



If this admin group doesn't exist on both grids before you create the tenant, the tenant isn't replicated to the destination.

Access the wizard

Steps

- 1. Select TENANTS.
- 2. Select Create.

Enter details

Steps

1. Enter details for the tenant.

Field	Description
Name	A name for the tenant account. Tenant names don't need to be unique. When the tenant account is created, it receives a unique, 20-digit account ID.
Description (optional)	A description to help identify the tenant. If you are creating a tenant that will use a grid federation connection, optionally, use this field to help identify which is the source tenant and which is the destination tenant. For example, this description for a tenant created on Grid 1 will also appear for the tenant replicated to Grid 2: "This tenant was created on Grid 1."
Client type	The type of client protocol this tenant will use, either S3 or Swift . Note : Support for Swift client applications has been deprecated and will be removed in a future release.
Storage quota (optional)	If you want this tenant to have a storage quota, a numerical value for the quota and the units.

2. Select Continue.

Select permissions

Steps

1. Optionally, select any permissions you want this tenant to have.



Some of these permissions have additional requirements. For details, select the help icon for each permission.

Permission	If selected
Allow platform services	The tenant can use S3 platform services such as CloudMirror. See Manage platform services for S3 tenant accounts.
Use own identity source	The tenant can configure and manage its own identity source for federated groups and users. This option is disabled if you have configured SSO for your StorageGRID system.

Permission	If selected					
Allow S3 Select	 The tenant can issue S3 SelectObjectContent API requests to filter and retrieve object data. See Manage S3 Select for tenant accounts. Important: SelectObjectContent requests can decrease load-balancer performance for all S3 clients and all tenants. Enable this feature only when required and only for trusted tenants. 					
Use grid federation connection	 The tenant can use a grid federation connection. Selecting this option: Causes this tenant and all tenant groups and users added to the account to be cloned from this grid (the <i>source grid</i>) to the other grid in the selected connection (the <i>destination grid</i>). Allows this tenant to configure cross-grid replication between corresponding buckets on each grid. See Manage the permitted tenants for grid federation. 					

2. If you selected **Use grid federation connection**, select one of the available grid federation connections.

Jse grid federation connection 💡		
Connection name 🛿 ≑	Remote grid hostname 👔 ≑	Connection status 👔 ≑
Grid A-Grid B	10.96.104.230	Connected

3. Select Continue.

Define root access and create tenant

Steps

1. Define root access for the tenant account, based on whether your StorageGRID system uses identity federation, single sign-on (SSO), or both.

Do this			
y the password to use when signing into the tenant as the local ser.			
lect an existing federated group to have Root access rmission for the tenant. tionally, specify the password to use when signing in to the nant as the local root user.			

Option	Do this
If both identity federation and single sign-on (SSO) are enabled	Select an existing federated group to have Root access permission for the tenant. No local users can sign in.

2. Select Create tenant.

A success message appears, and the new tenant is listed on the Tenants page. To learn how to view tenant details and monitor tenant activity, see Monitor tenant activity.

- 3. If you selected the **Use grid federation connection** permission for the tenant:
 - a. Confirm that an identical tenant was replicated to the other grid in the connection. The tenants on both grids will have the same 20-digit account ID, name, description, quota, and permissions.



If you see the error message "Tenant created without a clone," refer to the instructions in Troubleshoot grid federation errors.

b. If you provided a local root user password when defining root access, change the password for the local root user for the replicated tenant.



A local root user can't sign in to Tenant Manager on the destination grid until the password is changed.

Sign in to tenant (optional)

As required, you can sign in to the new tenant now to complete the configuration, or you can sign in to the tenant later. The sign-in steps depend on whether you are signed in to the Grid Manager using the default port (443) or a restricted port. See Control access at external firewall.

Sign in now

If you are using	Do this
Port 443 and you set a password for the local root user	 Select Sign in as root. When you sign in, links appear for configuring buckets, identity federation, groups, and users. Select the links to configure the tenant account. Each link opens the corresponding page in the Tenant Manager. To complete the page, see the instructions for using tenant accounts.
Port 443 and you did not set a password for the local root user	Select Sign in , and enter the credentials for a user in the Root access federated group.

If you are using	Do this
A restricted port	1. Select Finish
	2. Select Restricted in the Tenant table to learn more about accessing this tenant account.
	The URL for the Tenant Manager has this format:
	https://FQDN_or_Admin_Node_IP:port/?accountId=20-digit- account-id/
	 FQDN_or_Admin_Node_IP is a fully qualified domain name or the IP address of an Admin Node
	° port is the tenant-only port
	° 20-digit-account-id is the tenant's unique account ID

Sign in later

If you are using	Do one of these
Port 443	 From the Grid Manager, select TENANTS, and select Sign in to the right of the tenant name.
	• Enter the tenant's URL in a web browser:
	https://FQDN_or_Admin_Node_IP/?accountId=20-digit- account-id/
	 FQDN_or_Admin_Node_IP is a fully qualified domain name or the IP address of an Admin Node
	° 20-digit-account-id is the tenant's unique account ID
A restricted port	• From the Grid Manager, select TENANTS , and select Restricted .
	• Enter the tenant's URL in a web browser:
	https://FQDN_or_Admin_Node_IP:port/?accountId=20-digit- account-id
	 FQDN_or_Admin_Node_IP is a fully qualified domain name or the IP address of an Admin Node
	 port is the tenant-only restricted port
	° 20-digit-account-id is the tenant's unique account ID

Configure the tenant

Follow the instructions in Use a tenant account to manage tenant groups and users, S3 access keys, buckets, platform services, and account clone and cross-grid replication.

Edit tenant account

You can edit a tenant account to change the display name, storage quota, or tenant permissions.



If a tenant has the **Use grid federation connection** permission, you can edit tenant details from either grid in the connection. However, any changes you make on one grid in the connection will not be copied to the other grid. If you want to keep the tenant details exactly in sync between grids, make the same edits on both grids. See Manage the permitted tenants for grid federation connection.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access or Tenant accounts permission.

Steps

1. Select TENANTS.

Tenants							
View informat To view more	tion for each te recent values,	enant account. Depending or select the tenant name.	n the timing of ingests, network	< connectiv	vity, and node sta	tus, the usage data sho	wn might be out of date.
Create	xport to CSV	Actions 🛩 Search ten	ants by name or ID		Q		Displaying 5 results
Na	ame 😧 💠	Logical space used 👔 💠	Quota utilization 🧿 ≑		Quota 👔 💠	Object count 👔 💠	Sign in/Copy URL 👔
Te	enant 01	2.00 GB		10%	20.00 GB	100	→ □
Te	enant 02	85.00 GB		85%	100.00 GB	500	→ □
Te	enant 03	500.00 TB		50%	1.00 PB	10,000	→ □
Te	enant 04	475.00 TB		95%	500.00 TB	50,000	→ □
Te	enant 05	5.00 GB			E	500	→ □

2. Locate the tenant account you want to edit.

Use the search box to search for a tenant by name or tenant ID.

- 3. Select the tenant. You can do either of the following:
 - Select the checkbox for the tenant, and select Actions > Edit.
 - $\circ\,$ Select the tenant name to display the details page, and select Edit.
- 4. Optionally, change the values for these fields:
 - Name
 - Description
 - Storage quota

5. Select Continue.

- 6. Select or clear the permissions for the tenant account.
 - If you disable Platform services for a tenant who is already using them, the services that they have configured for their S3 buckets will stop working. No error message is sent to the tenant. For example, if the tenant has configured CloudMirror replication for an S3 bucket, they can still store objects in the bucket, but copies of those objects will no longer be made in the external S3 bucket that they have configured as an endpoint. See Manage platform services for S3 tenant accounts.
 - Change the setting of **Uses own identity source** to determine whether the tenant account will use its own identity source or the identity source that was configured for the Grid Manager.

If Uses own identity source is:

- Disabled and selected, the tenant has already enabled its own identity source. A tenant must disable its identity source before it can use the identity source that was configured for the Grid Manager.
- Disabled and not selected, SSO is enabled for the StorageGRID system. The tenant must use the identity source that was configured for the Grid Manager.
- Select or clear the Allow S3 Select permission as needed. See Manage S3 Select for tenant accounts.
- To remove the **Use grid federation connection** permission:
 - a. Go to the tenant's details page.
 - b. Select the Grid federation tab.
 - c. Select Remove permission.
- To add the Use grid federation connection permission:
 - a. Select the **Use grid federation connection** checkbox.
 - b. Optionally, select **Clone existing local users and groups** to clone them to the remote grid. If you want, you can stop the cloning in progress or retry cloning if some local users or groups failed to be cloned after the last clone operation was completed.

Change password for tenant's local root user

You might need to change the password for a tenant's local root user if the root user is locked out of the account.

Before you begin

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

About this task

If single sign-on (SSO) is enabled for your StorageGRID system, the local root user can't sign in to the tenant account. To perform root user tasks, users must belong to a federated group that has the Root access permission for the tenant.

Steps

1. Select TENANTS.

Ter	nants						
View infoi To view m	rmation for each t nore recent values	tenant account. Depending o 5, select the tenant name.	n the timing of ingests, network c	onnecti	ivity, and node sta	tus, the usage data sho	wn might be out of date.
Create	Export to CSV	Actions 🐱 Search ter	nants by name or ID		Q		Displaying 5 results
	Name 💡 💠	Logical space used 🧿 💠	Quota utilization 🧿 ≑		Quota 💡 💠	Object count 🧿 🜩	Sign in/Copy URL 💡
	Tenant 01	2.00 GB		10%	20.00 GB	100	→ □
	Tenant 02	85.00 GB		85%	100.00 GB	500	→] [
	Tenant 03	500.00 TB		50%	1.00 PB	10,000	→ □
	Tenant 04	475.00 TB		95%	500.00 TB	50,000	→ □
	Tenant 05	5.00 GB			i i i	500	→ 「

- 2. Select the tenant account. You can do either of the following:
 - Select the checkbox for the tenant, and select Actions > Change root password.
 - Select the tenant's name to display the details page, and select Actions > Change root password.
- 3. Enter the new password for the tenant account.
- 4. Select Save.

Delete tenant account

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

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- · You have specific access permissions.
- You have removed all buckets (S3), containers (Swift), and objects associated with the tenant account.
- If the tenant is permitted to use a grid federation connection, you have reviewed the considerations for deleting a tenant with the Use grid federation connection permission.

Steps

- 1. Select TENANTS.
- 2. Locate the tenant account or accounts you want to delete.

Use the search box to search for a tenant by name or tenant ID.

- 3. To delete multiple tenants, select the checkboxes and select Actions > Delete.
- 4. To delete a single tenant, do either of the following:
 - Select the checkbox, and select Actions > Delete.
 - Select the tenant name to display the details page, and then select Actions > Delete.

5. Select Yes.

Manage platform services

Manage platform services for tenants: Overview

If you enable platform services for S3 tenant accounts, you must configure your grid so that tenants can access the external resources necessary to use these services.

What are platform services?

Platform services include CloudMirror replication, event notifications, and the search integration service.

CloudMirror replication

The StorageGRID CloudMirror replication service is used to mirror specific objects from a StorageGRID bucket to a specified external destination.

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



CloudMirror replication has some important similarities and differences with the cross-grid replication feature. To learn more, see Compare cross-grid replication and CloudMirror replication.



CloudMirror replication is not supported if the source bucket has S3 Object Lock enabled.

Notifications

Per-bucket event notifications are used to send notifications about specific actions performed on objects to a specified external Kafka cluster or Amazon Simple Notification Service.

For example, you could configure alerts to be sent to administrators about each object added to a bucket, where the objects represent log files associated with a critical system event.



Although event notification can be configured on a bucket with S3 Object Lock enabled, the S3 Object Lock metadata (including Retain Until Date and Legal Hold status) of the objects will not be included in the notification messages.

Search integration service

The search integration service is used to send S3 object metadata to a specified Elasticsearch index where the metadata can be searched or analyzed using the external service.

For example, you could configure your buckets to send S3 object metadata to a remote Elasticsearch service. You could then use Elasticsearch to perform searches across buckets, and perform sophisticated analyses of patterns present in your object metadata.



Although Elasticsearch integration can be configured on a bucket with S3 Object Lock enabled, the S3 Object Lock metadata (including Retain Until Date and Legal Hold status) of the objects will not be included in the notification messages.

Platform services give tenants the ability to use external storage resources, notification services, and search or analysis services with their data. Because the target location for platform services is typically external to your

StorageGRID deployment, you must decide if you want to permit tenants to use these services. If you do, you must enable the use of platform services when you create or edit tenant accounts. You must also configure your network such that the platform services messages that tenants generate can reach their destinations.

Recommendations for using platform services

Before using platform services, be aware of the following recommendations:

- If an S3 bucket in the StorageGRID system has both versioning and CloudMirror replication enabled, you should also enable S3 bucket versioning for the destination endpoint. This allows CloudMirror replication to generate similar object versions on the endpoint.
- You should not use more than 100 active tenants with S3 requests requiring CloudMirror replication, notifications, and search integration. Having more than 100 active tenants can result in slower S3 client performance.
- Requests to an endpoint that can't be completed will be queued to a maximum of 500,000 requests. This limit is equally shared among active tenants. New tenants are allowed to temporarily exceed this 500,000 limit so that newly created tenants aren't unfairly penalized.

Related information

- Manage platform services
- Configure Storage proxy settings
- Monitor StorageGRID

Network and ports for platform services

If you allow an S3 tenant to use platform services, you must configure networking for the grid to ensure that platform services messages can be delivered to their destinations.

You can enable platform services for an S3 tenant account when you create or update the tenant account. If platform services are enabled, the tenant can create endpoints that serve as a destination for CloudMirror replication, event notifications, or search integration messages from its S3 buckets. These platform services messages are sent from Storage Nodes that run the ADC service to the destination endpoints.

For example, tenants might configure the following types of destination endpoints:

- · A locally-hosted Elasticsearch cluster
- A local application that supports receiving Amazon Simple Notification Service messages
- · A locally-hosted Kafka cluster
- · A locally-hosted S3 bucket on the same or another instance of StorageGRID
- An external endpoint, such as an endpoint on Amazon Web Services.

To ensure that platform services messages can be delivered, you must configure the network or networks containing the ADC Storage Nodes. You must ensure that the following ports can be used to send platform services messages to the destination endpoints.

By default, platform services messages are sent on the following ports:

- 80: For endpoint URIs that begin with http (most endpoints)
- 443: For endpoint URIs that begin with https (most endpoints)
- 9092: For endpoint URIs that begin with http or https (Kafka endpoints only)

Tenants can specify a different port when they create or edit an endpoint.



If a StorageGRID deployment is used as the destination for CloudMirror replication, replication messages might be received on a port other than 80 or 443. Ensure that the port being used for S3 by the destination StorageGRID deployment is specified in the endpoint.

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

Related information

Use a tenant account

Per-site delivery of platform services messages

All platform services operations are performed on a per-site basis.

That is, if a tenant uses a client to perform an S3 API Create operation on an object by connecting to a Gateway Node at Data Center Site 1, the notification about that action is triggered and sent from Data Center Site 1.



If the client subsequently performs an S3 API Delete operation on that same object from Data Center Site 2,

the notification about the delete action is triggered and sent from Data Center Site 2.



Make sure that the networking at each site is configured such that platform services messages can be delivered to their destinations.

Troubleshoot platform services

The endpoints used in platform services are created and maintained by tenant users in the Tenant Manager; however, if a tenant has issues configuring or using platform services, you might be able to use the Grid Manager to help resolve the issue.

Issues with new endpoints

Before a tenant can use platform services, they must create one or more endpoints using the Tenant Manager. Each endpoint represents an external destination for one platform service, such as a StorageGRID S3 bucket, an Amazon Web Services bucket, an Amazon Simple Notification Service topic, a Kafka topic, or an Elasticsearch cluster hosted locally or on AWS. Each endpoint includes both the location of the external resource and the credentials needed to access that resource.

When a tenant creates an endpoint, the StorageGRID system validates that the endpoint exists and that it can be reached using the credentials that were specified. The connection to the endpoint is validated from one

node at each site.

If endpoint validation fails, an error message explains why endpoint validation failed. The tenant user should resolve the issue, then try creating the endpoint again.



i.

Endpoint creation will fail if platform services aren't enabled for the tenant account.

Issues with existing endpoints

If an error occurs when StorageGRID tries to reach an existing endpoint, a message is displayed on the dashboard in the Tenant Manager.



Tenant users can go to the Endpoints page to review the most recent error message for each endpoint and to determine how long ago the error occurred. The **Last error** column displays the most recent error message for

each endpoint and indicates how long ago the error occurred. Errors that include the 😵 icon occurred within the past 7 days.

Pla	Platform services endpoints								
A platfor replicati	A platform services endpoint stores the information StorageGRID needs to use an external resource as a target for a platform service (CloudMirror replication, notifications, or search integration). You must configure an endpoint for each platform service you plan to use.								
😣 One	e or more endpoints have exper	ienced an error. Select the e	endpoint for more deta	ils about the error. Meanwhile, the platf	form service request will be retried automatically.				
5 endpoir	nts				Create endpoint				
Delete er	ndpoint								
	Display name 🥹 🗢	Last error 💡 🌻	Туре 🔮 🗘	URI 🖗 ≑	URN 🔮 ≑				
	my-endpoint-2	ጰ 2 hours ago	Search	http://10.96.104.30:9200	urn:sgws:es:::mydomain/sveloso/_doc				
	my-endpoint-3	😢 3 days ago	Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example1				
	my-endpoint-5	12 days ago	Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example3				
	my-endpoint-4 Notifications http://10.96.104.202:8080/ arn:aws:sns:us-west-2::example2								
	my-endpoint-1 S3 Bucket http://10.96.104.167:10443 urn:sgws:s3:::bucket1								

Some error messages in the **Last error** column might include a logID in parentheses. A grid administrator or technical support can use this ID to locate more detailed information about the error in the bycast.log.

Issues related to proxy servers

If you have configured a storage proxy between Storage Nodes and platform service endpoints, errors might occur if your proxy service does not allow messages from StorageGRID. To resolve these issues, check the settings of your proxy server to ensure that platform service-related messages aren't blocked.

Determine if an error has occurred

If any endpoint errors have occurred within the past 7 days, the dashboard in the Tenant Manager displays an alert message. You can go the Endpoints page to see more details about the error.

Client operations fail

Some platform services issues might cause client operations on the S3 bucket to fail. For example, S3 client operations will fail if the internal Replicated State Machine (RSM) service stops, or if there are too many platform services messages queued for delivery.

To check the status of services:

- 1. Select **SUPPORT > Tools > Grid topology**.
- 2. Select site > Storage Node > SSM > Services.

Recoverable and unrecoverable endpoint errors

After endpoints have been created, platform service request errors can occur for various reasons. Some errors are recoverable with user intervention. For example, recoverable errors might occur for the following reasons:

- The user's credentials have been deleted or have expired.
- The destination bucket does not exist.
- The notification can't be delivered.

If StorageGRID encounters a recoverable error, the platform service request will be retried until it succeeds.

Other errors are unrecoverable. For example, an unrecoverable error occurs if the endpoint is deleted.

If StorageGRID encounters an unrecoverable endpoint error, the Total Events (SMTT) legacy alarm is triggered in the Grid Manager. To view the Total Events legacy alarm:

- 1. Select SUPPORT > Tools > Grid topology.
- 2. Select *site > node > SSM > 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. Select the **Configuration** tab to 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 can't be delivered

If the destination encounters an issue that prevents it from accepting platform services messages, the client operation on the bucket succeeds, but the platform services message is not delivered. For example, this error might happen if credentials are updated on the destination such that StorageGRID can no longer authenticate to the destination service.

If platform services messages can't be delivered because of an unrecoverable error, the Total Events (SMTT) legacy 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 error rate chart.



Platform services unavailable alert

The **Platform services unavailable** alert indicates that no platform service operations can be performed at a site because too few Storage Nodes with the RSM service are running or available.

The RSM service ensures platform service requests are sent to their respective endpoints.

To resolve this alert, determine which Storage Nodes at the site include the RSM service. (The RSM service is present on Storage Nodes that also include the ADC service.) Then, ensure that a simple majority of those Storage Nodes are running and available.



If more than one Storage Node that contains the RSM service fails at a site, you lose any pending platform service requests for that site.

Additional troubleshooting guidance for platform services endpoints

For additional information see Use a tenant account > Troubleshoot platform services endpoints.

Related information

Troubleshoot StorageGRID system

Manage S3 Select for tenant accounts

You can allow certain S3 tenants to use S3 Select to issue SelectObjectContent requests on individual objects.

S3 Select provides an efficient way to search through large amounts of data without having to deploy a database and associated resources to enable searches. It also reduces the cost and latency of retrieving data.

What is S3 Select?

S3 Select allows S3 clients to use SelectObjectContent requests to filter and retrieve only the data needed from an object. The StorageGRID implementation of S3 Select includes a subset of S3 Select commands and features.

Considerations and requirements for using S3 Select

Grid administration requirements

The grid administrator must grant tenants S3 Select ability. Select **Allow S3 Select** when creating a tenant or editing a tenant.

Object format requirements

The object you want to query must be in one of the following formats:

- CSV. Can be used as is or compressed into GZIP or BZIP2 archives.
- Parquet. Additional requirements for Parquet objects:
 - S3 Select supports only columnar compression using GZIP or Snappy. S3 Select doesn't support whole-object compression for Parquet objects.
 - S3 Select doesn't support Parquet output. You must specify the output format as CSV or JSON.
 - The maximum uncompressed row group size is 512 MB.
 - You must use the data types specified in the object's schema.
 - You can't use INTERVAL, JSON, LIST, TIME, or UUID logical types.

Endpoint requirements

The SelectObjectContent request must be sent to a StorageGRID load balancer endpoint.

The Admin and Gateway Nodes used by the endpoint must be one of the following:

- A services appliance node
- A VMware-based software node
- A bare metal node running a kernel with cgroup v2 enabled

General considerations

Queries can't be sent directly to Storage Nodes.



SelectObjectContent requests can decrease load-balancer performance for all S3 clients and all tenants. Enable this feature only when required and only for trusted tenants.

See the instructions for using S3 Select.

To view Grafana charts for S3 Select operations over time, select **SUPPORT** > **Tools** > **Metrics** in the Grid Manager.

Configure client connections

Configure S3 and Swift client connections: Overview

As a grid administrator, you manage the configuration options that control how S3 and Swift client applications connect to your StorageGRID system to store and retrieve data.



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

Configuration workflow

As shown in the workflow diagram, there are four primary steps for connecting StorageGRID to any S3 or Swift application:

- 1. Perform prerequisite tasks in StorageGRID, based on how the client application will connect to StorageGRID.
- 2. Use StorageGRID to obtain the values the application needs to connect to the grid. You can either use the S3 setup wizard or configure each StorageGRID entity manually.
- 3. Use the S3 or Swift application to complete the connection to StorageGRID. Create DNS entries to associate IP addresses to any domain names you plan to use.
- 4. Perform ongoing tasks in the application and in StorageGRID to manage and monitor object storage over time.



Information needed to attach StorageGRID to a client application

Before you can attach StorageGRID to an S3 or Swift client application, you must perform configuration steps in StorageGRID and obtain certain value.

What values do I need?

The following table shows the values you must configure in StorageGRID and where those values are used by the S3 or Swift application and the DNS server.

Value	Where value is configured	Where value is used
Virtual IP (VIP) addresses	StorageGRID > HA group	DNS entry
Port	StorageGRID > Load balancer endpoint	Client application
SSL certificate	StorageGRID > Load balancer endpoint	Client application
Server name (FQDN)	StorageGRID > Load balancer endpoint	Client applicationDNS entry
S3 access key ID and secret access key	StorageGRID > Tenant and bucket	Client application
Bucket/Container name	StorageGRID > Tenant and bucket	Client application

How do I get these values?

Depending on your requirements, you can do either of the following to obtain the information you need:

• Use the S3 setup wizard. The S3 setup wizard helps you to quickly configure the required values in StorageGRID and outputs one or two files that you can use when you configure the S3 application. The wizard guides you through the required steps and helps to make sure your settings conform to StorageGRID best practices.



If you are configuring an S3 application, using the S3 setup wizard is recommended unless you know you have special requirements or your implementation will require significant customization.

• Use the FabricPool setup wizard. Similar to the S3 setup wizard, the FabricPool setup wizard helps you to quickly configure required values and outputs a file that you can use when you configure a FabricPool cloud tier in ONTAP.



If you plan to use StorageGRID as the object storage system for a FabricPool cloud tier, using the FabricPool setup wizard is recommended unless you know you have special requirements or your implementation will require significant customization.

- **Configure items manually**. If you are connecting to a Swift application (or you are connecting to an S3 application and prefer not to use the S3 setup wizard), you can obtain the required values by performing the configuration manually. Follow these steps:
 - 1. Configure the high availability (HA) group you want to use for the S3 or Swift application. See Configure high availability groups.
 - 2. Create the load balancer endpoint that the S3 or Swift application will use. See Configure load balancer endpoints.
 - 3. Create the tenant account that the S3 or Swift application will use. See Create a tenant account.
 - 4. For an S3 tenant, sign in to the tenant account, and generate an access key ID and secret access key

for each user that will access the application. See Create your own access keys.

- 5. Create one or more S3 buckets or Swift containers within the tenant account. For S3, see Create S3 bucket. For Swift, use the PUT container request.
- 6. To add specific placement instructions for the objects belonging to the new tenant or bucket/container, create a new ILM rule and activate a new ILM policy to use that rule. See Create ILM rule and Create ILM policy.

Security for S3 or Swift clients

StorageGRID tenant accounts use S3 or Swift client applications to save object data to StorageGRID. You should review the security measures implemented for client applications.

Summary

The following table summarizes how security is implemented for the S3 and Swift REST APIs:

Security issue	Implementation for REST API
Connection security	TLS
Server authentication	X.509 server certificate signed by system CA or custom server certificate supplied by administrator
Client authentication	S3S3 account (access key ID and secret access key)SwiftSwift account (user name and password)
Client authorization	 S3 Bucket ownership and all applicable access control policies Swift Administrator role access

How StorageGRID provides security for client applications

S3 and Swift client applications can connect to the Load Balancer service on Gateway Nodes or Admin Nodes or directly to Storage Nodes.

• Clients that connect to the Load Balancer service can use HTTPS or HTTP, based on how you configure the load balancer endpoint.

HTTPS provides secure, TLS-encrypted communication and is recommended. You must attach a security certificate to the endpoint.

HTTP provides less secure, unencrypted communication and should only be used for non-production or test grids.

· Clients that connect to Storage Nodes can also use HTTPS or HTTP.

HTTPS is the default and is recommended.

HTTP provides less secure, unencrypted communication but can be optionally enabled for non-production or test grids.

- Communications between StorageGRID and the client are encrypted using TLS.
- Communications between the Load Balancer service and Storage Nodes within the grid are encrypted whether the load balancer endpoint is configured to accept HTTP or HTTPS connections.
- Clients must supply HTTP authentication headers to StorageGRID to perform REST API operations. See Authenticate requests and Supported Swift API endpoints.

Security certificates and client applications

In all cases, client applications can make TLS connections using either a custom server certificate uploaded by the grid administrator or a certificate generated by the StorageGRID system:

• When client applications connect to the Load Balancer service, they use the certificate that was configured for the load balancer endpoint. Each load balancer endpoint has its own certificate—either a custom server certificate uploaded by the grid administrator or a certificate that the grid administrator generated in StorageGRID when configuring the endpoint.

See Considerations for load balancing.

• When client applications connect directly to a Storage Node, they use either the system-generated server certificates that were generated for Storage Nodes when the StorageGRID system was installed (which are signed by the system certificate authority), or a single custom server certificate that is supplied for the grid by a grid administrator. See add a custom S3 or Swift API certificate.

Clients should be configured to trust the certificate authority that signed whichever certificate they use to establish TLS connections.

Supported hashing and encryption algorithms for TLS libraries

The StorageGRID system supports a set of cipher suites that client applications can use when establishing a TLS session. To configure ciphers, go to **CONFIGURATION** > **Security** > **Security settings** and select **TLS and SSH policies**.

Supported versions of TLS

StorageGRID supports TLS 1.2 and TLS 1.3.



SSLv3 and TLS 1.1 (or earlier versions) are no longer supported.

Use S3 setup wizard

Use S3 setup wizard: Considerations and requirements

You can use the S3 setup wizard to configure StorageGRID as the object storage system for an S3 application.

When to use the S3 setup wizard

The S3 setup wizard guides you through each step of configuring StorageGRID for use with an S3 application. As part of completing the wizard, you download files that you can use to enter values into the S3 application. Use the wizard to configure your system more quickly and to make sure your settings conform to StorageGRID best practices.

If you have the Root access permission, you can complete the S3 setup wizard when you start using the StorageGRID Grid Manager, or you can access and complete the wizard at any later time. Depending on your requirements, you can also configure some or all of the required items manually and then use the wizard to assemble the values that an S3 application needs.

Before using the wizard

Before using the wizard, confirm you have completed these prerequisites.

Obtain IP addresses and set up VLAN interfaces

If you will configure a high availability (HA) group, you know which nodes the S3 application will connect to and which StorageGRID network will be used. You also know which values to enter for the subnet CIDR, gateway IP address, and virtual IP (VIP) addresses.

If you plan to use a virtual LAN to segregate the traffic from the S3 application, you have already configured the VLAN interface. See Configure VLAN interfaces.

Configure identity federation and SSO

If you plan to use identity federation or single sign-on (SSO) for your StorageGRID system, you have enabled these features. You also know which federated group should have root access for the tenant account that the S3 application will use. See Use identity federation and Configure single sign-on.

Obtain and configure domain names

You know which fully qualified domain name (FQDN) to use for StorageGRID. Domain name server (DNS) entries will map this FQDN to the virtual IP (VIP) addresses of the HA group that you create using the wizard.

If you plan to use S3 virtual hosted-style requests, you should have configured S3 endpoint domain names. Using virtual hosted-style requests is recommended.

Review load balancer and security certificate requirements

If you plan to use the StorageGRID load balancer, you have reviewed the general considerations for load balancing. You have the certificates you will upload or the values you need to generate a certificate.

If you plan to use an external (third-party) load balancer endpoint, you have the fully qualified domain name (FQDN), port, and certificate for that load balancer.

Configure any grid federation connections

If you want to allow the S3 tenant to clone account data and replicate bucket objects to another grid using a grid federation connection, confirm the following before starting the wizard:

- You have configured the grid federation connection.
- The status of the connection is **Connected**.
- You have Root access permission.

You can use the S3 setup wizard to configure StorageGRID for use with an S3 application. The setup wizard provides the values the application needs to access a StorageGRID bucket and to save objects.

Before you begin

- You have the Root access permission.
- You have reviewed the considerations and requirements for using the wizard.

Access the wizard

Steps

- 1. Sign in to the Grid Manager using a supported web browser.
- If the FabricPool and S3 setup wizard banner appears on the dashboard, select the link in the banner. If the banner no longer appears, select the help icon from the header bar in the Grid Manager and select FabricPool and S3 setup wizard.



3. In the S3 application section of the FabricPool and S3 setup wizard page, select **Configure now**.

Step 1 of 6: Configure HA group

An HA group is a collection of nodes that each contain the StorageGRID Load Balancer service. An HA group can contain Gateway Nodes, Admin Nodes, or both.

You can use an HA group to help keep the S3 data connections available. If the active interface in the HA group fails, a backup interface can manage the workload with little impact to S3 operations.

For details about this task, see Manage high availability groups.

Steps

- 1. If you plan to use an external load balancer, you don't need to create an HA group. Select **Skip this step** and go to Step 2 of 6: Configure load balancer endpoint.
- 2. To use the StorageGRID load balancer, you can create a new HA group or use an existing HA group.

Create HA group

- a. To create a new HA group, select Create HA group.
- b. For the Enter details step, complete the following fields.

Field	Description
HA group name	A unique display name for this HA group.
Description (optional)	The description of this HA group.

c. For the Add interfaces step, select the node interfaces you want to use in this HA group.

Use the column headers to sort the rows, or enter a search term to locate interfaces more quickly.

You can select one or more nodes, but you can select only one interface for each node.

d. For the **Prioritize interfaces** step, determine the Primary interface and any backup interfaces for this HA group.

Drag rows to change the values in the **Priority order** column.

The first interface in the list is the Primary interface. The Primary interface is the active interface unless a failure occurs.

If the HA group includes more than one interface and the active interface fails, the virtual IP (VIP) addresses move to the first backup interface in the priority order. If that interface fails, the VIP addresses move to the next backup interface, and so on. When failures are resolved, the VIP addresses move back to the highest priority interface available.

e. For the Enter IP addresses step, complete the following fields.

Field	Description
Subnet CIDR	The address of the VIP subnet in CIDR notation — an IPv4 address followed by a slash and the subnet length (0-32). The network address must not have any host bits set. For example, 192.16.0.0/22.
Gateway IP address (optional)	If the S3 IP addresses used to access StorageGRID aren't on the same subnet as the StorageGRID VIP addresses, enter the StorageGRID VIP local gateway IP address. The local gateway IP address must be within the VIP subnet.
Virtual IP address	Enter at least one and no more than ten VIP addresses for the active interface in the HA group. All VIP addresses must be within the VIP subnet. At least one address must be IPv4. Optionally, you can specify additional IPv4 and IPv6 addresses.

- f. Select Create HA group and then select Finish to return to the S3 setup wizard.
- g. Select **Continue** to go to the load balancer step.

Use existing HA group

- a. To use an existing HA group, select the HA group name from the Select an HA group.
- b. Select **Continue** to go to the load balancer step.

Step 2 of 6: Configure load balancer endpoint

StorageGRID uses a load balancer to manage the workload from client applications. Load balancing maximizes speed and connection capacity across multiple Storage Nodes.

You can use the StorageGRID Load Balancer service, which exists on all Gateway and Admin Nodes, or you can connect to an external (third-party) load balancer. Using the StorageGRID load balancer is recommended.

For details about this task, see Considerations for load balancing.

To use the StorageGRID Load Balancer service, select the **StorageGRID load balancer** tab and then create or select the load balancer endpoint you want to use. To use an external load balancer, select the **External load balancer** tab and provide details about the system you have already configured.

Create endpoint

Steps

- 1. To create a load balancer endpoint, select **Create endpoint**.
- 2. For the Enter endpoint details step, complete the following fields.

Field	Description
Name	A descriptive name for the endpoint.
Port	The StorageGRID port you want to use for load balancing. This field defaults to 10433 for the first endpoint you create, but you can enter any unused external port. If you enter 80 or 443, the endpoint is configured only on Gateway Nodes, because these ports are reserved on Admin Nodes. Note: Ports used by other grid services aren't permitted. See the Network port reference.
Client type	Must be S3 .
Network protocol	Select HTTPS . Note : Communicating with StorageGRID without TLS encryption is supported but not recommended.

3. For the **Select binding mode** step, specify the binding mode. The binding mode controls how the endpoint is accessed using any IP address or using specific IP addresses and network interfaces.

Mode	Description
Global (default)	Clients can access the endpoint using the IP address of any Gateway Node or Admin Node, the virtual IP (VIP) address of any HA group on any network, or a corresponding FQDN.
	Use the Global setting (default) unless you need to restrict the accessibility of this endpoint.
Virtual IPs of HA groups	Clients must use a virtual IP address (or corresponding FQDN) of an HA group to access this endpoint.
	Endpoints with this binding mode can all use the same port number, as long as the HA groups you select for the endpoints don't overlap.
Node interfaces	Clients must use the IP addresses (or corresponding FQDNs) of selected node interfaces to access this endpoint.

Mode	Description
Node type	Based on the type of node you select, clients must use either the IP address (or corresponding FQDN) of any Admin Node or the IP address (or corresponding FQDN) of any Gateway Node to access this endpoint.

4. For the Tenant access step, select one of the following:

Field	Description
Allow all tenants (default)	All tenant accounts can use this endpoint to access their buckets.
Allow selected tenants	Only the selected tenant accounts can use this endpoint to access their buckets.
Block selected tenants	The selected tenant accounts can't use this endpoint to access their buckets. All other tenants can use this endpoint.

5. For the Attach certificate step, select one of the following:

Field	Description
Upload certificate (recommended)	Use this option to upload a CA-signed server certificate, certificate private key, and optional CA bundle.
Generate certificate	Use this option to generate a self-signed certificate. See Configure load balancer endpoints for details of what to enter.
Use StorageGRID S3 and Swift certificate	Use this option only if you have already uploaded or generated a custom version of the StorageGRID global certificate. See Configure S3 and Swift API certificates for details.

- 6. Select **Finish** to return to the S3 setup wizard.
- 7. Select **Continue** to go to the tenant and bucket step.



Changes to an endpoint certificate can take up to 15 minutes to be applied to all nodes.

Use existing load balancer endpoint

Steps

- 1. To use an existing endpoint, select its name from the **Select a load balancer endpoint**.
- 2. Select **Continue** to go to the tenant and bucket step.

Use external load balancer

Steps

1. To use an external load balancer, complete the following fields.

Field	Description
FQDN	The fully qualified domain name (FQDN) of the external load balancer.
Port	The port number that the S3 application will use to connect to the external load balancer.
Certificate	Copy the server certificate for the external load balancer and paste it into this field.

Step 3 of 6: Create tenant and bucket

A tenant is an entity that can use S3 applications to store and retrieve objects in StorageGRID. Each tenant has its own users, access keys, buckets, objects, and a specific set of capabilities. You must create the tenant before you can create the bucket that the S3 application will use to store its objects.

A bucket is a container used to store a tenant's objects and object metadata. Although some tenants might have many buckets, the wizard helps you to create a tenant and a bucket in the quickest and easiest way. You can use the Tenant Manager later to add any additional buckets you need.

You can create a new tenant for this S3 application to use. Optionally, you can also create a bucket for the new tenant. Finally, you can allow the wizard to create the S3 access keys for the tenant's root user.

For details about this task, see Create tenant account and Create S3 bucket.

Steps

- 1. Select Create tenant.
- 2. For the Enter details steps, enter the following information.

Field	Description
Name	A name for the tenant account. Tenant names don't need to be unique. When the tenant account is created, it receives a unique, numeric account ID.
Description (optional)	A description to help identify the tenant.
Client type	The type of client protocol this tenant will use. For the S3 setup wizard, S3 is selected and the field is disabled.
Storage quota (optional)	If you want this tenant to have a storage quota, a numerical value for the quota and the units.

3. Select Continue.

4. Optionally, select any permissions you want this tenant to have.



Some of these permissions have additional requirements. For details, select the help icon for each permission.

Permission	If selected
Allow platform services	The tenant can use S3 platform services such as CloudMirror. See Manage platform services for S3 tenant accounts.
Use own identity source	The tenant can configure and manage its own identity source for federated groups and users. This option is disabled if you have configured SSO for your StorageGRID system.
Allow S3 Select	The tenant can issue S3 SelectObjectContent API requests to filter and retrieve object data. See Manage S3 Select for tenant accounts. Important : SelectObjectContent requests can decrease load-balancer performance for all S3 clients and all tenants. Enable this feature only when required and only for trusted tenants.
Use grid federation connection	 The tenant can use a grid federation connection. Selecting this option: Causes this tenant and all tenant groups and users added to the account to be cloned from this grid (the <i>source grid</i>) to the other grid in the selected connection (the <i>destination grid</i>). Allows this tenant to configure cross-grid replication between corresponding buckets on each grid. See Manage the permitted tenants for grid federation.

- 5. If you selected **Use grid federation connection**, select one of the available grid federation connections.
- 6. Define root access for the tenant account, based on whether your StorageGRID system uses identity federation, single sign-on (SSO), or both.

Option	Do this
If identity federation is not enabled	Specify the password to use when signing into the tenant as the local root user.
If identity federation is enabled	 Select an existing federated group to have Root access permission for the tenant.
	Optionally, specify the password to use when signing in to the tenant as the local root user.
If both identity federation and single sign-on (SSO) are enabled	Select an existing federated group to have Root access permission for the tenant. No local users can sign in.

7. If you want the wizard to create the access key ID and secret access key for the root user, select Create

root user S3 access key automatically.



Select this option if the only user for the tenant will be the root user. If other users will use this tenant, use Tenant Manager to configure keys and permissions.

- 8. Select Continue.
- 9. For the Create bucket step, optionally create a bucket for the tenant's objects. Otherwise, select **Create tenant without bucket** to go to the download data step.



If S3 Object Lock is enabled for the grid, the bucket created in this step doesn't have S3 Object Lock enabled. If you need to use an S3 Object Lock bucket for this S3 application, select **Create tenant without bucket**. Then, use Tenant Manager to create the bucket instead.

a. Enter the name of the bucket that the S3 application will use. For example, S3-bucket.



You can't change the bucket name after creating the bucket.

b. Select the Region for this bucket.

Use the default region (us-east-1) unless you expect to use ILM in the future to filter objects based on the bucket's region.

- c. Select Enable object versioning if you want to store each version of each object in this bucket.
- d. Select Create tenant and bucket and go to the download data step.

Step 4 of 6: Download data

In the download data step, you can download one or two files to save the details of what you just configured.

Steps

1. If you selected Create root user S3 access key automatically, do one or both of the following:

 Select Download access keys to download a .csv file containing the tenant account name, access key ID, and secret access key.

- Select the copy icon (()) to copy the access key ID and secret access key to the clipboard.
- 2. Select **Download configuration values** to download a .txt file containing the settings for the load balancer endpoint, tenant, bucket, and the root user.
- 3. Save this information to a secure location.



Don't close this page until you have copied both access keys. The keys will not be available after you close this page. Make sure to save this information in a secure location because it can be used to obtain data from your StorageGRID system.

- 4. If prompted, select the checkbox to confirm that you have downloaded or copied the keys.
- 5. Select **Continue** to go to the ILM rule and policy step.

Step 5 of 6: Review ILM rule and ILM policy for S3

Information lifecycle management (ILM) rules control the placement, duration, and ingest behavior of all objects in your StorageGRID system. The ILM policy included with StorageGRID makes two replicated copies of all objects. This policy is in effect until you activate at least one new policy.

Steps

- 1. Review the information provided on the page.
- 2. If you want to add specific instructions for the objects belonging to the new tenant or bucket, create a new rule and a new policy. See Create ILM rule and ILM policies: Overview.
- 3. Select I have reviewed these steps and understand what I need to do.
- 4. Select the checkbox to indicate that you understand what to do next.
- 5. Select **Continue** to go to **Summary**.

Step 6 of 6: Review summary

Steps

- 1. Review the summary.
- 2. Make note of the details in the next steps, which describe the additional configuration that might be needed before you connect to the S3 client. For example, selecting **Sign in as root** takes you to the Tenant Manager, where you can add tenant users, create additional buckets, and update bucket settings.
- 3. Select Finish.
- 4. Configure the application using the file you downloaded from StorageGRID or the values you obtained manually.

Manage HA groups

Manage high availability (HA) groups: Overview

You can group the network interfaces of multiple Admin and Gateway Nodes into a high availability (HA) group. If the active interface in the HA group fails, a backup interface can manage the workload.

What is an HA group?

You can use high availability (HA) groups to provide highly available data connections for S3 and Swift clients or to provide highly available connections to the Grid Manager and the Tenant Manager.

Each HA group provides access to the shared services on the selected nodes.

- HA groups that include Gateway Nodes, Admin Nodes, or both provide highly available data connections for S3 and Swift clients.
- HA groups that include only Admin Nodes provide highly available connections to the Grid Manager and the Tenant Manager.
- An HA group that includes only services appliances and VMware-based software nodes can provide highly available connections for S3 tenants that use S3 Select.
 HA groups are recommended when using S3 Select, but not required.

How do you create an HA group?

1. You select a network interface for one or more Admin Nodes or Gateway Nodes. You can use a Grid Network (eth0) interface, Client Network (eth2) interface, VLAN interface, or an access interface you have added to the node.



You can't add an interface to an HA group if it has a DHCP-assigned IP address.

- 2. You specify one interface to be the Primary interface. The Primary interface is the active interface unless a failure occurs.
- 3. You determine the priority order for any Backup interfaces.
- 4. You assign one to 10 virtual IP (VIP) addresses to the group. Clients applications can use any of these VIP addresses to connect to StorageGRID.

For instructions, see Configure high availability groups.

What is the active interface?

During normal operation, all of the VIP addresses for the HA group are added to the Primary interface, which is the first interface in the priority order. As long as the Primary interface remains available, it is used when clients connect to any VIP address for the group. That is, during normal operation, the Primary interface is the "active" interface for the group.

Similarly, during normal operation, any lower priority interfaces for the HA group act as "backup" interfaces. These backup interfaces aren't used unless the Primary (currently active) interface becomes unavailable.

View the current HA group status of a node

To see if a node is assigned to an HA group and determine its current status, select NODES > node.

If the **Overview** tab includes an entry for **HA groups**, the node is assigned to the HA groups listed. The value after the group name is the current status of the node in the HA group:

- Active: The HA group is currently being hosted on this node.
- Backup: The HA group is not currently using this node; this is a backup interface.
- **Stopped**: The HA group can't be hosted on this node because the High Availability (keepalived) service has been stopped manually.
- Fault: The HA group can't be hosted on this node because of one or more of the following:
 - The Load Balancer (nginx-gw) service is not running on the node.
 - The node's eth0 or VIP interface is down.
 - The node is down.

In this example, the primary Admin Node has been added to two HA groups. This node is currently the active interface for the Admin clients group and a backup interface for the FabricPool clients group.

DC1-ADM1 (Primary Admin Node) 🗹				
Overview	Hardware Network Storage Load balancer Tasks			
Node information 🥝				
Name:	DC1-ADM1			
Type:	Primary Admin Node			
ID:	ce00d9c8-8a79-4742-bdef-c9c658db5315			
Connection state	Connected			
Software version	11.6.0 (build 20211207.1804.614bc17)			
HA groups:	Admin clients (Active)			
	FabricPool clients (Backup)			
IP addresses:	172.16.1.225 - eth0 (Grid Network)			
	10.224.1.225 - eth1 (Admin Network)			
	47.47.0.2, 47.47.1.225 - eth2 (Client Network)			
	Show additional IP addresses 🗸			

What happens when the active interface fails?

The interface that currently hosts the VIP addresses is the active interface. If the HA group includes more than one interface and the active interface fails, the VIP addresses move to the first available backup interface in the priority order. If that interface fails, the VIP addresses move to the next available backup interface, and so on.

Failover can be triggered for any of these reasons:

- The node on which the interface is configured goes down.
- The node on which the interface is configured loses connectivity to all other nodes for at least 2 minutes.
- The active interface goes down.
- The Load Balancer service stops.
- The High Availability service stops.



Failover might not be triggered by network failures external to the node that hosts the active interface. Similarly, failover is not triggered by the services for the Grid Manager or the Tenant Manager.

The failover process generally takes only a few seconds and is fast enough that client applications should experience little impact and can rely on normal retry behaviors to continue operation.

When failure is resolved and a higher priority interface becomes available again, the VIP addresses are automatically moved to the highest priority interface that is available.

How are HA groups used?

You can use high availability (HA) groups to provide highly available connections to StorageGRID for object data and for administrative use.

- An HA group can provide highly available administrative connections to the Grid Manager or the Tenant Manager.
- An HA group can provide highly available data connections for S3 and Swift clients.
- An HA group that contains only one interface allows you to provide many VIP addresses and to explicitly set IPv6 addresses.

An HA group can provide high availability only if all nodes included in the group provide the same services. When you create an HA group, add interfaces from the types of nodes that provide the services you require.

- Admin Nodes: Include the Load Balancer service and enable access to the Grid Manager or the Tenant Manager.
- Gateway Nodes: Include the Load Balancer service.

Purpose of HA group	Add nodes of this type to the HA group
Access to Grid Manager	 Primary Admin Node (Primary) Non-primary Admin Nodes Note: The primary Admin Node must be the Primary interface. Some maintenance procedures can only be performed from the primary Admin Node.
Access to Tenant Manager only	Primary or non-primary Admin Nodes
S3 or Swift client access — Load Balancer service	Admin NodesGateway Nodes
S3 client access for S3 Select	 Services appliances VMware-based software nodes Note: HA groups are recommended when using S3 Select, but not required.

Limitations of using HA groups with Grid Manager or Tenant Manager

If a Grid Manager or Tenant Manager service fails, HA group failover is not triggered.

If you are signed in to the Grid Manager or the Tenant Manager when failover occurs, you are signed out and must sign in again to resume your task.

Some maintenance procedures can't be performed when the primary Admin Node is unavailable. During failover, you can use the Grid Manager to monitor your StorageGRID system.

The following diagrams provide examples of different ways you can configure HA groups. Each option has advantages and disadvantages.

In the diagrams, blue indicates the primary interface in the HA group and yellow indicates the backup interface in the HA group.





The table summarizes the benefits of each HA configuration shown in the diagram.

Configuration	Advantages	Disadvantages
Active-Backup HA	 Managed by StorageGRID with no external dependencies. Fast failover. 	 Only one node in an HA group is active. At least one node per HA group will be idle.
Configuration	Advantages	Disadvantages
------------------	--	---
DNS Round Robin	Increased aggregate throughput.No idle hosts.	 Slow failover, which could depend on client behavior. Requires configuration of hardware outside of StorageGRID. Needs a customer-implemented health check.
Active-Active HA	 Traffic is distributed across multiple HA groups. High aggregate throughput that scales with the number of HA groups. Fast failover. 	 More complex to configure. Requires configuration of hardware outside of StorageGRID. Needs a customer-implemented health check.

Configure high availability groups

You can configure high availability (HA) groups to provide highly available access to the services on Admin Nodes or Gateway Nodes.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.
- If you plan to use a VLAN interface in an HA group, you have created the VLAN interface. See Configure VLAN interfaces.
- If you plan to use an access interface for a node in an HA group, you have created the interface:
 - Red Hat Enterprise Linux (before installing the node): Create node configuration files
 - Ubuntu or Debian (before installing the node): Create node configuration files
 - Linux (after installing the node): Linux: Add trunk or access interfaces to a node
 - VMware (after installing the node): VMware: Add trunk or access interfaces to a node

Create a high availability group

When you create a high availability group, you select one or more interfaces and organize them in priority order. Then, you assign one or more VIP addresses to the group.

An interface must be for a Gateway Node or an Admin Node to be included in an HA group. An HA group can only use one interface for any given node; however, other interfaces for the same node can be used in other HA groups.

Access the wizard

Steps

- 1. Select CONFIGURATION > Network > High availability groups.
- 2. Select Create.

Enter details for the HA group

Steps

- 1. Provide a unique name for the HA group.
- 2. Optionally, enter a description for the HA group.
- 3. Select Continue.

Add interfaces to the HA group

Steps

1. Select one or more interfaces to add to this HA group.

Use the column headers to sort the rows, or enter a search term to locate interfaces more quickly.

earch	2	Q			Total interface cour
	Node ≑	Interface 🥝 ≑	Site 😧 💠	IPv4 subnet ≑	Node type 😧 💠
	DC1-ADM1-104-96	eth0 🕜	DC1	10.96.104.0/22	Primary Admin Node
	DC1-ADM1-104-96	eth2 😧	DC1	-	Primary Admin Node
	DC2-ADM1-104-103	eth0 💡	DC2	10.96.104.0/22	Admin Node
	DC2-ADM1-104-103	eth2 🔞	DC2		Admin Node



After creating a VLAN interface, wait up to 5 minutes for the new interface to appear in the table.

Guidelines for selecting interfaces

- You must select at least one interface.
- $\,\circ\,$ You can select only one interface for a node.
- If the HA group is for HA protection of Admin Node services, which include the Grid Manager and the Tenant Manager, select interfaces on Admin Nodes only.
- If the HA group is for HA protection of S3 or Swift client traffic, select interfaces on Admin Nodes, Gateway Nodes, or both.
- If you select interfaces on different types of nodes, an informational note appears. You are reminded that if a failover occurs, services provided by the previously active node might not be available on the newly active node. For example, a backup Gateway Node can't provide HA protection of Admin Node services. Similarly, a backup Admin Node can't perform all of the maintenance procedures that the primary Admin Node can provide.

• If you can't select an interface, its checkbox is disabled. The tool tip provides more information.



- You can't select an interface if its subnet value or gateway conflicts with another selected interface.
- You can't select a configured interface if it does not have a static IP address.

2. Select Continue.

Determine the priority order

If the HA group includes more than one interface, you can determine which is the Primary interface and which are the Backup (failover) interfaces. If the Primary interface fails, the VIP addresses move to the highest priority interface that is available. If that interface fails, the VIP addresses move to the next highest priority interface that is available, and so on.

Steps

1. Drag rows in the **Priority order** column to determine the Primary interface and any Backup interfaces.

The first interface in the list is the Primary interface. The Primary interface is the active interface unless a failure occurs.

Determine the priority order Determine the primary interface and the backup (failover) interfaces for this HA group. Drag and drop rows or select the Inrows.			
Priority order 💡	Node	Interface 👔	Node type 👔
1 (Primary interface)	DC1-ADM1-104-96	eth2	Primary Admin Node
2	DC2-ADM1-104-103	eth2	Admin Node



If the HA group provides access to the Grid Manager, you must select an interface on the primary Admin Node to be the Primary interface. Some maintenance procedures can only be performed from the primary Admin Node.

2. Select Continue.

Enter IP addresses

Steps

1. In the **Subnet CIDR** field, specify the VIP subnet in CIDR notation—an IPv4 address followed by a slash and the subnet length (0-32).

The network address must not have any host bits set. For example, 192.16.0.0/22.



If you use a 32-bit prefix, the VIP network address also serves as the gateway address and the VIP address.

Inter details for the HA group
ubnet CIDR 🥹
pecify the subnet in CIDR notation. The optional gateway IP and all VIPs must be in this subnet.
Pv4 address followed by a slash and the subnet length (0-32)
ateway IP address (optional)
ptionally specify the IP address of the gateway, which must be in the subnet. If the subnet address length is 32, the gateway IP ddress is automatically set to the subnet IP.
Virtual IP address 💡
pecify at least 1 and no more than 10 virtual IPs for the HA group. All virtual IPs must be in the same subnet. If the subnet length is
2, only one VIP is allowed, which is automatically set to the subnet/gateway IP.
1.2.3.4
dd another IP address

2. Optionally, if any S3, Swift, administrative or tenant clients will access these VIP addresses from a different subnet, enter the **Gateway IP address**. The gateway address must be within the VIP subnet.

Client and admin users will use this gateway to access the virtual IP addresses.

3. Enter at least one and no more than ten VIP addresses for the active interface in the HA group. All VIP addresses must be within the VIP subnet and all will be active at the same time on the active interface.

You must provide at least one IPv4 address. Optionally, you can specify additional IPv4 and IPv6 addresses.

4. Select Create HA group and select Finish.

The HA Group is created, and you can now use the configured virtual IP addresses.

Next steps

If you will use this HA group for load balancing, create a load balancer endpoint to determine the port and network protocol and to attach any required certificates. See Configure load balancer endpoints.

Edit a high availability group

You can edit a high availability (HA) group to change its name and description, add or remove interfaces, change the priority order, or add or update virtual IP addresses.

For example, you might need to edit an HA group if you want to remove the node associated with a selected interface in a site or node decommission procedure.

Steps

1. Select CONFIGURATION > Network > High availability groups.

The High availability groups page shows all existing HA groups.

- 2. Select the checkbox for the HA group you want to edit.
- 3. Do one of the following, based on what you want to update:
 - Select Actions > Edit virtual IP address to add or remove VIP addresses.
 - Select **Actions** > **Edit HA group** to update the group's name or description, add or remove interfaces, change the priority order, or add or remove VIP addresses.

4. If you selected Edit virtual IP address:

- a. Update the virtual IP addresses for the HA group.
- b. Select Save.
- c. Select Finish.
- 5. If you selected Edit HA group:
 - a. Optionally, update the group's name or description.
 - b. Optionally, select or clear the checkboxes to add or remove interfaces.



If the HA group provides access to the Grid Manager, you must select an interface on the primary Admin Node to be the Primary interface. Some maintenance procedures can only be performed from the primary Admin Node

- c. Optionally, drag rows to change the priority order of the Primary interface and any Backup interfaces for this HA group.
- d. Optionally, update the virtual IP addresses.
- e. Select Save and then select Finish.

Remove a high availability group

You can remove one or more high availability (HA) groups at a time.



You can't remove an HA group if it is bound to a load balancer endpoint. To delete an HA group, you must remove it from any load balancer endpoints that use it.

To prevent client disruptions, update any affected S3 or Swift client applications before you remove an HA group. Update each client to connect using another IP address, for example, the virtual IP address of a different HA group or the IP address that was configured for an interface during installation.

Steps

1. Select CONFIGURATION > Network > High availability groups.

- 2. Review the **Load balancer endpoints** column for each HA group you want to remove. If any load balancer endpoints are listed:
 - a. Go to CONFIGURATION > Network > Load balancer endpoints.
 - b. Select the checkbox for the endpoint.
 - c. Select Actions > Edit endpoint binding mode.
 - d. Update the binding mode to remove the HA group.
 - e. Select Save changes.
- 3. If no load balancer endpoints are listed, select the checkbox for each HA group you want to remove.
- 4. Select Actions > Remove HA group.
- 5. Review the message and select **Delete HA group** to confirm your selection.

All HA groups you selected are removed. A green success banner appears on the High availability groups page.

Manage load balancing

Considerations for load balancing

You can use load balancing to handle ingest and retrieval workloads from S3 and Swift clients.

What is load balancing?

When a client application saves or retrieves data from a StorageGRID system, StorageGRID uses a load balancer to manage the ingest and retrieval workload. Load balancing maximizes speed and connection capacity by distributing the workload across multiple Storage Nodes.

The StorageGRID Load Balancer service is installed on all Admin Nodes and all Gateway Nodes and provides Layer 7 load balancing. It performs Transport Layer Security (TLS) termination of client requests, inspects the requests, and establishes new secure connections to the Storage Nodes.

The Load Balancer service on each node operates independently when forwarding client traffic to the Storage Nodes. Through a weighting process, the Load Balancer service routes more requests to Storage Nodes with higher CPU availability.



Although the StorageGRID Load Balancer service is the recommended load balancing mechanism, you might want to integrate a third-party load balancer instead. For information, contact your NetApp account representative or refer to TR-4626: StorageGRID third-party and global load balancers.

How many load balancing nodes do I need?

As a general best practice, each site in your StorageGRID system should include two or more nodes with the Load Balancer service. For example, a site might include two Gateway Nodes or both an Admin Node and a Gateway Node. Make sure that there is adequate networking, hardware, or virtualization infrastructure for each load-balancing node, whether you are using services appliances, bare metal nodes, or virtual machine (VM) based nodes.

What is a load balancer endpoint?

A load balancer endpoint defines the port and the network protocol (HTTPS or HTTP) that incoming and outgoing client application requests will use to access those nodes that contain the Load Balancer service. The endpoint also defines the client type (S3 or Swift), the binding mode, and optionally a list of allowed or blocked tenants.

To create a load balancer endpoint, either select **CONFIGURATION** > **Network** > **Load balancer endpoints** or complete the FabricPool and S3 setup wizard. For instructions:

- Configure load balancer endpoints
- Use the S3 setup wizard
- Use the FabricPool setup wizard

Considerations for the port

The port for a load balancer endpoint defaults to 10433 for the first endpoint you create, but you can specify any unused external port between 1 and 65535. If you use port 80 or 443, the endpoint will use the Load Balancer service on Gateway Nodes only. These ports are reserved on Admin Nodes. If you use the same port for more than one endpoint, you must specify a different binding mode for each endpoint.

Ports used by other grid services aren't permitted. See the Network port reference.

Considerations for the network protocol

In most cases, the connections between client applications and StorageGRID should use Transport Layer Security (TLS) encryption. Connecting to StorageGRID without TLS encryption is supported but not recommended, especially in production environments. When you select the network protocol for the StorageGRID load balancer endpoint, you should select **HTTPS**.

Considerations for load balancer endpoint certificates

If you select **HTTPS** as the network protocol for the load balancer endpoint, you must provide a security certificate. You can use any of these three options when you create the load balancer endpoint:

• Upload a signed certificate (recommended). This certificate can be signed by either a publicly trusted or a private certificate authority (CA). Using a publicly trusted CA server certificate to secure the connection is the best practice. In contrast to generated certificates, certificates signed by a CA can be rotated nondisruptively, which can help avoid expiration issues.

You must obtain the following files before you create the load balancer endpoint:

- The custom server certificate file.
- The custom server certificate private key file.
- Optionally, a CA bundle of the certificates from each intermediate issuing certificate authority.
- Generate a self-signed certificate.
- Use the global StorageGRID S3 and Swift certificate. You must upload or generate a custom version of this certificate before you can select it for the load balancer endpoint. See Configure S3 and Swift API certificates.

What values do I need?

To create the certificate, you must know all of the domain names and IP addresses that S3 or Swift client applications will use to access the endpoint.

The **Subject DN** (Distinguished Name) entry for the certificate must include the fully qualified domain name that the client application will use for StorageGRID. For example:

```
Subject DN:
/C=Country/ST=State/O=Company,Inc./CN=s3.storagegrid.example.com
```

As required, the certificate can use wildcards to represent the fully qualified domain names of all Admin Nodes and Gateway Nodes running the Load Balancer service. For example, *.storagegrid.example.com uses the * wildcard to represent adm1.storagegrid.example.com and gn1.storagegrid.example.com.

If you plan to use S3 virtual hosted-style requests, the certificate must also include an **Alternative Name** entry for each S3 endpoint domain name you have configured, including any wildcard names. For example:

Alternative Name: DNS:*.s3.storagegrid.example.com



If you use wildcards for domain names, review the Hardening guidelines for server certificates.

You must also define a DNS entry for each name in the security certificate.

How do I manage expiring certificates?



If the certificate used to secure the connection between the S3 application and StorageGRID expires, the application might temporarily lose access to StorageGRID.

To avoid certificate expiration issues, follow these best practices:

- Carefully monitor any alerts that warn of approaching certificate expiration dates, such as the **Expiration of load balancer endpoint certificate** and **Expiration of global server certificate for S3 and Swift API** alerts.
- Always keep the StorageGRID and S3 application's versions of the certificate in sync. If you replace or renew the certificate used for a load balancer endpoint, you must replace or renew the equivalent certificate used by the S3 application.
- Use a publicly signed CA certificate. If you use a certificate signed by a CA, you can replace soon-to-expire certificates nondisruptively.
- If you have generated a self-signed StorageGRID certificate and that certificate is about to expire, you must
 manually replace the certificate in both StorageGRID and in the S3 application before the existing
 certificate expires.

Considerations for the binding mode

The binding mode lets you control which IP addresses can be used to access a load balancer endpoint. If an endpoint uses a binding mode, client applications can only access the endpoint if they use an allowed IP address or its corresponding fully qualified domain name (FQDN). Client applications using any other IP address or FQDN can't access the endpoint.

You can specify any of the following binding modes:

- **Global** (default): Client applications can access the endpoint using the IP address of any Gateway Node or Admin Node, the virtual IP (VIP) address of any HA group on any network, or a corresponding FQDN. Use this setting unless you need to restrict the accessibility of an endpoint.
- Virtual IPs of HA groups. Client applications must use a virtual IP address (or corresponding FQDN) of an HA group.
- Node interfaces. Clients must use the IP addresses (or corresponding FQDNs) of selected node interfaces.
- **Node type**. Based on the type of node you select, clients must use either the IP address (or corresponding FQDN) of any Admin Node or the IP address (or corresponding FQDN) of any Gateway Node.

Considerations for tenant access

Tenant access is an optional security feature that lets you control which StorageGRID tenant accounts can use a load balancer endpoint to access their buckets. You can allow all tenants to access an endpoint (default), or you can specify a list of the allowed or blocked tenants for each endpoint.

You can use this feature to provide better security isolation between tenants and their endpoints. For example, you might use this feature to ensure that the top-secret or highly classified materials owned by one tenant remain completely inaccessible to other tenants.



For the purpose of access control, the tenant is determined from the access keys used in the client request, if no access keys are provided as part of the request (such as with anonymous access) the bucket owner is used to determine the tenant.

Tenant access example

To understand how this security feature works, consider the following example:

- 1. You have created two load balancer endpoints, as follows:
 - Public endpoint: Uses port 10443 and allows access to all tenants.
 - **Top secret** endpoint: Uses port 10444 and allows access to the **Top secret** tenant only. All other tenants are blocked from accessing this endpoint.
- 2. The top-secret.pdf is in a bucket owned by the Top secret tenant.

To access the top-secret.pdf, a user in the **Top secret** tenant can issue a GET request to https://w.x.y.z:10444/top-secret.pdf. Because this tenant is allowed to use the 10444 endpoint, the user can access the object. However, if a user belonging to any other tenant issues the same request to the same URL, they receive an immediate Access Denied message. Access is denied even if the credentials and signature are valid.

CPU availability

The Load Balancer service on each Admin Node and Gateway Node operates independently when forwarding S3 or Swift traffic to the Storage Nodes. Through a weighting process, the Load Balancer service routes more requests to Storage Nodes with higher CPU availability. Node CPU load information is updated every few minutes, but weighting might be updated more frequently. All Storage Nodes are assigned a minimal base weight value, even if a node reports 100% utilization or fails to report its utilization.

In some cases, information about CPU availability is limited to the site where the Load Balancer service is

located.

Configure load balancer endpoints

Load balancer endpoints determine the ports and network protocols S3 and Swift clients can use when connecting to the StorageGRID load balancer on Gateway and Admin Nodes. You can also use endpoints to access the Grid Manager, Tenant Manager, or both.



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

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.
- You have reviewed the considerations for load balancing.
- If you previously remapped a port you intend to use for the load balancer endpoint, you have removed the port remap.
- You have created any high availability (HA) groups you plan to use. HA groups are recommended, but not required. See Manage high availability groups.
- If the load balancer endpoint will be used by S3 tenants for S3 Select, it must not use the IP addresses or FQDNs of any bare-metal nodes. Only services appliances and VMware-based software nodes are allowed for the load balancer endpoints used for S3 Select.
- You have configured any VLAN interfaces you plan to use. See Configure VLAN interfaces.
- If you are creating an HTTPS endpoint (recommended), you have the information for the server certificate.



Changes to an endpoint certificate can take up to 15 minutes to be applied to all nodes.

- To upload a certificate, you need the server certificate, the certificate private key, and optionally, a CA bundle.
- To generate a certificate, you need all of the domain names and IP addresses that S3 or Swift clients will use to access the endpoint. You must also know the subject (Distinguished Name).
- If you want to use the StorageGRID S3 and Swift API certificate (which can also be used for connections directly to Storage Nodes), you have already replaced the default certificate with a custom certificate signed by an external certificate authority. See Configure S3 and Swift API certificates.

Create a load balancer endpoint

Each S3 or Swift client load balancer endpoint specifies a port, a client type (S3 or Swift), and a network protocol (HTTP or HTTPS). Management interface load balancer endpoints specifies a port, interface type, and untrusted Client Network.

Access the wizard

Steps

1. Select CONFIGURATION > Network > Load balancer endpoints.

- 2. To create an endpoint for an S3 or Swift client, select the S3 or Swift client tab.
- 3. To create an endpoint for access to the Grid Manager, Tenant Manager, or both, select the **Management interface** tab.
- 4. Select Create.

Enter endpoint details

Steps

1. Select the appropriate instructions to enter details for the type of endpoint you want to create.

S3 or Swift client			
Field	Description		
Name	A descriptive name for the endpoint, which will appear in the table on the Load balancer endpoints page.		
Port	 The StorageGRID port you want to use for load balancing. This field defaults to 10433 for the first endpoint you create, but you can enter any unused external port from 1 to 65535. If you enter 80 or 8443, the endpoint is configured only on Gateway Nodes, unless you have freed up port 8443. Then you can use port 8443 as an S3 endpoint, and the port will be configured on both Gateway and Admin Nodes. 		
Client type	The type of client application that will use this endpoint, either S3 or Swift .		
Network protocol	 The network protocol that clients will use when connecting to this endpoint. Select HTTPS for secure, TLS encrypted communication (recommended). You must attach a security certificate before you can save the endpoint. Select HTTP for less secure, unencrypted communication. Use HTTP only for a non-production grid. 		

Management interface

Field	Description
Name	A descriptive name for the endpoint, which will appear in the table on the Load balancer endpoints page.
Port	 The StorageGRID port you want to use to access the Grid Manager, Tenant Manager, or both. Grid Manager: 8443 Tenant Manager: 9443 Both Grid Manager and Tenant Manager: 443
	Note : You can use these preset ports or other available ports.
Interface type	Select the radio button for the StorageGRID interface you will access using this endpoint.
Untrusted Client Network	Select Yes if this endpoint should be accessible to untrusted Client Networks. Otherwise, select No . When you select Yes , the port is open on all untrusted Client Networks.
	Note : You can only configure a port to be open or closed to untrusted Client Networks when you are creating the load balancer endpoint.

2. Select Continue.

Select a binding mode

Steps

1. Select a binding mode for the endpoint to control how the endpoint is accessed using any IP address or using specific IP addresses and network interfaces.

Some binding modes are available for either client endpoints or management interface endpoints. All modes for both endpoint types are listed here.

Mode	Description
Global (default for client endpoints)	Clients can access the endpoint using the IP address of any Gateway Node or Admin Node, the virtual IP (VIP) address of any HA group on any network, or a corresponding FQDN. Use the Global setting unless you need to restrict the accessibility of this endpoint.
Virtual IPs of HA groups	Clients must use a virtual IP address (or corresponding FQDN) of an HA group to access this endpoint. Endpoints with this binding mode can all use the same port number, as long as the HA groups you select for the endpoints don't overlap.
Node interfaces	Clients must use the IP addresses (or corresponding FQDNs) of selected node interfaces to access this endpoint.
Node type (client endpoints only)	Based on the type of node you select, clients must use either the IP address (or corresponding FQDN) of any Admin Node or the IP address (or corresponding FQDN) of any Gateway Node to access this endpoint.
All Admin Nodes (default for management interface endpoints)	Clients must use the IP address (or corresponding FQDN) of any Admin Node to access this endpoint.

If more than one endpoint uses the same port, StorageGRID uses this priority order to decide which endpoint to use: Virtual IPs of HA groups > Node interfaces > Node type > Global.

If you are creating management interface endpoints, only Admin Nodes are allowed.

2. If you selected Virtual IPs of HA groups, select one or more HA groups.

If you are creating management interface endpoints, select VIPs associated only with Admin Nodes.

- 3. If you selected **Node interfaces**, select one or more node interfaces for each Admin Node or Gateway Node that you want to associate with this endpoint.
- 4. If you selected **Node type**, select either Admin Nodes, which includes both the primary Admin Node and any non-primary Admin Nodes, or Gateway Nodes.

Control tenant access



A management interface endpoint can control tenant access only when the endpoint has the interface type of Tenant Manager.

Steps

1. For the Tenant access step, select one of the following:

Field	Description
Allow all tenants (default)	All tenant accounts can use this endpoint to access their buckets. You must select this option if you have not yet created any tenant accounts. After you add tenant accounts, you can edit the load balancer endpoint to allow or block specific accounts.
Allow selected tenants	Only the selected tenant accounts can use this endpoint to access their buckets.
Block selected tenants	The selected tenant accounts can't use this endpoint to access their buckets. All other tenants can use this endpoint.

2. If you are creating an **HTTP** endpoint, you don't need to attach a certificate. Select **Create** to add the new load balancer endpoint. Then, go to After you finish. Otherwise, select **Continue** to attach the certificate.

Attach certificate

Steps

1. If you are creating an **HTTPS** endpoint, select the type of security certificate you want to attach to the endpoint.

The certificate secures the connections between S3 and Swift clients and the Load Balancer service on Admin Node or Gateway Nodes.

- Upload certificate. Select this option if you have custom certificates to upload.
- Generate certificate. Select this option if you have the values needed to generate a custom certificate.
- **Use StorageGRID S3 and Swift certificate**. Select this option if you want to use the global S3 and Swift API certificate, which can also be used for connections directly to Storage Nodes.

You can't select this option unless you have replaced the default S3 and Swift API certificate, which is signed by the grid CA, with a custom certificate signed by an external certificate authority. See Configure S3 and Swift API certificates.

- **Use management interface certificate**. Select this option if you want to use the global management interface certificate, which can also be used for direct connections to Admin Nodes.
- 2. If you aren't using the StorageGRID S3 and Swift certificate, upload or generate the certificate.

Upload certificate

- a. Select Upload certificate.
- b. Upload the required server certificate files:
 - Server certificate: The custom server certificate file in PEM encoding.
 - Certificate private key: The custom server certificate private key file (.key).



EC private keys must be 224 bits or larger. RSA private keys must be 2048 bits or larger.

- **CA bundle**: A single optional file containing the certificates from each intermediate issuing certificate authority (CA). The file should contain each of the PEM-encoded CA certificate files, concatenated in certificate chain order.
- c. Expand **Certificate details** to see the metadata for each certificate you uploaded. If you uploaded an optional CA bundle, each certificate displays on its own tab.
 - Select Download certificate to save the certificate file or select Download CA bundle to save the certificate bundle.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid_certificate.pem

- Select Copy certificate PEM or Copy CA bundle PEM to copy the certificate contents for pasting elsewhere.
- d. Select Create.

The load balancer endpoint is created. The custom certificate is used for all subsequent new connections between S3 and Swift clients or the management interface and the endpoint.

Generate certificate

- a. Select Generate certificate.
- b. Specify the certificate information:

Field	Description	
Domain name	One or more fully qualified domain names to include in the certificate. Use an * as a wildcard to represent multiple domain names.	
IP	One or more IP addresses to include in the certificate.	
Subject (optional)	X.509 subject or distinguished name (DN) of the certificate owner. If no value is entered in this field, the generated certificate uses the first domain name or IP address as the subject common name (CN).	
Days valid	Number of days after creation that the certificate expires.	

Field	Description	
Add key usage extensions	If selected (default and recommended), key usage and extended key usage extensions are added to the generated certificate.	
	These extensions define the purpose of the key contained in the certificate.	
	Note : Leave this checkbox selected unless you experience connection problems with older clients when certificates include these extensions.	

- c. Select Generate.
- d. Select Certificate details to see the metadata for the generated certificate.
 - Select **Download certificate** to save the certificate file.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid certificate.pem

- Select Copy certificate PEM to copy the certificate contents for pasting elsewhere.
- e. Select Create.

The load balancer endpoint is created. The custom certificate is used for all subsequent new connections between S3 and Swift clients or the management interface and this endpoint.

After you finish

Steps

1. If you use a DNS, ensure that the DNS includes a record to associate the StorageGRID fully qualified domain name (FQDN) to each IP address that clients will use to make connections.

The IP address you enter in the DNS record depends on whether you are using an HA group of loadbalancing nodes:

- If you have configured an HA group, clients will connect to the virtual IP addresses of that HA group.
- If you aren't using an HA group, clients will connect to the StorageGRID Load Balancer service using the IP address of a Gateway Node or Admin Node.

You must also ensure that the DNS record references all required endpoint domain names, including any wildcard names.

- 2. Provide S3 and Swift clients with the information needed to connect to the endpoint:
 - Port number
 - Fully qualified domain name or IP address
 - · Any required certificate details

View and edit load balancer endpoints

You can view details for existing load balancer endpoints, including the certificate metadata for a secured endpoint. You can change certain settings for an endpoint.

- To view basic information for all load balancer endpoints, review the tables on the Load balancer endpoints page.
- To view all details about a specific endpoint, including certificate metadata, select the endpoint's name in the table. The information shown varies depending on the endpoint type and how it's configured.

S3 load balancer endpoint 🧪			
Port:	10443		
Client type:	53		
Network protocol:	HTTPS		
Binding mode:	Global		
Endpoint ID:	3d02c126-9437-478c-8b24-08384401d3cb		
Remove Binding mod	le Certificate Tenant access (2 allowed)		
You can select a different binding mode or change IP addresses for the current binding mode. Edit binding mode Binding mode: Global			
This endpoint uses the Global binding mode. Unless there are one or more overriding endpoints for the same port, clients can access this endpoint using the IP address of any Gateway Node, any Admin Node, or the virtual IP of any HA group on any network.			

• To edit an endpoint, use the **Actions** menu on the Load balancer endpoints page.



If you lose access to Grid Manager while editing the port of a management interface endpoint, update the URL and port to regain access.



After editing an endpoint, you might need to wait up to 15 minutes for your changes to be applied to all nodes.

Task	Actions menu	Details page
Edit endpoint name	 a. Select the checkbox for the endpoint. b. Select Actions > Edit endpoint name. c. Enter the new name. d. Select Save. 	 a. Select the endpoint name to display the details. b. Select the edit icon . c. Enter the new name. d. Select Save.
Edit endpoint port	 a. Select the checkbox for the endpoint. b. Select Actions > Edit endpoint port c. Enter a valid port number. d. Select Save. 	n/a
Edit endpoint binding mode	 a. Select the checkbox for the endpoint. b. Select Actions > Edit endpoint binding mode. c. Update the binding mode as required. d. Select Save changes. 	 a. Select the endpoint name to display the details. b. Select Edit binding mode. c. Update the binding mode as required. d. Select Save changes.
Edit endpoint certificate	 a. Select the checkbox for the endpoint. b. Select Actions > Edit endpoint certificate. c. Upload or generate a new custom certificate or begin using the global S3 and Swift certificate, as required. d. Select Save changes. 	 a. Select the endpoint name to display the details. b. Select the Certificate tab. c. Select Edit certificate. d. Upload or generate a new custom certificate or begin using the global S3 and Swift certificate, as required. e. Select Save changes.
Edit tenant access	 a. Select the checkbox for the endpoint. b. Select Actions > Edit tenant access. c. Choose a different access option, select or remove tenants from the list, or do both. d. Select Save changes. 	 a. Select the endpoint name to display the details. b. Select the Tenant access tab. c. Select Edit tenant access. d. Choose a different access option, select or remove tenants from the list, or do both. e. Select Save changes.

Remove load balancer endpoints

You can remove one or more endpoints using the **Actions** menu, or you can remove a single endpoint from the details page.



To prevent client disruptions, update any affected S3 or Swift client applications before you remove a load balancer endpoint. Update each client to connect using a port assigned to another load balancer endpoint. Be sure to update any required certificate information as well.



If you lose access to Grid Manager while removing a management interface endpoint, update the URL.

- To remove one or more endpoints:
 - a. From the Load balancer page, select the checkbox for each endpoint you want to remove.
 - b. Select **Actions** > **Remove**.
 - c. Select OK.
- To remove one endpoint from the details page:
 - a. From the Load balancer page. select the endpoint name.
 - b. Select **Remove** on the details page.
 - c. Select OK.

Configure S3 endpoint domain names

To support S3 virtual-hosted-style requests, you must use the Grid Manager to configure the list of S3 endpoint domain names that S3 clients connect to.



Using an IP address for an endpoint domain name is unsupported. Future releases will prevent this configuration.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.
- You have confirmed that a grid upgrade is not in progress.



Don't make any changes to the domain name configuration when a grid upgrade is in progress.

About this task

To enable clients to use S3 endpoint domain names, you must do all of the following:

- Use the Grid Manager to add the S3 endpoint domain names to the StorageGRID system.
- Ensure that the certificate the client uses for HTTPS connections to StorageGRID is signed for all domain names that the client requires.

For example, if the endpoint is s3.company.com, you must ensure that the certificate used for HTTPS connections includes the s3.company.com endpoint and the endpoint's wildcard Subject Alternative Name (SAN): *.s3.company.com.

• Configure the DNS server used by the client. Include DNS records for the IP addresses that clients use to make connections, and ensure that the records reference all required S3 endpoint domain names, including any wildcard names.



Clients can connect to StorageGRID using the IP address of a Gateway Node, an Admin Node, or a Storage Node, or by connecting to the virtual IP address of a high availability group. You should understand how client applications connect to the grid so you include the correct IP addresses in the DNS records.

Clients that use HTTPS connections (recommended) to the grid can use either of these certificates:

- Clients that connect to a load balancer endpoint can use a custom certificate for that endpoint. Each load balancer endpoint can be configured to recognize different S3 endpoint domain names.
- Clients that connect to a load balancer endpoint or directly to a Storage Node can customize the global S3 and Swift API certificate to include all required S3 endpoint domain names.



If you don't add S3 endpoint domain names and the list is empty, support for S3 virtual-hostedstyle requests is disabled.

Add an S3 endpoint domain name

Steps

- 1. Select CONFIGURATION > Network > S3 endpoint domain names.
- 2. Enter the domain name in the **Domain name 1** field. Select **Add another domain name** to add more domain names.
- 3. Select Save.
- 4. Ensure that the server certificates that clients use match the required S3 endpoint domain names.
 - If clients connect to a load balancer endpoint that uses its own certificate, update the certificate associated with the endpoint.
 - If clients connect to a load balancer endpoint that uses the global S3 and Swift API certificate or directly to Storage Nodes, update the global S3 and Swift API certificate.
- 5. Add the DNS records required to ensure that endpoint domain name requests can be resolved.

Result

Now, when clients use the endpoint *bucket.s3.company.com*, the DNS server resolves to the correct endpoint and the certificate authenticates the endpoint as expected.

Rename an S3 endpoint domain name

If you change a name used by S3 applications, virtual-hosted-style requests will fail.

Steps

1. Select CONFIGURATION > Network > S3 endpoint domain names.

- 2. Select the domain name field you want to edit and make the necessary changes.
- 3. Select Save.
- 4. Select Yes to confirm your change.

Delete an S3 endpoint domain name

If you remove a name used by S3 applications, virtual-hosted-style requests will fail.

Steps

- 1. Select CONFIGURATION > Network > S3 endpoint domain names.
- 2. Select the delete icon \times next to the domain name.
- 3. Select **Yes** to confirm the deletion.

Related information

- Use S3 REST API
- View IP addresses
- Configure high availability groups

Summary: IP addresses and ports for client connections

To store or retrieve objects, S3 and Swift client applications connect to the Load Balancer service, which is included on all Admin Nodes and Gateway Nodes, or to the Local Distribution Router (LDR) service, which is included on all Storage Nodes.

Client applications can connect to StorageGRID using the IP address of a grid node and the port number of the service on that node. Optionally, you can create high availability (HA) groups of load-balancing nodes to provide highly available connections that use virtual IP (VIP) addresses. If you want to connect to StorageGRID using a fully qualified domain name (FQDN) instead of an IP or VIP address, you can configure DNS entries.

This table summarizes the different ways that clients can connect to StorageGRID and the IP addresses and ports that are used for each type of connection. If you have already created load balancer endpoints and high availability (HA) groups, see Where to find IP addresses to locate these values in the Grid Manager.

Where connection is made	Service that client connects to	IP address	Port
HA group	Load Balancer	Virtual IP address of an HA group	Port assigned to the load balancer endpoint
Admin Node	Load Balancer	IP address of the Admin Node	Port assigned to the load balancer endpoint
Gateway Node	Load Balancer	IP address of the Gateway Node	Port assigned to the load balancer endpoint
Storage Node	LDR	IP address of Storage Node	Default S3 ports: • HTTPS: 18082 • HTTP: 18084 Default Swift ports: • HTTPS: 18083 • HTTP:18085

Example URLs

To connect a client application to the Load Balancer endpoint of an HA group of Gateway Nodes, use a URL structured as shown below:

https://VIP-of-HA-group:LB-endpoint-port

For example, if the virtual IP address of the HA group is 192.0.2.5 and the port number of the load balancer endpoint is 10443, then an application could use the following URL to connect to StorageGRID:

https://192.0.2.5:10443

Where to find IP addresses

- 1. Sign in to the Grid Manager using a supported web browser.
- 2. To find the IP address of a grid node:
 - a. Select NODES.
 - b. Select the Admin Node, Gateway Node, or Storage Node to which you want to connect.
 - c. Select the **Overview** tab.
 - d. In the Node Information section, note the IP addresses for the node.
 - e. Select **Show more** to view IPv6 addresses and interface mappings.

You can establish connections from client applications to any of the IP addresses in the list:

- eth0: Grid Network
- eth1: Admin Network (optional)
- eth2: Client Network (optional)



If you are viewing an Admin Node or a Gateway Node and it is the active node in a high availability group, the virtual IP address of the HA group is shown on eth2.

3. To find the virtual IP address of a high availability group:

a. Select CONFIGURATION > Network > High availability groups.

- b. In the table, note the virtual IP address of the HA group.
- 4. To find the port number of a Load Balancer endpoint:

a. Select CONFIGURATION > Network > Load balancer endpoints.

b. Note the port number for the endpoint you want to use.



If the port number is 80 or 443, the endpoint is configured only on Gateway Nodes, because those ports are reserved on Admin Nodes. All other ports are configured on both Gateway Nodes and Admin Nodes.

- c. Select the name of the endpoint from the table.
- d. Confirm that the Client type (S3 or Swift) matches the client application that will use the endpoint.

Manage networks and connections

Configure network settings: Overview

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

Configure VLAN interfaces

You can create virtual LAN (VLAN) interfaces to isolate and partition traffic for security, flexibility, and performance. Each VLAN interface is associated with one or more parent interfaces on Admin Nodes and Gateway Nodes. You can use VLAN interfaces in HA groups and in load balancer endpoints to segregate client or admin traffic by application or tenant.

Traffic classification policies

You can use traffic classification policies to identify and handle different types of network traffic, including traffic related to specific buckets, tenants, client subnets, or load balancer endpoints. These policies can assist with traffic limiting and monitoring.

Guidelines for StorageGRID networks

You can use the Grid Manager to configure and manage StorageGRID networks and connections.

See Configure S3 and Swift client connections to learn how to connect S3 or Swift clients.

Default StorageGRID networks

By default, StorageGRID supports three network interfaces per grid node, allowing you to configure the networking for each individual grid node to match your security and access requirements.

For more information about network topology, see Networking guidelines.

Grid Network

Required. The Grid Network is used for all internal StorageGRID traffic. It provides connectivity between all nodes in the grid, across all sites and subnets.

Admin Network

Optional. The Admin Network is typically used for system administration and maintenance. It can also be used for client protocol access. The Admin Network is typically a private network and does not need to be routable between sites.

Client Network

Optional. The Client Network is an open network typically used to provide access to S3 and Swift client applications, so the Grid Network can be isolated and secured. The Client Network can communicate with any subnet reachable through the local gateway.

Guidelines

- Each StorageGRID node requires a dedicated network interface, IP address, subnet mask, and gateway for each network it is assigned to.
- A grid node can't have more than one interface on a network.
- A single gateway, per network, per grid node is supported, and it must be on the same subnet as the node. You can implement more complex routing in the gateway, if required.
- On each node, each network maps to a specific network interface.

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

- If the node is connected to a StorageGRID appliance, specific ports are used for each network. For details, see the installation instructions for your appliance.
- The default route is generated automatically, per node. If eth2 is enabled, then 0.0.0.0/0 uses the Client Network on eth2. If eth2 is not enabled, then 0.0.0.0/0 uses the Grid Network on eth0.
- The Client Network does not become operational until the grid node has joined the grid
- The Admin Network can be configured during grid node deployment to allow access to the installation user interface before the grid is fully installed.

Optional interfaces

Optionally, you can add extra interfaces to a node. For example, you might want to add a trunk interface to an Admin or Gateway Node, so you can use VLAN interfaces to segregate the traffic belonging to different applications or tenants. Or, you might want to add an access interface to use in a high availability (HA) group.

To add trunk or access interfaces, see the following:

- VMware (after installing the node): VMware: Add trunk or access interfaces to a node
 - Red Hat Enterprise Linux (before installing the node): Create node configuration files
 - Ubuntu or Debian (before installing the node): Create node configuration files
 - RHEL, Ubuntu, or Debian (after installing the node): Linux: Add trunk or access interfaces to a node

View IP addresses

You can view the IP address for each grid node in your StorageGRID system. You can then use this IP address to log in to the grid node at the command line and perform various maintenance procedures.

Before you begin

You are signed in to the Grid Manager using a supported web browser.

About this task

For information about changing IP addresses, see Configure IP addresses.

Steps

- 1. Select NODES > grid node > Overview.
- 2. Select **Show more** to the right of the IP Addresses title.

DC2-SGA-010-096-106-021 (Storage Node)

The IP addresses for that grid node are listed in a table.

DC2-SGA-010	J-096-106-021 (Storage Node)	× X
Overview	Hardware Network Storage Obje	cts ILM Tasks
Node informatio	n 🕜	
Name:	DC2-SGA-010-096-106-021	
Гуре:	Storage Node	
D:	f0890e03-4c72-401f-ae92-245511a38e51	
Connection state:	Connected	
storage used:	Object data	7%
	Object metadata	5%
Software version:	11.6.0 (build 20210915.1941.afce2d9)	
P addresses:	10.96.106.21 - eth0 (Grid Network)	
	Hide additional IP addresses 🔨	
	Interface 🗢	IP address 🗢
	eth0 (Grid Network)	10.96.106.21
	eth0 (Grid Network)	fe80::2a0:98ff:fe64:6582
	hic2	10.96.106.21
	hic4	10.96.106.21
	mtc2	169.254.0.1
Alerts		
Alert name 🔶		Severity 🤪 💠 Time triggered 💠 Current values
ILM placement u	nachievable 🗹	Major 2 hours ago
A placement inst	truction in an ILM rule cannot be achieved for ce	ertain objects.

Configure VLAN interfaces

You can create virtual LAN (VLAN) interfaces on Admin Nodes and Gateway Nodes and use them in HA groups and load balancer endpoints to isolate and partition traffic for security, flexibility, and performance.

Considerations for VLAN interfaces

- You create a VLAN interface by entering a VLAN ID and choosing a parent interface on one or more nodes.
- A parent interface must be configured as a trunk interface at the switch.
- A parent interface can be the Grid Network (eth0), the Client Network (eth2), or an additional trunk interface for the VM or bare-metal host (for example, ens256).
- For each VLAN interface, you can select only one parent interface for a given node. For example, you can't use both the Grid Network interface and the Client Network interface on the same Gateway Node as the parent interface for the same VLAN.
- If the VLAN interface is for Admin Node traffic, which includes traffic related to the Grid Manager and the Tenant Manager, select interfaces on Admin Nodes only.
- If the VLAN interface is for S3 or Swift client traffic, select interfaces on either Admin Nodes or Gateway Nodes.
- If you need to add trunk interfaces, see the following for details:
 - VMware (after installing the node): VMware: Add trunk or access interfaces to a node
 - RHEL (before installing the node): Create node configuration files
 - Ubuntu or Debian (before installing the node): Create node configuration files
 - RHEL, Ubuntu, or Debian (after installing the node): Linux: Add trunk or access interfaces to a node

Create a VLAN interface

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.
- A trunk interface has been configured in the network and attached to the VM or Linux node. You know the name of the trunk interface.
- You know the ID of the VLAN you are configuring.

About this task

Your network administrator might have configured one or more trunk interfaces and one or more VLANs to segregate the client or admin traffic belonging to different applications or tenants. Each VLAN is identified by a numeric ID or tag. For example, your network might use VLAN 100 for FabricPool traffic and VLAN 200 for an archive application.

You can use the Grid Manager to create VLAN interfaces that allow clients to access StorageGRID on a specific VLAN. When you create VLAN interfaces, you specify the VLAN ID and select parent (trunk) interfaces on one or more nodes.

Access the wizard

Steps

- 1. Select CONFIGURATION > Network > VLAN interfaces.
- 2. Select Create.

Enter details for the VLAN interfaces

Steps

1. Specify the ID of the VLAN in your network. You can enter any value between 1 and 4094.

VLAN IDs don't need to be unique. For example, you might use VLAN ID 200 for admin traffic at one site and the same VLAN ID for client traffic at another site. You can create separate VLAN interfaces with different sets of parent interfaces at each site. However, two VLAN interfaces with the same ID can't share the same interface on a node.

If you specify an ID that has already been used, a message appears.

- 2. Optionally, enter a short description for the VLAN interface.
- 3. Select Continue.

Choose parent interfaces

The table lists the available interfaces for all Admin Nodes and Gateway Nodes at each site in your grid. Admin Network (eth1) interfaces can't be used as parent interfaces and aren't shown.

Steps

1. Select one or more parent interfaces to attach this VLAN to.

For example, you might want to attach a VLAN to the Client Network (eth2) interface for a Gateway Node and an Admin Node.

earch			Q			
	Site 😧 💠	Node name 🔞 🖨	Interface 😧 💠	Description 🔗 🛟	Node type 💡 💲	Attached VLANs 🥝
	Data Center 2	DC2-ADM1	eth0	Grid Network	Non-primary Admin	_
~	Data Center 2	DC2-ADM1	eth2	Client Network	Non-primary Admin	
	Data Center 1	DC1-G1	eth0	Grid Network	Gateway	—
~	Data Center 1	DC1-G1	eth2	Client Network	Gateway	
	Data Center 1	DC1-ADM1	eth0	Grid Network	Primary Admin	-

2. Select Continue.

Confirm the settings

Steps

- 1. Review the configuration and make any changes.
 - If you need to change the VLAN ID or description, select Enter VLAN details at the top of the page.
 - If you need to change a parent interface, select **Choose parent interfaces** at the top of the page or select **Previous**.
 - If you need to remove a parent interface, select the trash can
- 2. Select Save.
- 3. Wait up to 5 minutes for the new interface to appear as a selection on the High availability groups page and to be listed in the **Network interfaces** table for the node (**NODES** > *parent interface node* > **Network**).

Edit a VLAN interface

When you edit a VLAN interface, you can make the following types of changes:

- Change the VLAN ID or description.
- Add or remove parent interfaces.

For example, you might want to remove a parent interface from a VLAN interface if you plan to decommission the associated node.

Note the following:

- You can't change a VLAN ID if the VLAN interface is used in an HA group.
- You can't remove a parent interface if that parent interface is used in an HA group.

For example, suppose VLAN 200 is attached to parent interfaces on Nodes A and B. If an HA group uses the VLAN 200 interface for Node A and the eth2 interface for Node B, you can remove the unused parent interface for Node B, but you can't remove the used parent interface for Node A.

Steps

- 1. Select CONFIGURATION > Network > VLAN interfaces.
- 2. Select the checkbox for the VLAN interface you want to edit. Then, select Actions > Edit.
- 3. Optionally, update the VLAN ID or the description. Then, select Continue.

You can't update a VLAN ID if the VLAN is used in an HA group.

- 4. Optionally, select or clear the checkboxes to add parent interfaces or to remove unused interfaces. Then, select **Continue**.
- 5. Review the configuration and make any changes.
- 6. Select Save.

Remove a VLAN interface

You can remove one or more VLAN interfaces.

You can't remove a VLAN interface if it is currently used in an HA group. You must remove the VLAN interface from the HA group before you can remove it.

To avoid any disruptions in client traffic, consider doing one of the following:

- Add a new VLAN interface to the HA group before removing this VLAN interface.
- Create a new HA group that does not use this VLAN interface.
- If the VLAN interface you want to remove is currently the active interface, edit the HA group. Move the VLAN interface you want to remove to the bottom of the priority list. Wait until communication is established on the new primary interface and then remove the old interface from the HA group. Finally, delete the VLAN interface on that node.

Steps

- 1. Select CONFIGURATION > Network > VLAN interfaces.
- 2. Select the checkbox for each VLAN interface you want to remove. Then, select Actions > Delete.
- 3. Select **Yes** to confirm your selection.

All VLAN interfaces you selected are removed. A green success banner appears on the VLAN interfaces page.

Manage traffic classification policies

Manage traffic classification policies: Overview

To enhance your quality-of-service (QoS) offerings, you can create traffic classification policies to identify and monitor different types of network traffic. These policies can assist with traffic limiting and monitoring.

Traffic classification policies are applied to endpoints on the StorageGRID Load Balancer service for Gateway Nodes and Admin Nodes. To create traffic classification policies, you must have already created load balancer endpoints.

Matching rules

Each traffic classification policy contains one or more matching rules to identify the network traffic related to one or more of the following entities:

- Buckets
- Subnet
- Tenant
- · Load balancer endpoints

StorageGRID monitors traffic that matches any rule within the policy according to the objectives of the rule. Any traffic that matches any rule for a policy is handled by that policy. Conversely, you can set rules to match all traffic except a specified entity.

Traffic limiting

Optionally, you can add the following limit types to a policy:

- Aggregate bandwidth
- Per-request bandwidth
- Concurrent requests

Request rate

Limit values are enforced on a per load balancer basis. If traffic is distributed simultaneously across multiple load balancers, the total maximum rates are a multiple of the rate limits you specify.



You can create policies to limit aggregate bandwidth or to limit per-request bandwidth. However, StorageGRID can't limit both types of bandwidth at the same time. Aggregate bandwidth limits might impose an additional minor performance impact on non-limited traffic.

For aggregate or per-request bandwidth limits, the requests stream in or out at the rate you set. StorageGRID can only enforce one speed, so the most specific policy match, by matcher type, is the one enforced. The bandwidth consumed by the the request does not count against other less specific matching policies containing aggregate bandwidth limit policies. For all other limit types, client requests are delayed by 250 milliseconds and receive a 503 Slow Down response for requests that exceed any matching policy limit.

In the Grid Manager, you can view traffic charts and verify that the polices are enforcing the traffic limits you expect.

Use traffic classification policies with SLAs

You can use traffic classification policies in conjunction with capacity limits and data protection to enforce service-level agreements (SLAs) that provide specifics for capacity, data protection, and performance.

The following example shows three tiers of an SLA. You can create traffic classification policies to achieve the performance objectives of each SLA tier.

Service Level Tier	Capacity	Data Protection	Maximum performance allowed	Cost
Gold	1 PB storage allowed	3 copy ILM rule	25 K requests/sec 5 GB/sec (40 Gbps) bandwidth	\$\$\$ per month
Silver	250 TB storage allowed	2 copy ILM rule	10 K requests/sec 1.25 GB/sec (10 Gbps) bandwidth	\$\$ per month
Bronze	100 TB storage allowed	2 copy ILM rule	5 K requests/sec 1 GB/sec (8 Gbps) bandwidth	\$ per month

Create traffic classification policies

You can create traffic classification policies if you want to monitor, and optionally limit network traffic by bucket, bucket regex, CIDR, load balancer endpoint, or tenant. Optionally, you can set limits for a policy based on bandwidth, the number of concurrent requests, or the request rate.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.
- You have created any load balancer endpoints you want to match.
- · You have created any tenants you want to match.

Steps

- 1. Select CONFIGURATION > Network > Traffic classification.
- 2. Select Create.
- 3. Enter a name and a description (optional) for the policy and select **Continue**.

For example, describe what this traffic classification policy applies to and what it will limit.

4. Select **Add rule** and specify the following details to create one or more matching rules for the policy. Any policy that you create should have at least one matching rule. Select **Continue**.

Field	Description
Туре	Select the types of traffic that the matching rule applies to. Traffic types are bucket, bucket regex, CIDR, load balancer endpoint, and tenant.
Match value	 Enter the value that matches the selected Type. Bucket: Enter one or more bucket names
	 Bucket regex: Enter one or more regular expressions used to match a set of bucket names.
	The regular expression is unanchored. Use the ^ anchor to match at the beginning of the bucket name, and use the \$ anchor to match at the end of the name. Regular expression matching supports a subset of PCRE (Perl compatible regular expression) syntax.
	 CIDR: Enter one or more IPv4 subnets, in CIDR notation, that matches the desired subnet.
	 Load balancer endpoint: Select an endpoint name. These are the load balancer endpoints you defined on the Configure load balancer endpoints.
	 Tenant: Tenant matching uses the access key ID. If the request does not contain an access key ID (for example, anonymous access), then the ownership of the bucket accessed is used to determine the tenant.

Field	Description
Inverse match	If you want to match all network traffic <i>except</i> traffic consistent with the Type and Match Value just defined, select the Inverse match checkbox. Otherwise, leave the checkbox cleared.
	For example, if you want this policy to apply to all but one of the load balancer endpoints, specify the load balancer endpoint to be excluded, and select Inverse match .
	For a policy containing multiple matchers where at least one is an inverse matcher, be careful not to create a policy that matches all requests.

5. Optionally, select **Add a limit** and select the following details to add one or more limits to control the network traffic matched by a rule.



StorageGRID collects metrics even if you don't add any limits, so you can understand traffic trends.

Field	Description
Туре	The type of limit you want to apply to the network traffic matched by the rule. For example, you can limit bandwidth or request rate.
	Note : You can create policies to limit aggregate bandwidth or to limit per- request bandwidth. However, StorageGRID can't limit both types of bandwidth at the same time. When aggregate bandwidth is in use, per-request bandwidth is unavailable. Conversely, when per-request bandwidth is in use, aggregate bandwidth is unavailable. Aggregate bandwidth limits might impose an additional minor performance impact on non-limited traffic.
	For bandwidth limits, StorageGRID applies the policy that best matches the type of limit set. For example, if you have a policy that limits traffic in only one direction, then traffic in the opposite direction will be unlimited, even if there is traffic that matches additional policies that have bandwidth limits. StorageGRID implements "best" matches for bandwidth limits in the following order:
	• Exact IP address (/32 mask)
	Exact bucket name
	Bucket regex
	• Tenant
	• Endpoint
	Non-exact CIDR matches (not /32)
	Inverse matches
Applies to	Whether this limit applies to client read requests (GET or HEAD) or write requests (PUT, POST, or DELETE).

Field	Description
Value	The value that network traffic will be limited to, based on the Unit you select. For example, enter 10 and select MiB/s to prevent the network traffic matched by this rule from exceeding 10 MiB/s.
	Note : Depending on the units setting, the available units will be either binary (for example, GiB) or decimal (for example, GB). To change the units setting, select the user drop-down in the upper right of the Grid Manager, then select User Preferences .
Unit	The unit that describes the value you entered.

For example, if you want to create a 40 GB/s bandwidth limit for an SLA tier, create two Aggregate bandwidth limits: GET/HEAD at 40 GB/s and PUT/POST/DELETE at 40 GB/s.

- 6. Select Continue.
- 7. Read and review the Traffic classification policy. Use the **Previous** button to go back and make changes as required. When you are satisfied with the policy, select **Save and continue**.

S3 and Swift client traffic is now handled according to the traffic classification policy.

After you finish

View network traffic metrics to verify that the polices are enforcing the traffic limits you expect.

Edit traffic classification policy

You can edit a traffic classification policy to change its name or description, or to create, edit, or delete any rules or limits for the policy.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.

Steps

1. Select **CONFIGURATION > Network > Traffic classification**.

The Traffic classification policies page appears and the existing policies are listed in a table.

2. Edit the policy using the Actions menu or the details page. See create traffic classification policies for what to enter.

Actions menu

- a. Select the checkbox for the policy.
- b. Select Actions > Edit.

Details page

- a. Select the policy name.
- b. Select the **Edit** button beside the policy name.
- 3. For the Enter policy name step, optionally edit the policy name or description, and select Continue.
- 4. For the Add matching rules step, optionally add a rule or edit the **Type** and **Match value** of the existing rule, and select **Continue**.
- 5. For the Set limits step, optionally add, edit, or delete a limit, and select Continue.
- 6. Review the updated policy, and select Save and continue.

The changes you made to the policy are saved, and network traffic is now handled according to the traffic classification policies. You can view traffic charts and verify that the polices are enforcing the traffic limits you expect.

Delete a traffic classification policy

You can delete a traffic classification policy if you no longer need it. Make sure you delete the right policy because a policy can't be retrieved when deleted.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.

Steps

1. Select CONFIGURATION > Network > Traffic classification.

The Traffic classification policies page appears with the existing policies listed in a table.

2. Delete the policy using the Actions menu or the details page.

Actions menu

- a. Select the checkbox for the policy.
- b. Select Actions > Remove.

Policy details page

- a. Select the policy name.
- b. Select the **Remove** button beside the policy name.
- 3. Select **Yes** to confirm that you want to delete the policy.

The policy is deleted.

View network traffic metrics

You can monitor network traffic by viewing the graphs that are available from the Traffic classification policies page.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access or Tenant accounts permission.

About this task

For any existing traffic classification policy, you can view metrics for the load balancer service to determine if the policy is successfully limiting traffic across the network. The data in the graphs can help you determine if you need to adjust the policy.

Even if no limits are set for a traffic classification policy, metrics are collected and the graphs provide useful information for understanding traffic trends.

Steps

1. Select CONFIGURATION > Network > Traffic classification.

The Traffic classification policies page appears, and the existing policies are listed in the table.

- 2. Select the traffic classification policy name for which you want to view metrics.
- 3. Select the Metrics tab.

The traffic classification policy graphs appear. The graphs display metrics only for the traffic that matches the selected policy.

The following graphs are included on the page.

 Request rate: This graph provides the amount of bandwidth matching this policy handled by all load balancers. Received data includes request headers for all requests and body data size for responses that have body data. Sent includes response headers for all requests and response body data size for requests that include body data in the response.



When requests are complete, this chart only shows bandwidth usage. For slow or large object requests the actual instantaneous bandwidth might differ from the values reported in this graph.

- Error response rate: This graph provides an approximate rate at which requests matching this policy are returning errors (HTTP status code >= 400) to clients.
- Average request duration (non-error): This graph provides an average duration of successful requests matching this policy.
- Policy bandwidth usage: This graph provides the amount of bandwidth matching this policy handled by all load balancers. Received data includes request headers for all requests and body data size for responses that have body data. Sent includes response headers for all requests and response body data size for requests that include body data in the response.
- 4. Position the cursor over a line graph to see a pop-up of values on a specific part of the graph.
- 5. Select **Grafana dashboard** right below the Metrics title to view all the graphs for a policy. In addition to the four graphs from the **Metrics** tab, you can view two more graphs:

- Write request rate by object size: The rate for PUT/POST/DELETE requests matching this policy.
 Positioning on an individual cell shows per second rates. Rates shown in the hover view are truncated to integer counts and might report 0 when there are non-zero requests in the bucket.
- Read request rate by object size: The rate for GET/HEAD requests matching this policy. Positioning on an individual cell shows per second rates. Rates shown in the hover view are truncated to integer counts and might report 0 when there are non-zero requests in the bucket.
- 6. Alternatively, access the graphs from the **SUPPORT** menu.
 - a. Select **SUPPORT > Tools > Metrics**.
 - b. Select Traffic Classification Policy from the Grafana section.
 - c. Select the policy from the menu on the upper left of the page.
 - d. Position the cursor over a graph to see a pop-up that shows the date and time of the sample, object sizes that are aggregated into the count, and the number of requests per second during that time period.

Traffic classification policies are identified by their ID. Policy IDs are listed on the Traffic classification policies page.

7. Analyze the graphs to determine how often the policy is limiting traffic and whether you need to adjust the policy.

Supported ciphers for outgoing TLS connections

The StorageGRID system supports a limited set of cipher suites for Transport Layer Security (TLS) connections to the external systems used for identity federation and Cloud Storage Pools.

Supported versions of TLS

StorageGRID supports TLS 1.2 and TLS 1.3 for connections to external systems used for identity federation and Cloud Storage Pools.

The TLS ciphers that are supported for use with external systems have been selected to ensure compatibility with a range of external systems. The list is larger than the list of ciphers that are supported for use with S3 or Swift client applications. To configure ciphers, go to **CONFIGURATION** > **Security** > **Security settings** and select **TLS and SSH policies**.



TLS configuration options such as protocol versions, ciphers, key exchange algorithms, and MAC algorithms aren't configurable in StorageGRID. Contact your NetApp account representative if you have specific requests about these settings.

Benefits of active, idle, and concurrent HTTP connections

How you configure HTTP connections can impact the performance of the StorageGRID system. Configurations differ depending on whether the HTTP connection is active or idle or you have concurrent multiple connections.

You can identify the performance benefits for the following types of HTTP connections:

• Idle HTTP connections
- Active HTTP connections
- Concurrent HTTP connections

Benefits of keeping idle HTTP connections open

You should keep HTTP connections open even when client applications are idle to allow client applications to perform subsequent transactions over the open connection. Based on system measurements and integration experience, you should keep an idle HTTP connection open for a maximum of 10 minutes. StorageGRID might automatically close an HTTP connection that is kept open and idle for longer than 10 minutes.

Open and idle HTTP connections provide the following benefits:

• Reduced latency from the time that the StorageGRID system determines it has to perform an HTTP transaction to the time that the StorageGRID system can perform the transaction

Reduced latency is the main advantage, especially for the amount of time required to establish TCP/IP and TLS connections.

- Increased data transfer rate by priming the TCP/IP slow-start algorithm with previously performed transfers
- Instantaneous notification of several classes of fault conditions that interrupt connectivity between the client application and the StorageGRID system

Determining how long to keep an idle connection open is a trade-off between the benefits of slow start that is associated with the existing connection and the ideal allocation of the connection to internal system resources.

Benefits of active HTTP connections

For connections directly to Storage Nodes, you should limit the duration of an active HTTP connection to a maximum of 10 minutes, even if the HTTP connection continuously performs transactions.

Determining the maximum duration that a connection should be held open is a trade-off between the benefits of connection persistence and the ideal allocation of the connection to internal system resources.

For client connections to Storage Nodes, limiting active HTTP connections provides the following benefits:

• Enables optimal load balancing across the StorageGRID system.

Over time, an HTTP connection might no longer be optimal as load balancing requirements change. The system performs its best load balancing when client applications establish a separate HTTP connection for each transaction, but this negates the much more valuable gains associated with persistent connections.

- Allows client applications to direct HTTP transactions to LDR services that have available space.
- · Allows maintenance procedures to start.

Some maintenance procedures start only after all the in-progress HTTP connections are complete.

For client connections to the Load Balancer service, limiting the duration of open connections can be useful for allowing some maintenance procedures to start promptly. If the duration of client connections is not limited, it might take several minutes for active connections to be automatically terminated.

Benefits of concurrent HTTP connections

You should keep multiple TCP/IP connections to the StorageGRID system open to allow parallelism, which

increases performance. The optimal number of parallel connections depends on a variety of factors.

Concurrent HTTP connections provide the following benefits:

Reduced latency

Transactions can start immediately instead of waiting for other transactions to be completed.

Increased throughput

The StorageGRID system can perform parallel transactions and increase aggregate transaction throughput.

Client applications should establish multiple HTTP connections. When a client application has to perform a transaction, it can select and immediately use any established connection that is not currently processing a transaction.

Each StorageGRID system's topology has different peak throughput for concurrent transactions and connections before performance begins to degrade. Peak throughput depends on factors such as computing resources, network resources, storage resources, and WAN links. The number of servers and services and the number of applications that the StorageGRID system supports are also factors.

StorageGRID systems often support multiple client applications. You should keep this in mind when you determine the maximum number of concurrent connections used by a client application. If the client application consists of multiple software entities that each establish connections to the StorageGRID system, you should add up all the connections across the entities. You might have to adjust the maximum number of concurrent connections in the following situations:

- The StorageGRID system's topology affects the maximum number of concurrent transactions and connections that the system can support.
- Client applications that interact with the StorageGRID system over a network with limited bandwidth might have to reduce the degree of concurrency to ensure that individual transactions are completed in a reasonable time.
- When many client applications share the StorageGRID system, you might have to reduce the degree of concurrency to avoid exceeding the limits of the system.

Separation of HTTP connection pools for read and write operations

You can use separate pools of HTTP connections for read and write operations and control how much of a pool to use for each. Separate pools of HTTP connections enable you to better control transactions and balance loads.

Client applications can create loads that are retrieve-dominant (read) or store-dominant (write). With separate pools of HTTP connections for read and write transactions, you can adjust how much of each pool to dedicate for read or write transactions.

Manage link costs

Link costs let you prioritize which data center site provides a requested service when two or more data center sites exist. You can adjust link costs to reflect latency between sites.

What are link costs?

- Link costs are used to prioritize which object copy is used to fulfill object retrievals.
- Link costs are used by the Grid Management API and the Tenant Management API to determine which internal StorageGRID services to use.
- Link costs are used by the Load Balancer service on Admin Nodes and Gateway Nodes to direct client connections. See Considerations for load balancing.

The diagram shows a three site grid that has link costs configured between sites:



• The Load Balancer service on Admin Nodes and Gateway Nodes equally distributes client connections to all Storage Nodes at the same data center site and to any data center sites with a link cost of 0.

In the example, a Gateway Node at data center site 1 (DC1) equally distributes client connections to Storage Nodes at DC1 and to Storage Nodes at DC2. A Gateway Node at DC3 sends client connections only to Storage Nodes at DC3.

• When retrieving an object that exists as multiple replicated copies, StorageGRID retrieves the copy at the data center that has the lowest link cost.

In the example, if a client application at DC2 retrieves an object that is stored both at DC1 and DC3, the object is retrieved from DC1, because the link cost from DC1 to DC2 is 0, which is lower than the link cost from DC3 to DC2 (25).

Link costs are arbitrary relative numbers with no specific unit of measure. For example, a link cost of 50 is used less preferentially than a link cost of 25. The table shows commonly used link costs.

Link	Link cost	Notes
Between physical data center sites	25 (default)	Data centers connected by a WAN link.

Link	Link cost	Notes
Between logical data center sites at the same physical location	0	Logical data centers in the same physical building or campus connected by a LAN.

Update link costs

You can update the link costs between data center sites to reflect latency between sites.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Grid topology page configuration permission.

Steps

1. Select **SUPPORT > Other > Link cost**.

Site Names (1 - 3 of	3)			
Site ID		Site Name	Action	ıs
10		Data Center 1	0	
20		Data Center 2	1	
30		Data Center 3	0	
Show 50 V Records Po Link Costs	er Page Re	fresh	Previous	n 1 x Next
		Link Destination		

2. Select a site under Link Source and enter a cost value between 0 and 100 under Link Destination.

You can't change the link cost if the source is the same as the destination.

To cancel changes, select 🕥 Revert.

3. Select Apply Changes.

Use AutoSupport

Use AutoSupport: Overview

The AutoSupport feature enables StorageGRID to send health and status packages to NetApp technical support.

Using AutoSupport can significantly speed up problem determination and resolution. Technical support can also monitor the storage needs of your system and help you determine if you need to add new nodes or sites. Optionally, you can configure AutoSupport packages to be sent to one additional destination.

StorageGRID has two types of AutoSupport:

StorageGRID AutoSupport

Reports StorageGRID software issues. Enabled by default when you first install StorageGRID. You can change the default AutoSupport configuration if needed.



If StorageGRID AutoSupport is not enabled, a message appears on the Grid Manager dashboard. The message includes a link to the AutoSupport configuration page. If you close the message, it will not appear again until your browser cache is cleared, even if AutoSupport remains disabled.

Appliance hardware AutoSupport

Reports StorageGRID appliance issues. You must configure hardware AutoSupport on each appliance.

What is Active IQ?

Active IQ is a cloud-based digital advisor that leverages predictive analytics and community wisdom from NetApp's installed base. Its continuous risk assessments, predictive alerts, prescriptive guidance, and automated actions help you prevent problems before they occur, leading to improved system health and higher system availability.

If you want to use the Active IQ dashboards and functionality on the NetApp Support Site, you must enable AutoSupport.

Active IQ Digital Advisor Documentation

Information included in AutoSupport package

An AutoSupport package contains the following XML files and details.

File name	Fields	Description
AUTOSUPPORT- HISTORY.XML	AutoSupport Sequence Number Destination for this AutoSupport Trigger Event Status of Delivery Delivery Attempts AutoSupport Subject Delivery URI Last error AutoSupport PUT Filename Time of Generation Autosupport Compressed Size Autosupport Decompressed Size Total Collection Time (ms)	AutoSupport history file
AUTOSUPPORT.XML	Node Protocol to contact support Support URL for HTTP/HTTPS Support Address AutoSupport OnDemand State AutoSupport OnDemand Server URL AutoSupport OnDemand Polling Interval	AutoSupport status file. Provides details of protocol used, technical support URL and address, polling interval, and OnDemand AutoSupport if enabled or disabled.

File name	Fields	Description
BUCKETS.XML	Bucket ID Account ID Build Version Location Constraint Compliance Enabled Compliance Enabled Compliance Configuration S3 Object Lock Enabled S3 Object Lock Configuration Consistency Configuration CORS Enabled CORS Configuration Last Access Time Enabled Policy Enabled Policy Configuration Notifications Enabled Notifications Configuration Cloud Mirror Enabled Cloud Mirror Configuration Search Enabled Search Configuration Swift Read ACL Enabled Swift Read ACL Enabled Swift Write ACL Enabled Swift Write ACL Configuration Bucket Tagging Enabled Bucket Tagging Configuration	Provides configuration details and statistics at the bucket level. Example of bucket configurations include platform services, compliance, and bucket consistency.
GRID- CONFIGURATIONS.XML	Attribute ID Attribute Name Value Index Table ID Table Name	Grid-wide configuration information file. Contains information about grid certificates, metadata reserved space, grid-wide configuration settings (compliance, S3 Object Lock, object compression, alerts, syslog, and ILM configuration), erasure-coding profile details, DNS name, NMS name, and more.
GRID-SPEC.XML	Grid specifications, raw XML	Used for configuring and deploying StorageGRID. Contains grid specifications, NTP server IP, DNS server IP, network topology, and hardware profiles of the nodes.
GRID-TASKS.XML	Node Service Path Attribute ID Attribute name Value Index Table ID Table name	Grid tasks (maintenance procedures) status file. Provides details of the grid's active, terminated, completed, failed, and pending tasks.

File name	Fields	Description
ILM-STATUS.XML	Node Service path Attribute ID Attribute name Value Index Table ID Table name	ILM metrics information file. Contains ILM evaluation rates for each node and grid-wide metrics.
ILM.XML	ILM raw XML	ILM active policy file. Contains details about the active ILM policies, such as storage pool ID, ingest behavior, filters, rules, and description.
LOG.TGZ	n/a	Downloadable log file. Contains bycast- err.log and servermanager.log from each node.
MANIFEST.XML	Collection order AutoSupport content filename for this data Description of this data item Number of bytes collected Time spent collecting Status of this data item Description of the error AutoSupport content type for this data +	Contains AutoSupport metadata and brief descriptions of all AutoSupport XML files.
NMS-ENTITIES.XML	Attribute index Entity OID Node ID Device model ID Device model version Entity name	Group and service entities in the NMS tree. Provides grid topology details. The node can be determined based on the services running on the node.
OBJECTS-STATUS.XML	Node Service path Attribute ID Attribute name Value Index Table ID Table name	Object status, including background scan status, active transfer, transfer rate, total transfers, delete rate, corrupted fragments, lost objects, missing objects, repair attempted, scan rate, estimated scan period, repair completion status, and more.

File name	Fields	Description
SERVER-STATUS.XML	Node Service path Attribute ID Attribute name Value Index Table ID Table name	Server configurations and events file. Contains these details for each node: platform type, operating system, installed memory, available memory, storage connectivity, storage appliance chassis serial number, storage controller failed drive count, compute controller chassis temperature, compute hardware, compute controller serial number, power supply, drive size, drive type, and more.
SERVICE-STATUS.XML	Node Service path Attribute ID Attribute name Value Index Table ID Table name	Service node information file. Contains details such as allocated table space, free table space, Reaper metrics of the database, segment repair duration, repair job duration, auto job restarts, auto job termination, and more.
STORAGE-GRADES.XML	Storage grade ID Storage grade name Storage node ID Storage node path	Storage grade definitions file for each Storage Node.
SUMMARY- ATTRIBUTES.XML	Group OID Group Path Summary attribute ID Summary attribute name Value Index Table ID Table name	High-level system status data that summarizes StorageGRID usage information. Provides details such as name of grid, names of sites, number of Storage Nodes per grid and per site, license type, license capacity and usage, software support terms, and details of S3 and Swift operations.
SYSTEM-ALARMS.XML	Node Service path Severity Alarmed attribute Attribute name Status Value Trigger time Acknowledge time	System level alarms (deprecated) and status data used to indicate abnormal activities or potential problems.

File name	Fields	Description
SYSTEM-ALERTS.XML	Name Severity Node name Alert Status Site name Alert triggered time Alert resolved time Rule ID Node ID Site ID Silenced Other annotations Other labels	Current system alerts that indicate potential problems in the StorageGRID system.
USERAGENTS.XML	User agent Number of days Total HTTP requests Total bytes ingested Total bytes retrieved PUT requests GET requests DELETE requests HEAD requests POST requests OPTIONS requests Average request time (ms) Average QET request time (ms) Average DELETE request time (ms) Average HEAD request time (ms) Average POST request time (ms) Average OPTIONS request time (ms)	Statistics based on the application user agents. For example, the number of PUT/GET/DELETE/HEAD operations per user agent and total bytes size of each operation.
X-HEADER-DATA	X-Netapp-asup-generated-on X-Netapp-asup-hostname X-Netapp-asup-os-version X-Netapp-asup-serial-num X-Netapp-asup-subject X-Netapp-asup-system-id X-Netapp-asup-model-name +	AutoSupport header data.

Configure AutoSupport

By default, the StorageGRID AutoSupport feature is enabled when you first install StorageGRID. However, you must configure hardware AutoSupport on each appliance. As needed, you can change the AutoSupport configuration.

If you want to change the configuration of StorageGRID AutoSupport, make your changes only on the primary Admin Node. You must configure hardware AutoSupport on each appliance.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.
- If you will use HTTPS for sending AutoSupport packages, you have provided outbound internet access to the primary Admin Node, either directly or using a proxy server (inbound connections not required).
- If HTTP is selected on the StorageGRID AutoSupport page, you have configured a proxy server to forward AutoSupport packages as HTTPS. NetApp's AutoSupport servers will reject packages sent using HTTP.

Learn about configuring admin proxy settings.

• If you will use SMTP as the protocol for AutoSupport packages, you have configured an SMTP mail server. The same mail server configuration is used for alarm email notifications (legacy system).

About this task

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

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

Specify the protocol for AutoSupport packages

You can use any of the following protocols for sending AutoSupport packages:

- **HTTPS**: This is the default and recommended setting for new installations. This protocol uses port 443. If you want to enable the AutoSupport on Demand feature, you must use HTTPS.
- **HTTP**: If you select HTTP, you must configure a proxy server to forward AutoSupport packages as HTTPS. NetApp's AutoSupport servers reject packages sent using HTTP. This protocol uses port 80.
- **SMTP**: Use this option if you want AutoSupport packages to be emailed. If you use SMTP as the protocol for AutoSupport packages, you must configure an SMTP mail server on the Legacy Email Setup page (**SUPPORT** > **Alarms (legacy)** > **Legacy email setup**).

The protocol you set is used for sending all types of AutoSupport packages.

Steps

1. Select SUPPORT > Tools > AutoSupport > Settings.

- 2. Select the protocol you want to use to send AutoSupport packages.
- 3. If you selected **HTTPS**, select whether to use a NetApp support certificate (TLS certificate) to secure the connection to the technical support server.
 - **Verify certificate** (default): Ensures that the transmission of AutoSupport packages is secure. The NetApp support certificate is already installed with the StorageGRID software.
 - **Do not verify certificate**: Select this option only when you have a good reason not to use certificate validation, such as when there is a temporary problem with a certificate.
- 4. Select Save. All weekly, user-triggered, and event-triggered packages are sent using the selected protocol.

Disable weekly AutoSupport

By default, the StorageGRID system is configured to send an AutoSupport package to technical support once a week.

To determine when the weekly AutoSupport package will be sent, go to the **AutoSupport** > **Results** tab. In the **Weekly AutoSupport** section, look at the value for **Next Scheduled Time**.

You can disable the automatic sending of weekly AutoSupport packages at any time.

Steps

1. Select SUPPORT > Tools > AutoSupport > Settings.

- 2. Clear the Enable Weekly AutoSupport checkbox.
- 3. Select Save.

Disable event-triggered AutoSupport

By default, the StorageGRID system is configured to send an AutoSupport package to technical support every hour or when an important alert or other significant system event occurs.

You can disable event-triggered AutoSupport at any time.

Steps

- 1. Select SUPPORT > Tools > AutoSupport > Settings.
- 2. Clear the Enable Event-Triggered AutoSupport checkbox.
- 3. Select Save.

Enable AutoSupport on Demand

AutoSupport on Demand can assist in solving issues that technical support is actively working on.

By default, AutoSupport on Demand is disabled. Enabling this feature allows technical support to request that your StorageGRID system send AutoSupport packages automatically. Technical support can also set the polling time interval for AutoSupport on Demand queries.

Technical support can't enable or disable AutoSupport on Demand.

Steps

- 1. Select SUPPORT > Tools > AutoSupport > Settings.
- 2. Select the HTTPS for the protocol.
- 3. Select the Enable Weekly AutoSupport checkbox.

- 4. Select the Enable AutoSupport on Demand checkbox.
- 5. Select Save.

AutoSupport on Demand is enabled, and technical support can send AutoSupport on Demand requests to StorageGRID.

Disable checks for software updates

By default, StorageGRID contacts NetApp to determine if software updates are available for your system. If a StorageGRID hotfix or new version is available, the new version is shown on the StorageGRID Upgrade page.

As required, you can optionally disable the check for software updates. For example, if your system does not have WAN access, you should disable the check to avoid download errors.

Steps

- 1. Select SUPPORT > Tools > AutoSupport > Settings.
- 2. Clear the Check for software updates checkbox.
- 3. Select Save.

Add an additional AutoSupport destination

When you enable AutoSupport, heath and status packages are sent to technical support. You can specify one additional destination for all AutoSupport packages.

To verify or change the protocol used to send AutoSupport packages, see the instructions to specify the protocol for AutoSupport packages.



You can't use the SMTP protocol to send AutoSupport packages to an additional destination.

Steps

- 1. Select SUPPORT > Tools > AutoSupport > Settings.
- 2. Select Enable Additional AutoSupport Destination.
- 3. Specify the following:

Hostname

The server hostname or IP address of an additional AutoSupport destination server.



You can enter only one additional destination.

Port

The port used to connect to an additional AutoSupport destination server. The default is port 80 for HTTP or port 443 for HTTPS.

Certificate validation

Whether a TLS certificate is used to secure the connection to the additional destination.

- Select Verify certificate to use certificate validation.
- Select **Do not verify certificate** to send your AutoSupport packages without certificate validation.

Select this choice only when you have a good reason not to use certificate validation, such as when

there is a temporary problem with a certificate.

- 4. If you selected Verify certificate, do the following:
 - a. Browse to the location of the CA certificate.
 - b. Upload the CA certificate file.

The CA certificate metadata appears.

5. Select Save.

All future weekly, event-triggered, and user-triggered AutoSupport packages will be sent to the additional destination.

Configure AutoSupport for appliances

AutoSupport for appliances reports StorageGRID hardware issues, and StorageGRID AutoSupport reports StorageGRID software issues, with one exception: for the SGF6112, StorageGRID AutoSupport reports both hardware and software issues. You must configure AutoSupport on each appliance except the SGF6112, which does not require additional configuration. AutoSupport is implemented differently for services appliances and storage appliances.

You use SANtricity to enable AutoSupport for each storage appliance. You can configure SANtricity AutoSupport during initial appliance setup or after an appliance has been installed:

• For SG6000 and SG5700 appliances, configure AutoSupport in SANtricity System Manager

AutoSupport packages from E-Series appliances can be included in StorageGRID AutoSupport if you configure AutoSupport delivery by proxy in SANtricity System Manager.

StorageGRID AutoSupport does not report hardware issues, such as DIMM or host interface card (HIC) faults. However, some component failures might trigger hardware alerts. For StorageGRID appliances with a baseboard management controller (BMC), such as the SG100, SG1000, SG6060, or SGF6024, you can configure email and SNMP traps to report hardware failures:

- · Set up email notifications for BMC alerts
- Configure SNMP settings for BMC for the SG6000-CN controller or the SG100 and SG1000 services appliances

Related information

NetApp Support

Manually trigger an AutoSupport package

To assist technical support in troubleshooting issues with your StorageGRID system, you can manually trigger an AutoSupport package to be sent.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You must have the Root access or Other grid configuration permission.

Steps

- 1. Select SUPPORT > Tools > AutoSupport.
- 2. On the Actions tab, select Send User-Triggered AutoSupport.

StorageGRID attempts to send an AutoSupport package to the NetApp Support Site. If the attempt is successful, the **Most Recent Result** and **Last Successful Time** values on the **Results** tab are updated. If there is a problem, the **Most Recent Result** value updates to "Failed," and StorageGRID does not try to send the AutoSupport package again.



After sending an User-triggered AutoSupport package, refresh the AutoSupport page in your browser after 1 minute to access the most recent results.

Troubleshoot AutoSupport packages

If an attempt to send an AutoSupport package fails, the StorageGRID system takes different actions depending on the type of AutoSupport package. You can check the status of AutoSupport packages by selecting **SUPPORT** > **Tools** > **AutoSupport** > **Results**.

When the AutoSupport package fails to send, "Failed" appears on the **Results** tab of the **AutoSupport** page.



If you configured a proxy server to forward AutoSupport packages to NetApp, you should verify that the proxy server configuration settings are correct.

Weekly AutoSupport package failure

If a weekly AutoSupport package fails to send, the StorageGRID system takes the following actions:

- 1. Updates the Most Recent Result attribute to Retrying.
- 2. Attempts to resend the AutoSupport package 15 times every four minutes for one hour.
- 3. After one hour of send failures, updates the Most Recent Result attribute to Failed.
- 4. Attempts to send an AutoSupport package again at the next scheduled time.
- 5. Maintains the regular AutoSupport schedule if the package fails because the NMS service is unavailable, and if a package is sent before seven days pass.
- 6. When the NMS service is available again, sends an AutoSupport package immediately if a package has not been sent for seven days or more.

User-triggered or event-triggered AutoSupport package failure

If a user-triggered or an event-triggered AutoSupport package fails to send, the StorageGRID system takes the following actions:

- Displays an error message if the error is known. For example, if a user selects the SMTP protocol without providing correct email configuration settings, the following error is displayed: AutoSupport packages cannot be sent using SMTP protocol due to incorrect settings on the E-mail Server page.
- 2. Does not attempt to send the package again.
- 3. Logs the error in nms.log.

If a failure occurs and SMTP is the selected protocol, verify that the StorageGRID system's email server is correctly configured and that your email server is running (SUPPORT > Alarms (legacy) > > Legacy Email Setup). The following error message might appear on the AutoSupport page: AutoSupport packages cannot be sent using SMTP protocol due to incorrect settings on the E-mail Server page.

Learn how to configure email server settings.

Correct an AutoSupport package failure

If a failure occurs and SMTP is the selected protocol, verify that the StorageGRID system's email server is correctly configured and that your email server is running. The following error message might appear on the AutoSupport page: AutoSupport packages cannot be sent using SMTP protocol due to incorrect settings on the E-mail Server page.

Send E-Series AutoSupport packages through StorageGRID

You can send E-Series SANtricity System Manager AutoSupport packages to technical support through a StorageGRID Admin Node rather than the storage appliance management port.

See E-Series hardware AutoSupport for more information about using AutoSupport with E-Series appliances.

Before you begin

- You are signed into the Grid Manager using a supported web browser.
- You have the Storage appliance administrator or Root access permission.
- You have configured SANtricity AutoSupport:
 - For SG6000 and SG5700 appliances, configure AutoSupport in SANtricity System Manager



You must have SANtricity firmware 8.70 or higher to access SANtricity System Manager using the Grid Manager.

About this task

E-Series AutoSupport packages contain details of the storage hardware and are more specific than other AutoSupport packages sent by the StorageGRID system.

You can configure a special proxy server address in SANtricity System Manager to transmit AutoSupport packages through a StorageGRID Admin Node without the use of the appliance's management port. AutoSupport packages transmitted in this way are sent by the preferred sender Admin Node, and they use any admin proxy settings that have been configured in the Grid Manager.



This procedure is only for configuring a StorageGRID proxy server for E-Series AutoSupport packages. For additional details on E-Series AutoSupport configuration, see the NetApp E-Series and SANtricity Documentation.

Steps

- 1. In the Grid Manager, select **NODES**.
- 2. From the list of nodes on the left, select the storage appliance node you want to configure.
- 3. Select SANtricity System Manager.

The SANtricity System Manager home page appears.



4. Select SUPPORT > Support center > AutoSupport.

The AutoSupport operations page appears.

				Technical Support
			Chassis s	erial number: 031517000693
				占 NetApp My Support 🖸
	Support Resources	Diagnostics	Auto Support	US/Canada 888.463.8277
	Support Resources	Diagnostics	Autosupport	Other Contacts
oSupport operations			A	utoSupport status: Enabled 😯
AutoSupport proactively support team. Configure AutoSupport Del Connect to the support te Schedule AutoSupport Disp AutoSupport dispatches	ivery Method earn via HTTPS, HTTP or Ma patches are sent daily at 03:06 PM UT	orage array and auto il (SMTP) server deli TC and weekly at 07	omatically sends supp ivery methods. :39 AM UTC on Thurs	ort data ("dispatches") to the sday.
Send AutoSupport Dispatch Automatically sends the	n support team a dispatch to tro	oubleshoot system is	sues without waiting f	or periodic dispatches.
View AutoSupport Log The AutoSupport log pro AutoSupport dispatches.	vides information about statu	s, dispatch history, a	nd errors encountered	I during delivery of
Enable AutoSupport Mainte Enable AutoSupport Mai generating support cases	enance Window ntenance window to allow ma s.	intenance activities	to be performed on the	e storage array without
Disable AutoSupport Mainto Disable AutoSupport Mainto other destructive actions	enance Window intenance window to allow the	e storage array to ge	nerate support cases	on component failures and

5. Select Configure AutoSupport Delivery Method.

The Configure AutoSupport Delivery Method page appears.

Configure AutoSupport Delivery Method	×
Select AutoSupport dispatch delivery method HTTPS HTTP Email 	
HTTPS delivery settings	Show destination address
Connect to support team Directly via Proxy server Host address Host address Tunnel-host Port number 10225 My proxy server requires authentication via Proxy auto-configuration script (PAC) Y	
Save Test Co	nfiguration

6. Select **HTTPS** for the delivery method.



The certificate that enables HTTPS is pre-installed.

- 7. Select via Proxy server.
- 8. Enter tunnel-host for the Host address.

tunnel-host is the special address to use an Admin Node to send E-Series AutoSupport packages.

9. Enter 10225 for the Port number.

10225 is the port number on the StorageGRID proxy server that receives AutoSupport packages from the E-Series controller in the appliance.

10. Select **Test Configuration** to test the routing and configuration of your AutoSupport proxy server.

If correct, a message in a green banner appears: "Your AutoSupport configuration has been verified."

If the test fails, an error message appears in a red banner. Check your StorageGRID DNS settings and

networking, ensure the preferred sender Admin Node can connect to the NetApp Support Site, and try the test again.

11. Select Save.

The configuration is saved, and a confirmation message appears: "AutoSupport delivery method has been configured."

Manage Storage Nodes

Manage Storage Nodes: Overview

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

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

Use Storage options

What is object segmentation?

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

When an object is ingested into the StorageGRID system, the LDR service splits the object into segments, and creates a segment container that lists the header information of all segments as content.



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

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

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

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

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

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

What are storage volume watermarks?

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



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

What is the Soft Read-Only Watermark?

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

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

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

What is the Hard Read-Only Watermark?

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

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

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

i

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

What is the Read-Write Watermark?

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

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

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

View storage volume watermarks

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

Before you begin

- You have completed the upgrade to StorageGRID 11.6 or higher.
- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.

View current watermark settings

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

Steps

- 1. Select SUPPORT > Other > Storage watermarks.
- 2. On the Storage Watermarks page, look at the Use optimized values checkbox.
 - If the checkbox is selected, all three watermarks are optimized for every storage volume on every Storage Node, based on the size of the Storage Node and the relative capacity of the volume.

This is the default and recommended setting. Do not update these values. Optionally, you can View optimized storage watermarks.

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

When you specify custom watermark settings, you must enter values greater than 0.

View optimized storage watermarks

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

1. Select **SUPPORT > Tools > Metrics**.

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

storagegrid storage volume minimum optimized soft readonly watermark

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

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

storagegrid storage volume maximum optimized soft readonly watermark

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

Manage object metadata storage

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

What is object metadata?

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

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

- System metadata, including a unique ID for each object (UUID), the object name, the name of the S3 bucket or Swift container, the tenant account name or ID, the logical size of the object, the date and time the object was first created, and the date and time the object was last modified.
- · Any custom user metadata key-value pairs associated with the object.
- For S3 objects, any object tag key-value pairs associated with the object.
- For replicated object copies, the current storage location of each copy.
- For erasure-coded object copies, the current storage location of each fragment.
- For object copies in a Cloud Storage Pool, the location of the object, including the name of the external bucket and the object's unique identifier.
- For segmented objects and multipart objects, segment identifiers and data sizes.

How is object metadata stored?

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

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



Where is object metadata stored?

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



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

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

Metadata reserved space setting

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

setting is based on:

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

Version used for initial StorageGRID installation	Amount of RAM on Storage Nodes	Default Metadata reserved space setting
11.5 to 11.8	128 GB or more on each Storage Node in the grid	8 TB (8,000 GB)
	Less than 128 GB on any Storage Node in the grid	3 TB (3,000 GB)
11.1 to 11.4	128 GB or more on each Storage Node at any one site	4 TB (4,000 GB)
	Less than 128 GB on any Storage Node at each site	3 TB (3,000 GB)
11.0 or earlier	Any amount	2 TB (2,000 GB)

View Metadata reserved space setting

Follow these steps to view the Metadata reserved space setting for your StorageGRID system.

Steps

- 1. Select CONFIGURATION > System > Storage settings.
- 2. On the Storage settings page, expand the Metadata reserved space section.

For StorageGRID 11.8 or higher, the Metadata reserved space value must be at least 100 GB and no more than 1 PB.

The default setting for a new StorageGRID 11.6 or higher installation in which each Storage Node has 128 GB or more of RAM is 8,000 GB (8 TB).

Actual reserved space for metadata

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

Size of volume 0 for the node	Actual reserved space for metadata
Less than 500 GB (non-production use)	10% of volume 0

Size of volume 0 for the node	Actual reserved space for metadata
500 GB or more or Metadata-only Storage Nodes	 The smaller of these values: Volume 0 Metadata reserved space setting Note: Only one rangedb is required for metadata-only Storage Nodes.

View actual reserved space for metadata

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

Steps

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



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

Example for actual reserved metadata space

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

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

Allowed metadata space

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



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

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

View allowed metadata space

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

Steps

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



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

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

```
storagegrid storage utilization metadata allowed bytes
```

Example for allowed metadata space

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

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

How Storage Nodes of different sizes affect object capacity

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

Consider the following example:

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

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

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



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

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

Related information

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

Increase Metadata Reserved Space setting

You might be able to increase the Metadata Reserved Space system setting if your Storage Nodes meet specific requirements for RAM and available space.

What you'll need

• You are signed in to the Grid Manager using a supported web browser.

• You have the Root Access permission or the Grid Topology Page Configuration and Other Grid Configuration permissions.

About this task

You might be able to manually increase the system-wide Metadata Reserved Space setting up to 8 TB.

You can only increase the value of the system-wide Metadata Reserved Space setting if both of these statements are true:

- The Storage Nodes at any site in your system each have 128 GB or more RAM.
- The Storage Nodes at any site in your system each have sufficient available space on storage volume 0.

Be aware that if you increase this setting, you will simultaneously reduce the space available for object storage on storage volume 0 of all Storage Nodes. For this reason, you might prefer to set the Metadata Reserved Space to a value smaller than 8 TB, based on your expected object metadata requirements.



In general, it is better to use a higher value instead of a lower value. If the Metadata Reserved Space setting is too large, you can decrease it later. In contrast, if you increase the value later, the system might need to move object data to free up space.

For a detailed explanation of how the Metadata Reserved Space setting affects the allowed space for object metadata storage on a particular Storage Node, see Manage object metadata storage.

Steps

- 1. Determine the current Metadata Reserved Space setting.
 - a. Select CONFIGURATION > System > Storage options.
 - b. In the Storage Watermarks section, note the value of Metadata Reserved Space.
- 2. Ensure you have enough available space on storage volume 0 of each Storage Node to increase this value.

a. Select NODES.

- b. Select the first Storage Node in the grid.
- c. Select the Storage tab.
- d. In the Volumes section, locate the /var/local/rangedb/0 entry.
- e. Confirm that the Available value is equal to or greater than difference between the new value you want to use and the current Metadata Reserved Space value.

For example, if the Metadata Reserved Space setting is currently 4 TB and you want to increase it to 6 TB, the Available value must be 2 TB or greater.

- f. Repeat these steps for all Storage Nodes.
 - If one or more Storage Nodes do not have enough available space, the Metadata Reserved Space value cannot be increased. Do not continue with this procedure.
 - If each Storage Node has enough available space on volume 0, go to the next step.
- 3. Ensure you have at least 128 GB of RAM on each Storage Node.
 - a. Select NODES.
 - b. Select the first Storage Node in the grid.
 - c. Select the Hardware tab.

- d. Hover your cursor over the Memory Usage chart. Ensure that **Total Memory** is at least 128 GB.
- e. Repeat these steps for all Storage Nodes.
 - If one or more Storage Nodes do not have enough available total memory, the Metadata Reserved Space value cannot be increased. Do not continue with this procedure.
 - If each Storage Node has at least 128 GB of total memory, go to the next step.
- 4. Update the Metadata Reserved Space setting.
 - a. Select CONFIGURATION > System > Storage options.
 - b. Select the Configuration tab.
 - c. In the Storage Watermarks section, select Metadata Reserved Space.
 - d. Enter the new value.

For example, to enter 8 TB, which is the maximum supported value, enter **800000000000** (8, followed by 12 zeros)

	Updated: 2021-12-10 13:48:23 MST	options
verview		
onfiguration	Object Segmentation	
	Description	Settings
	Segmentation	Enabled
	Maximum Segment Size	100000000
	Storage Watermarks	
	Description	Settings
	Storage Volume Read-Write Watermark Override	0
	Storage Volume Soft Read-Only Watermark Override	0
	Storage Volume Hard Read-Only Watermark Override	0
	Metadata Reserved Space	800000000000

e. Select Apply Changes.

Compress stored objects

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

Before you begin

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

About this task

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



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

Before enabling object compression, be aware of the following:

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

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



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

Steps

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

Storage Node configuration settings

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



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

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

Steps

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

The following tables summarize Storage Node configuration settings.

	n	D
L	υ	Г

Attribute Name	Code	Description
HTTP State	HSTE	 The current state of HTTP for S3, Swift, and other internal StorageGRID traffic: 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	 If selected, the state of the system on restart depends on the state of the LDR > Storage component. If the LDR > Storage component is Read-only on restart, the HTTP interface is also Read-only. If the LDR > Storage component is Online, then HTTP is also Online. Otherwise, the HTTP interface remains in the Offline state. If not selected, the HTTP interface remains Offline until explicitly enabled.

LDR > Data Store

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

LDR > Storage

Attribute Name	Code	Description
Storage State — Desired	SSDS	 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. However, the LDR service continues to accept S3 or ILM driven purge and delete requests. 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 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

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

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

LDR > Erasure Coding

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

LDR > Replication
Attribute Name	Code	Description
Reset Inbound Replication Failure Count	RICR	Reset the counter for inbound replication failures. This can be used to clear the RIRF (Inbound Replication — Failed) alarm.
Reset Outbound Replication Failure Count	ROCR	Reset the counter for outbound replication failures. This can be used to clear the RORF (Outbound Replications — Failed) alarm.
Disable Inbound Replication	DSIR	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 can't 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 can't be retrieved from the Storage Node to be copied to other locations in the StorageGRID system. The LDR service is write-only.

Manage full Storage Nodes

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

Add storage volumes

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

Add storage expansion shelves

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

Add Storage Nodes

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

Manage Admin Nodes

Use multiple Admin Nodes

A StorageGRID system can include multiple Admin Nodes to enable you to continuously monitor and configure your StorageGRID system even if one Admin Node fails.

If an Admin Node becomes unavailable, attribute processing continues, alerts and alarms (legacy system) are still triggered, and email notifications and AutoSupport packages are still sent. However, having multiple Admin Nodes does not provide failover protection except for notifications and AutoSupport packages. In particular, alarm acknowledgments made from one Admin Node aren't copied to other Admin Nodes.



There are two options for continuing to view and configure the StorageGRID system if an Admin Node fails:

- Web clients can reconnect to any other available Admin Node.
- If a system administrator has configured a high availability group of Admin Nodes, web clients can continue

to access the Grid Manager or the Tenant Manager using the virtual IP address of the HA group. See Manage high availability groups.



When using an HA group, access is interrupted if the active Admin Node fails. Users must sign in again after the virtual IP address of the HA group fails over to another Admin Node in the group.

Some maintenance tasks can only be performed using the primary Admin Node. If the primary Admin Node fails, it must be recovered before the StorageGRID system is fully functional again.

Identify the primary Admin Node

The primary Admin Node hosts the CMN service. Some maintenance procedures can only be performed using the primary Admin Node.

Before you begin

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

Steps

- 1. Select SUPPORT > Tools > Grid topology.
- Select site > Admin Node, and then select + to expand the topology tree and show the services hosted on this Admin Node.

The primary Admin Node hosts the CMN service.

3. If this Admin Node does not host the CMN service, check the other Admin Nodes.

View notification status and queues

The Network Management System (NMS) service on Admin Nodes sends notifications to the mail server. You can view the current status of the NMS service and the size of its notifications queue on the Interface Engine page.

To access the Interface Engine page, select **SUPPORT > Tools > Grid topology**. Finally, select **site > Admin Node > NMS > Interface Engine**.

Overview Alarms	Reports	Configuration	
Main			
Overview	: NMS (170-	176) - Interface Engine	
	345 10.12.17 POT		
NMS Interface Engine Statu	8. ⁰	Connected	
Connected Services:	Connected Services 15		<u>r</u>
E-mail Notification Ev	ents		
E-mail Notifications Status.		No Errors	E 9
E-mail Notifications Queued		0	2 8
Database Connection	Pool		
Maximum Supported Capac	tv:	100	14

95 %

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. If a notification is dropped from the queue without being sent, the MINS (E-mail Notification Status) minor alarm is triggered.

How Admin Nodes show acknowledged alarms (legacy system)

When you acknowledge an alarm on one Admin Node, the acknowledged alarm is not copied to any other Admin Node. Because acknowledgments aren't copied to other Admin Nodes, the Grid Topology tree might not look the same for each Admin Node.

This difference can be useful when connecting web clients. Web clients can have different views of the StorageGRID system based on the administrator needs.



Remaining Capacity:

Active Connections:

Note that notifications are sent from the Admin Node where the acknowledgment occurs.

Configure audit client access

Configure audit client access for NFS

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. The audit share is automatically enabled as a read-only share.



Support for NFS has been deprecated and will be removed in a future release.

To access audit logs, you can configure client access to audit shares for NFS. Or, you can use an external syslog server.

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 Review audit logs.

Before you begin

- You have the Passwords.txt file with the root/admin password.
- You have the Configuration.txt file (available in the Recovery Package).
- The audit client is using NFS Version 3 (NFSv3).

About this task

Perform this procedure for each Admin Node in a StorageGRID deployment from which you want to retrieve audit messages.

Steps

- 1. Log in to the primary Admin Node:
 - a. Enter the following command: ssh admin@primary_Admin_Node_IP
 - b. Enter the password listed in the Passwords.txt file.
 - c. Enter the following command to switch to root: su -
 - d. Enter the password listed in the <code>Passwords.txt</code> file.

When you are logged in as root, the prompt changes from \$ to #.

2. Confirm that all services have a state of Running or Verified. Enter: storagegrid-status

If any services aren't 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 	<pre> validate-config refresh-config help exit</pre>

- 5. Add the audit client: add-audit-share
 - a. When prompted, enter the audit client's IP address or IP address range for the audit share: client IP address
 - b. When prompted, press Enter.
- 6. If more than one audit client is permitted to access the audit share, add the IP address of the additional user: add-ip-to-share
 - a. Enter the number of the audit share: audit share number
 - b. When prompted, enter the audit client's IP address or IP address range for the audit share: *client IP address*
 - c. When prompted, press Enter.

The NFS configuration utility is displayed.

- d. Repeat these substeps for each additional audit client that has access to the audit share.
- 7. Optionally, verify your configuration.
 - a. Enter the following: validate-config

The services are checked and displayed.

b. When prompted, press Enter.

The NFS configuration utility is displayed.

- c. Close the NFS configuration utility: exit
- 8. Determine if you must enable audit shares at other sites.
 - If the StorageGRID deployment is a single site, go to the next step.
 - If the StorageGRID deployment includes Admin Nodes at other sites, enable these audit shares as required:
 - a. Remotely log in to the site's Admin Node:
 - i. Enter the following command: ssh admin@grid node IP
 - ii. Enter the password listed in the Passwords.txt file.
 - iii. Enter the following command to switch to root: su -
 - iv. Enter the password listed in the Passwords.txt file.

- b. Repeat these steps to configure the audit shares for each additional Admin Node.
- c. Close the remote secure shell login to the remote Admin Node. Enter: \mathtt{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.

Add 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.



Support for NFS has been deprecated and will be removed in a future release.

Before you begin

- You have the Passwords.txt file with the root/admin account password.
- You have the Configuration.txt file (available in the Recovery Package).
- The audit client is using NFS Version 3 (NFSv3).

Steps

- 1. Log in to the primary Admin Node:
 - a. Enter the following command: ssh admin@primary_Admin_Node_IP
 - b. Enter the password listed in the Passwords.txt file.
 - c. Enter the following command to switch to root: su -
 - d. Enter the password listed in the <code>Passwords.txt</code> file.

When you are logged in as root, the prompt changes from \$ to #.

2. Start the NFS configuration utility: config_nfs.rb

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/log

- 4. Enter the number of the audit share: audit share number
- 5. When prompted, enter the audit client's IP address or IP address range for the audit share: *client_IP_address*

The audit client is added to the audit share.

6. When prompted, press Enter.

The NFS configuration utility is displayed.

- 7. Repeat the steps for each audit client that should be added to the audit share.
- 8. Optionally, verify your configuration: validate-config

The services are checked and displayed.

a. When prompted, press Enter.

The NFS configuration utility is displayed.

- 9. Close the NFS configuration utility: exit
- 10. If the StorageGRID deployment is a single site, go to the next step.

Otherwise, if the StorageGRID deployment includes Admin Nodes at other sites, optionally enable these audit shares as required:

- a. Remotely log in to a site's Admin Node:
 - i. Enter the following command: ssh admin@grid node IP
 - ii. Enter the password listed in the Passwords.txt file.
 - iii. Enter the following command to switch to root: su -
 - iV. Enter the password listed in the Passwords.txt file.
- b. Repeat these steps to configure the audit shares for each Admin Node.
- c. Close the remote secure shell login to the remote Admin Node: exit
- 11. Log out of the command shell: exit

Verify 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.



Support for NFS has been deprecated and will be removed in a future release.

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. An example Linux command is (enter on one line):

mount -t nfs -o hard, intr Admin Node IP address:/var/local/log 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.

Remove 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 have the Passwords.txt file with the root/admin account password.
- You have the Configuration.txt file (available in the Recovery Package).

About this task

You can't remove the last IP address permitted to access the audit share.

Steps

- 1. Log in to the primary Admin Node:
 - a. Enter the following command: ssh admin@primary_Admin_Node_IP
 - b. Enter the password listed in the Passwords.txt file.
 - c. Enter the following command to switch to root: su -
 - d. Enter the password listed in the <code>Passwords.txt</code> file.

When you are logged in as root, the prompt changes from \$ to #.

2. Start the NFS configuration utility: config nfs.rb

Shares	Clients	Config
add-audit-share	add-ip-to-share	validate-config
enable-disable-share	remove-ip-from-share	refresh-config
		help
		exit

3. Remove the IP address from the audit share: remove-ip-from-share

A numbered list of audit shares configured on the server is displayed. The audit share is listed as: /var/local/log

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 ${\tt Passwords.txt}$ file.
 - iii. Enter the following command to switch to root: ${\tt su}\,$ –
 - iv. Enter the password listed in the Passwords.txt file.
 - b. Repeat these steps to configure the audit shares for each additional Admin Node.
 - c. Close the remote secure shell login to the remote Admin Node: exit
- 9. Log out of the command shell: exit

Change the IP address of an NFS audit client

Complete these steps if you need to change 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 information

- Add an NFS audit client to an audit share
- Remove an NFS audit client from the audit share

Manage Archive Nodes

Archive to the cloud through the S3 API

You can configure an Archive Node to connect directly to Amazon Web Services (AWS) or to any other system that can interface to the StorageGRID system through the S3 API.

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



The Cloud Tiering - Simple Storage Service (S3) option is also deprecated. If you are currently using an Archive Node with this option, migrate your objects to a Cloud Storage Pool instead.

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

Configure connection settings for the S3 API

If you are connecting to an Archive Node using the S3 interface, you must configure the connection settings for the S3 API. Until these settings are configured, the ARC service remains in a Major alarm state as it is unable to communicate with the external archival storage system.

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



The Cloud Tiering - Simple Storage Service (S3) option is also deprecated. If you are currently using an Archive Node with this option, migrate your objects to a Cloud Storage Pool instead.

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

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.
- You have created a bucket on the target archival storage system:
 - The bucket is dedicated to a single Archive Node. It can't be used by other Archive Nodes or other applications.
 - The bucket has the appropriate region selected for your location.
 - The bucket should be configured with versioning suspended.
- Object Segmentation is enabled and the Maximum Segment Size is less than or equal to 4.5 GiB (4,831,838,208 bytes). S3 API requests that exceed this value will fail if S3 is used as the external archival storage system.

- 1. Select SUPPORT > Tools > Grid topology.
- 2. Select Archive Node > ARC > Target.
- 3. Select **Configuration > Main**.



Target Type:

Cloud Tiering - Simple Storage Service (S3)

Cloud Tiering (S3) Account

Bucket Name:	name		
Region:	Virginia or Pacific Northwest (us-east-1)		-
Endpoint:	https://10.10.10.123:8082	Use AWS	
Endpoint Authentication: Access Key:	ABCD123EFG45AB		
Secret Access Key:	•••••		
Storage Class:	Standard (Default)		•

Apply Changes

4. Select Cloud Tiering - Simple Storage Service (S3) from the Target Type drop-down list.



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. If the network to the external archival storage system is trusted, you can clear the checkbox to disable endpoint SSL certificate and hostname verification for the targeted external archival storage system. If another instance of a StorageGRID system is the target archival storage device and the system is configured with publicly signed certificates, you can keep the checkbox selected.

- Storage Class: Select Standard (Default) for regular storage. Select Reduced Redundancy only for objects that can be easily recreated. Reduced Redundancy provides lower cost storage with less reliability. If the targeted archival storage system is another instance of the StorageGRID system, Storage Class controls how many interim copies of the object are made at ingest on the target system, if dual commit is used when objects are ingested there.
- 6. Select Apply Changes.

The specified configuration settings are validated and applied to your StorageGRID system. After the settings are applied, the target can't be changed.

Modify 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 are signed in to the Grid Manager using a supported web 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.

- 1. Select SUPPORT > Tools > Grid topology.
- 2. Select Archive Node > ARC > Target.
- 3. Select Configuration > Main.



Target Type:

Cloud Tiering - Simple Storage Service (S3)

Cloud Tiering (S3) Account

Bucket Name:	name		
Region:	Virginia or Pacific Northwest (us-east-1)		_
Endpoint:	https://10.10.10.123:8082	Use AWS	
Endpoint Authentication: Access Key:	ABCD123EFG45AB		
Secret Access Key:	•••••		
Storage Class:	Standard (Default)		•

Apply Changes

4. Modify account information, as necessary.

If you change the storage class, new object data is stored with the new storage class. Existing object continue to be stored under the storage class set when ingested.



Bucket Name, Region, and Endpoint, use AWS values and can't be changed.

5. Select Apply Changes.

Modify 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 web 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**.

- 1. Select **SUPPORT > Tools > Grid topology**.
- 2. Select Archive Node > ARC.
- 3. Select **Configuration > Main**.

Overview	Alarms	Reports	Configuration			
Main	Alarms					
	Configuration Updated: 2015-09-24 1	1: ARC (98- 17:18:29 PDT	127) - ARC			
ARC State			Online			•
Cloud Tiering	Service State	F	Read-Write Enable	ed		•
					Apply Changes	

- 4. Select a Cloud Tiering Service State.
- 5. Select Apply Changes.

Reset the Store Failure Count for S3 API connection

If your Archive Node connects to an archival storage system through the S3 API, you can reset the Store Failure Count, which can be used to clear the ARVF (Store Failures) alarm.

Before you begin

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

- 1. Select **SUPPORT > Tools > Grid topology**.
- 2. Select Archive Node > ARC > Store.
- 3. Select Configuration > Main.

Overview	Alarms	Reports	Configuration			
Main	Alarms					
00	Configuration	n: ARC (98- 17:54:42 PDT	127) - Store			
Reset Store F	Failure Count	Γ				
					Apply Changes	2

- 4. Select Reset Store Failure Count.
- 5. Select Apply Changes.

The Store Failures attribute resets to zero.

Migrate 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, you should migrate 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 are signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.
- You have already stored objects in the S3 bucket configured for Cloud Tiering.



Before migrating object data, contact your NetApp account representative to understand and manage any associated costs.

About this task

From an ILM perspective, a Cloud Storage Pool is similar to a storage pool. However, while storage pools consist of Storage Nodes or Archive Nodes within the StorageGRID system, a Cloud Storage Pool consists of an external S3 bucket.

Before migrating objects from Cloud Tiering - S3 to a Cloud Storage Pool, you must first create an S3 bucket and then create the Cloud Storage Pool in StorageGRID. Then, you can create a new ILM policy and replace the ILM rule used to store objects in the Cloud Tiering bucket with a cloned ILM rule that stores the same objects in the Cloud Storage Pool.



When objects are stored in a Cloud Storage Pool, copies of those objects can't 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 policies that cause objects to be stored in the Cloud Tiering bucket.
- 3. Clone each of these rules.
- 4. In the cloned rules, change the placement location to the new Cloud Storage Pool.
- 5. Save the cloned rules.
- 6. Create a new policy that uses the new rules.
- 7. Simulate and activate the new policy.

When the new policy is activated and ILM evaluation occurs, the objects are moved from the S3 bucket configured for Cloud Tiering to the S3 bucket configured for the Cloud Storage Pool. The usable space on

the grid is not affected. After the objects are moved to the Cloud Storage Pool, they are removed from the Cloud Tiering bucket.

Related information

Manage objects with ILM

Archive to tape through TSM middleware

You can configure an Archive Node to target a Tivoli Storage Manager (TSM) server that provides a logical interface for storing and retrieving object data to random or sequential access storage devices, including tape libraries.

The Archive Node's ARC service acts as a client to the TSM server, using Tivoli Storage Manager as middleware for communicating with the archival storage system.

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

The Cloud Tiering - Simple Storage Service (S3) option is also deprecated. If you are currently using an Archive Node with this option, migrate your objects to a Cloud Storage Pool instead.

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

TSM management classes

Management classes defined by the TSM middleware outline how the TSM's backup and archive operations function, and can be used to specify rules for content that are applied by the TSM server. Such rules operate independently of the StorageGRID system's ILM policy, and must be consistent with the StorageGRID system's requirement that objects are stored permanently and are always available for retrieval by the Archive Node. After object data is sent to a TSM server by the Archive Node, the TSM lifecycle and retention rules are applied while the object data is stored to tape managed by the TSM server.

The TSM management class is used by the TSM server to apply rules for data location or retention after objects are sent to the TSM server by the Archive Node. For example, objects identified as database backups (temporary content that can be overwritten with newer data) could be treated differently than application data (fixed content that must be retained indefinitely).

Configure connections to TSM middleware

Before the Archive Node can communicate with Tivoli Storage Manager (TSM) middleware, you must configure several settings.

Before you begin

- You are signed in to the Grid Manager using a supported web 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.

Steps

- 1. Select SUPPORT > Tools > Grid topology.
- 2. Select Archive Node > ARC > Target.
- 3. Select Configuration > Main.

Overview Alarms Reports	Configuration
Main Alarms	
Configuration: ARC (Updated: 2015-09-28 09:56:36 PDT	DC1-ARC1-98-165) - Target
Target Type:	Tivoli Storage Manager (TSM)
Tivoli Storage Manager State:	Online
Target (TSM) Account	
Server IP or Hostname:	10.10.123
Server Port:	1500
Node Name:	ARC-USER
User Name:	arc-user
Password:	•••••
Management Class:	sg-mgmtclass
Number of Sessions:	2
Maximum Retrieve Sessions:	1
Maximum Store Sessions:	1



- 4. From the Target Type drop-down list, select Tivoli Storage Manager (TSM).
- 5. For the **Tivoli Storage Manager State**, select **Offline** to prevent retrievals from the TSM middleware server.

By default, the Tivoli Storage Manager State is set to Online, which means that the Archive Node is able to retrieve object data from the TSM middleware server.

- 6. Complete the following information:
 - **Server IP or Hostname**: Specify the IP address or fully qualified domain name of the TSM middleware server used by the ARC service. The default IP address is 127.0.0.1.
 - **Server Port**: Specify the port number on the TSM middleware server that the ARC service will connect to. The default is 1500.
 - **Node Name**: Specify the name of the Archive Node. You must enter the name (arc-user) that you registered on the TSM middleware server.
 - **User Name**: Specify the user name the ARC service uses to log in to the TSM server. Enter the default user name (arc-user) or the administrative user you specified for the Archive Node.
 - Password: Specify the password used by the ARC service to log in to the TSM server.
 - Management Class: Specify the default management class to use if a management class is not

specified when the object is being saved to the StorageGRID system, or the specified management class is not defined on the TSM middleware server.

 Number of Sessions: Specify the number of tape drives on the TSM middleware server that are dedicated to the Archive Node. The Archive Node concurrently creates a maximum of one session per mount point plus a small number of additional sessions (less than five).

You must change this value to be the same as the value set for MAXNUMMP (maximum number of mount points) when the Archive Node was registered or updated. (In the register command, the default value of MAXNUMMP used is 1, if no value is set.)

You must also change the value of MAXSESSIONS for the TSM server to a number that is at least as large as the Number of Sessions set for the ARC service. The default value of MAXSESSIONS on the TSM server is 25.

- Maximum Retrieve Sessions: Specify the maximum number of sessions that the ARC service can open to the TSM middleware server for retrieve operations. In most cases, the appropriate value is Number of Sessions minus Maximum Store Sessions. If you need to share one tape drive for storage and retrieval, specify a value equal to the Number of Sessions.
- **Maximum Store Sessions**: Specify the maximum number of concurrent sessions that the ARC service can open to the TSM middleware server for archive operations.

This value should be set to one except when the targeted archival storage system is full and only retrievals can be performed. Set this value to zero to use all sessions for retrievals.

7. Select Apply Changes.

Optimize an Archive Node for TSM middleware sessions

You can optimize the performance of an Archive Node that connects to Tivoli Server Manager (TSM) by configuring the Archive Node's sessions.

Before you begin

- You are signed in to the Grid Manager using a supported web 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 > Tools > Grid topology.
- 2. Select Archive Node > ARC > Target.
- 3. Select Configuration > Main.
- 4. Change Maximum Retrieve Sessions to be the same as Number of Sessions.

Overview Alarms Reports	Configuration	
Main Alarms		
Configuration: ARC (I Updated: 2015-09-28 09:56:36 PDT	DC1-ARC1-98-165) - Target	
Target Type:	Tivoli Storage Manager (TSM)	
Tivoli Storage Manager State:	Online	
Target (TSM) Account		
Server IP or Hostname:	10.10.123	
Server Port:	1500	
Node Name:	ARC-USER	
User Name:	arc-user	
Password:	•••••	
Management Class:	sg-mgmtclass	

Maximum Retrieve Sessions: Maximum Store Sessions:

Number of Sessions:

Apply Changes

5. Select Apply Changes.

Configure the archive state and counters for TSM

If your Archive Node connects to a TSM middleware server, you can configure an Archive Node's archive store state to Online or Offline. You can also disable the archive store when the Archive Node first starts up, or reset the failure count being tracked for the associated alarm.

Before you begin

• You are signed in to the Grid Manager using a supported web browser.

2

2

1

• You have specific access permissions.

- 1. Select SUPPORT > Tools > Grid topology.
- 2. Select Archive Node > ARC > Store.
- 3. Select **Configuration > Main**.

Overview Alarms Repo	orts Configuration	
Main Alarms		
Configuration: ARC Updated: 2015-09-29 17:10:12 F	С (DC1-ARC1-98-165) - Store	
Store State	Online	•
Archive Store Disabled on Startup		
Reset Store Failure Count		



- 4. Modify the following settings, as necessary:
 - Store State: Set the component state to either:
 - Online: The Archive Node is available to process object data for storage to the archival storage system.
 - Offline: The Archive Node is not available to process object data for storage to the archival storage system.
 - Archive Store Disabled on Startup: When selected, the Archive Store component remains in the Readonly state when restarted. Used to persistently disable storage to the targeted the archival storage system. Useful when the targeted 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. Select Apply Changes.

Related information

Manage an Archive Node when TSM server reaches capacity

Manage 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. This situation can be avoided through proactive monitoring of the TSM server.

Before you begin

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

About this task

The Archive Node continues to accept object data for transfer to the TSM server after the TSM server stops accepting new content. This content can't be written to media managed by the TSM server. An alarm is triggered if this happens.

Prevent ARC service from sending content to TSM server

To prevent the ARC service from sending further content to the TSM server, you can take the Archive Node

offline by taking its **ARC** > **Store** component offline. This procedure can also be useful in preventing alarms when the TSM server is unavailable for maintenance.

Steps

- 1. Select **SUPPORT > Tools > Grid topology**.
- 2. Select Archive Node > ARC > Store.
- 3. Select Configuration > Main.

Overview	Alarms	Reports	Configuration	
Main	Alarms			
00 0	Configuration	n: ARC (DC	1-ARC1-98-165) - Store	
Store State		[Offline	
Archive Store D	isabled on Start	up	3	lines
Reset Store Fai	lure Count	l		
				Apply Changes

- 4. Change Store State to Offline.
- 5. Select Archive Store Disabled on Startup.
- 6. Select Apply Changes.

Set 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.

Steps

- 1. Select SUPPORT > Tools > Grid topology.
- 2. Select Archive Node > ARC > Target.
- 3. Select Configuration > Main.
- 4. Change Maximum Retrieve Sessions to be the same as the number of concurrent sessions listed in Number of Sessions.
- 5. Change Maximum Store Sessions to 0.



Changing Maximum Store Sessions to 0 is not necessary if the Archive Node is Read-only. Store sessions will not be created.

6. Select Apply Changes.

Configure Archive Node retrieve settings

You can configure the retrieve settings for an Archive Node to set the state to Online or Offline, or reset the failure counts being tracked for the associated alarms.

Before you begin

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

Steps

- 1. Select SUPPORT > Tools > Grid topology.
- 2. Select Archive Node > ARC > Retrieve.
- 3. Select Configuration > Main.

Overview	Alarms	Reports	Configuration	
Main	Alarms			
	onfiguratior	: ARC (DC	;1-ARC1-98-165) - Retrieve	
Up	dated: 2015-05-07 12	:24:45 PDT		
Retrieve State		F	Online	Y
Reset Request Failure Count		[
Reset Verification Failure Count		(
				-
			Apply Cha	nges 📖

- 4. Modify the following settings, as necessary:
 - Retrieve State: Set the component state to either:
 - Online: The grid node is available to retrieve object data from the archival media device.
 - Offline: The grid node is not available to retrieve object data.
 - Reset Request Failures Count: Select the 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. Select Apply Changes.

Configure 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 are signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

- 1. Select SUPPORT > Tools > Grid topology.
- 2. Select Archive Node > ARC > Replication.
- 3. Select **Configuration > Main**.

Overview Alarms Reports	Configuration
Main Alarms	
Configuration: ARC (E Updated: 2015-05-07 12:21:53 PDT	OC1-ARC1-98-165) - Replication
Reset Inbound Replication Failure Count	
Reset Outbound Replication Failure Count	
Inbound Replication	
Disable Inbound Replication	
Outbound Replication	
Disable Outbound Replication	
	Apply Changes

- 4. Modify the following settings, as necessary:
 - **Reset Inbound Replication Failure Count**: Select to reset the counter for inbound replication failures. This can be used to clear the RIRF (Inbound Replications Failed) alarm.
 - **Reset Outbound Replication Failure Count**: Select to reset the counter for outbound replication failures. This can be used to clear the RORF (Outbound Replications Failed) alarm.
 - **Disable Inbound Replication**: Select to disable inbound replication as part of a maintenance or testing procedure. Leave cleared during normal operation.

When inbound replication is disabled, object data can be retrieved from the ARC service for replication to other locations in the StorageGRID system, but objects can't 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 can't be retrieved from the ARC service to be copied to other locations in the StorageGRID system. The ARC service is write-only.

5. Select Apply Changes.

Set 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. See the instructions for managing alarms (legacy system).

Integrate Tivoli Storage Manager

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 don't 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 can't 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 can't 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.

Complete 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.

Refer to the following IBM documentation, as necessary, as you prepare your TSM server for integration with the Archive Node in a StorageGRID system:

- IBM Tape Device Drivers Installation and User's Guide
- IBM Tape Device Drivers Programming Reference

Install a new TSM server

You can integrate the Archive Node with either a new or an existing TSM server. If you are installing a new TSM server, follow the instructions in your TSM documentation to complete the installation.



An Archive Node can't be co-hosted with a TSM server.

Configure the TSM server

This section includes example instructions for preparing a TSM server following TSM best practices.

The following instructions guide you through the process of:

• Defining a disk storage pool, and a tape storage pool (if required) on the TSM server

• Defining a domain policy that uses the TSM management class for the data saved from the Archive Node, and registering a node to use this domain policy

These instructions are provided for your guidance only; they aren't intended to replace TSM documentation, or to provide complete and comprehensive instructions suitable for all configurations. Deployment specific instructions should be provided by a TSM administrator who is familiar both with your detailed requirements, and with the complete set of TSM Server documentation.

Define 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 several steps on your TSM server before you can create a tape storage pool. (Create a tape library and at least one drive in the tape library. Define a path from the server to the library and from the server to the drives, and then define a device class for the drives.) The details of these steps can vary depending upon the hardware configuration and storage requirements of the site. For more information, see the TSM documentation.

The following set of instructions illustrates the process. You should be aware that the requirements for your site could be different depending on the requirements of your deployment. For configuration details and for instructions, see the TSM documentation.



You must log in to 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.

define path *servername tapelibrary* srctype=server desttype=library device=*lib-devicename*

- ° servername is the name of the TSM server
- ° tapelibrary is the tape library name you defined
- ° *lib-devicename* is the device name for the tape library
- 3. Define a drive for the library.

define drive tapelibrary drivename

- ° drivename is the name you want to specify for the drive
- ° tapelibrary is the tape library name you defined

You might want to configure an additional drive or drives, depending upon your hardware configuration. (For example, if the TSM server is connected to a Fibre Channel switch that has two inputs from a tape library, you might want to define a drive for each input.)

4. Define a path from the server to the drive you defined.

```
define path servername drivename srctype=server desttype=drive
library=tapelibrary device=drive-dname
```

- ° drive-dname is the device name for the drive
- ° tapelibrary is the tape library name you defined

Repeat for each drive that you have defined for the tape library, using a separate *drivename* and *drive-dname* for each drive.

5. Define a device class for the drives.

```
define devclass DeviceClassName devtype=lto library=tapelibrary
format=tapetype
```

- ° DeviceClassName is the name of the device class
- ° 1 to 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.
- $^\circ\,$ 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).

Define a domain policy and register a node

You need to define a domain policy that uses the TSM management class for the data saved from the Archive Node, and then register a node to use this domain policy.



Archive Node processes can leak memory if the client password for the Archive Node in Tivoli Storage Manager (TSM) expires. Ensure that the TSM server is configured so the client username/password for the Archive Node never expires.

When registering a node on the TSM server for the use of the Archive Node (or updating an existing node), you must specify the number of mount points that the node can use for write operations by specifying the MAXNUMMP parameter to the REGISTER NODE command. The number of mount points is typically equivalent to the number of tape drive heads allocated to the Archive Node. The number specified for MAXNUMMP on the TSM server must be at least as large as the value set for the **ARC** > **Target** > **Configuration** > **Main** > **Maximum Store Sessions** for the Archive Node, which is set to a value of 0 or 1, as concurrent store sessions aren't supported by the Archive Node.

The value of MAXSESSIONS set for the TSM server controls the maximum number of sessions that can be opened to the TSM server by all client applications. The value of MAXSESSIONS specified on the TSM must be at least as large as the value specified for **ARC** > **Target** > **Configuration** > **Main** > **Number of Sessions** in the Grid Manager for the Archive Node. The Archive Node concurrently creates at most one session per mount point plus a small number (< 5) of additional sessions.

The TSM node assigned to the Archive Node uses a custom domain policy tsm-domain. The tsm-domain domain policy is a modified version of the "standard" domain policy, configured to write to tape and with the archive destination set to be the StorageGRID system's storage pool (*SGWSDiskPool*).



You must log in to the TSM server with administrative privileges and use the dsmadmc tool to create and activate the domain policy.

Create and activate 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 aren't using an existing management class, enter one of the following:

define policyset tsm-domain standard

define mgmtclass tsm-domain standard default

default is the default management class for the deployment.

3. Create a copygroup to the appropriate storage pool. Enter (on one line):

define copygroup tsm-domain standard default type=archive
destination=SGWSDiskPool retinit=event retmin=0 retver=0

default is the default Management Class for the Archive Node. The values of retinit, retmin, and retver have been chosen to reflect the retention behavior currently used by the Archive Node



Don't set retinit to retinit=create. Setting retinit=create blocks the Archive Node from deleting content, because retention events are used to remove content from the TSM server.

4. Assign the management class to be the default.

assign defmgmtclass tsm-domain standard default

5. Set the new policy set as active.

activate policyset tsm-domain standard

Ignore the "no backup copy group" warning that appears when you enter the activate command.

6. Register a node to use the new policy set on the TSM server. On the TSM server, enter (on one line):

```
register node arc-user arc-password passexp=0 domain=tsm-domain
MAXNUMMP=number-of-sessions
```

arc-user and arc-password are same client node name and password as you define on the Archive Node, and the value of MAXNUMMP is set to the number of tape drives reserved for Archive Node store sessions.



By default, registering a node creates an administrative user ID with client owner authority, with the password defined for the node.

Migrate data into StorageGRID

You can migrate large amounts of data to the StorageGRID system while simultaneously using the StorageGRID system for day-to-day operations.

Use this guide as you plan a migration of large amounts of data into the StorageGRID system. It is not a general guide to data migration, and it does not include detailed steps for performing a migration. Follow the guidelines and instructions in this section to ensure that data is migrated efficiently into the StorageGRID system without interfering with day-to-day operations, and that the migrated data is handled appropriately by the StorageGRID system.

Confirm 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 Manage Storage Nodes and the instructions for monitoring StorageGRID.

Determine 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. See Manage objects with ILM.



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.

Assess 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 guide 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 aren't exceeded.

Schedule and monitor data migration

Data migration must be scheduled and monitored as necessary to ensure data is placed according to the ILM policy within the required timeframe.

Schedule 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, don't 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.

Monitor data migration

This table lists the attributes you must monitor during data migration, and the issues that they represent.

If you use traffic classification policies with rate limits to throttle ingest, you can monitor the observed rate in conjunction with the statistics described in the following table and reduce the limits if necessary.

Monitor	Description	
Number of objects waiting for ILM	1. Select SUPPORT > Tools > Grid topology.	
evaluation	2. Select <i>deployment</i> > Overview > Main.	
	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. 	
	 Awaiting - Client (XCQZ): The total number of objects awaiting ILM evaluation from client operations (for example, ingest). 	
	 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	If an alarm for the Store Failures (ARVF) 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.	

Manage objects with ILM

Manage objects with ILM

The information lifecycle management (ILM) rules in an ILM policy instruct StorageGRID how to create and distribute copies of object data and how to manage those copies over time.

About these instructions

Designing and implementing ILM rules and policies requires careful planning. You must understand your operational requirements, the topology of your StorageGRID system, your object protection needs, and the available storage types. Then, you must determine how you want different types of objects to be copied, distributed, and stored.

Use these instructions to:

- Learn about StorageGRID ILM, including how ILM operates throughout an object's life.
- Learn how to configure storage pools, Cloud Storage Pools, and ILM rules.
- Learn how to create, simulate, and activate an ILM policy that will protect object data across one or more sites.
- Learn how to manage objects with S3 Object Lock, which helps to ensure that objects in specific S3 buckets aren't deleted or overwritten for a specified amount of time.

Learn more

To learn more, review these videos:

• Video: Information lifecycle management rules in StorageGRID 11.8.



• Video: Information lifecycle management policies in StorageGRID 11.8



ILM and object lifecycle

How ILM operates throughout an object's life

Understanding how StorageGRID uses ILM to manage objects during every stage of their life can help you design a more effective policy.

- Ingest: Ingest begins when an S3 or Swift client application establishes a connection to save an object to the StorageGRID system, and is complete when StorageGRID returns an "ingest successful" message to the client. Object data is protected during ingest either by applying ILM instructions immediately (synchronous placement) or by creating interim copies and applying ILM later (dual commit), depending on how the ILM requirements were specified.
- **Copy management**: After creating the number and type of object copies that are specified in the ILM's placement instructions, StorageGRID manages object locations and protects objects against loss.
 - ILM scanning and evaluation: StorageGRID continuously scans the list of objects stored in the grid and checks if the current copies meet ILM requirements. When different types, numbers, or locations of object copies are required, StorageGRID creates, deletes, or moves copies as needed.
 - Background verification: StorageGRID continuously performs background verification to check the integrity of object data. If a problem is found, StorageGRID automatically creates a new object copy or a replacement erasure-coded object fragment in a location that meets current ILM requirements. See Verify object integrity.
- **Object deletion**: Management of an object ends when all copies are removed from the StorageGRID system. Objects can be removed as a result of a delete request by a client, or as a result of deletion by ILM or deletion caused by the expiration of an S3 bucket lifecycle.



Objects in a bucket that has S3 Object Lock enabled can't be deleted if they are under a legal hold or if a retain-until-date has been specified but not yet met.

The diagram summarizes how ILM operates throughout an object's lifecycle.



How objects are ingested

Ingest options

When you create an ILM rule, you specify one of three options for protecting objects at ingest: Dual commit, Strict, or Balanced.

Depending on your choice, StorageGRID makes interim copies and queues the objects for ILM evaluation
later, or it uses synchronous placement and immediately makes copies to meet ILM requirements.

Flowchart of ingest options

The flowchart shows what happens when objects are matched by an ILM rule that uses each of the three ingest options.



Dual commit

When you select the Dual commit option, StorageGRID immediately makes interim object copies on two different Storage Nodes and returns an "ingest successful" message to the client. The object is queued for ILM evaluation, and copies that meet the rule's placement instructions are made later. If the ILM policy can't be processed immediately after the dual commit, site-loss protection could take time to achieve.

Use the Dual commit option in either of these cases:

• You are using multi-site ILM rules and client ingest latency is your primary consideration. When using Dual

commit, you must ensure your grid can perform the additional work of creating and removing the dualcommit copies if they don't satisfy ILM. Specifically:

- $\circ\,$ The load on the grid must be low enough to prevent an ILM backlog.
- The grid must have excess hardware resources (IOPS, CPU, memory, network bandwidth, and so on).
- You are using multi-site ILM rules and the WAN connection between the sites usually has high latency or limited bandwidth. In this scenario, using the Dual commit option can help prevent client timeouts. Before choosing the Dual commit option, you should test the client application with realistic workloads.

Balanced (default)

When you select the Balanced option, StorageGRID also uses synchronous placement on ingest and immediately makes all copies specified in the rule's placement instructions. In contrast with the Strict option, if StorageGRID can't immediately make all copies, it uses Dual commit instead. If the ILM policy uses placements on multiple sites and immediate site-loss protection can't be achieved, the **ILM placement unachievable** alert is triggered.

Use the Balanced option to achieve the best combination of data protection, grid performance, and ingest success. Balanced is the default option in the Create ILM rule wizard.

Strict

When you select the Strict option, StorageGRID uses synchronous placement on ingest and immediately makes all object copies specified in the rule's placement instructions. Ingest fails if StorageGRID can't create all copies, for example, because a required storage location is temporarily unavailable. The client must retry the operation.

Use the Strict option if you have an operational or regulatory requirement to immediately store objects only in the locations outlined in the ILM rule. For example, to satisfy a regulatory requirement, you might need to use the Strict option and a Location Constraint advanced filter to guarantee that objects are never stored at certain data centers.

See Example 5: ILM rules and policy for Strict ingest behavior.

Advantages, disadvantages, and limitations of the ingest options

Understanding the advantages and disadvantages of each of the three options for protecting data at ingest (Balanced, Strict, or Dual commit) can help you decide which one to select for an ILM rule.

For an overview of ingest options, see Ingest options.

Advantages of the Balanced and Strict options

When compared to Dual commit, which creates interim copies during ingest, the two synchronous placement options can provide the following advantages:

- **Better data security**: Object data is immediately protected as specified in the ILM rule's placement instructions, which can be configured to protect against a wide variety of failure conditions, including the failure of more than one storage location. Dual commit can only protect against the loss of a single local copy.
- More efficient grid operation: Each object is processed only once, as it is ingested. Because the StorageGRID system does not need to track or delete interim copies, there is less processing load and less

database space is consumed.

- (Balanced) Recommended: The Balanced option provides optimal ILM efficiency. Using the Balanced option is recommended unless Strict ingest behavior is required or the grid meets all of the criteria for using Dual commit.
- (Strict) Certainty about object locations: The Strict option guarantees that objects are immediately stored according to the placement instructions in the ILM rule.

Disadvantages of the Balanced and Strict options

When compared to Dual commit, the Balanced and Strict options have some disadvantages:

- Longer client ingests: Client ingest latencies might be longer. When you use the Balanced or Strict options, an "ingest successful" message is not returned to the client until all erasure-coded fragments or replicated copies are created and stored. However, object data will most likely reach its final placement much faster.
- (Strict) Higher rates of ingest failure: With the Strict option, ingest fails whenever StorageGRID can't immediately make all copies specified in the ILM rule. You might see high rates of ingest failure if a required storage location is temporarily offline or if network issues cause delays in copying objects between sites.
- (Strict) S3 multipart upload placements might not be as expected in some circumstances: With Strict, you expect objects either to be placed as described by the ILM rule or for ingest to fail. However, with an S3 multipart upload, ILM is evaluated for each part of the object as it is ingested, and for the object as a whole when the multipart upload completes. In the following circumstances this might result in placements that are different than you expect:
 - If ILM changes while an S3 multipart upload is in progress: Because each part is placed according to the rule that is active when the part is ingested, some parts of the object might not meet current ILM requirements when the multipart upload completes. In these cases, ingest of the object does not fail. Instead, any part that is not placed correctly is queued for ILM re-evaluation and is moved to the correct location later.
 - When ILM rules filter on size: When evaluating ILM for a part, StorageGRID filters on the size of the part, not the size of the object. This means that parts of an object can be stored in locations that don't meet ILM requirements for the object as a whole. For example, if a rule specifies that all objects 10 GB or larger are stored at DC1 while all smaller objects are stored at DC2, at ingest each 1 GB part of a 10-part multipart upload is stored at DC2. When ILM is evaluated for the object, all parts of the object are moved to DC1.
- (Strict) Ingest does not fail when object tags or metadata are updated and newly required placements cannot be made: With Strict, you expect objects either to be placed as described by the ILM rule or for ingest to fail. However, when you update metadata or tags for an object that is already stored in the grid, the object is not re-ingested. This means that any changes to object placement that are triggered by the update aren't made immediately. Placement changes are made when ILM is re-evaluated by normal background ILM processes. If required placement changes can't be made (for example, because a newly required location is unavailable), the updated object retains its current placement until the placement changes are possible.

Limitations on object placements with the Balanced and Strict options

The Balanced or Strict options can't be used for ILM rules that have any of these placement instructions:

- Placement in a Cloud Storage Pool at day 0.
- Placement in an Archive Node at day 0.

• Placements in a Cloud Storage Pool or an Archive Node when the rule has a User defined creation time as its Reference time.

These restrictions exist because StorageGRID can't synchronously make copies to a Cloud Storage Pool or an Archive Node, and a User defined creation time could resolve to the present.

How ILM rules and consistency interact to affect data protection

Both your ILM rule and your choice of consistency affect how objects are protected. These settings can interact.

For example, the ingest behavior selected for an ILM rule affects the initial placement of object copies, while the consistency used when an object is stored affects the initial placement of object metadata. Because StorageGRID requires access to both an object's data and metadata to fulfill client requests, selecting matching levels of protection for the consistency and ingest behavior can provide better initial data protection and more predictable system responses.

Here is a brief summary of the consistency values that are available in StorageGRID:

- All: All nodes receive object metadata immediately or the request will fail.
- **Strong-global**: Object metadata is immediately distributed to all sites. Guarantees read-after-write consistency for all client requests across all sites.
- **Strong-site**: Object metadata is immediately distributed to other nodes at the site. Guarantees read-afterwrite consistency for all client requests within a site.
- **Read-after-new-write**: Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees. Recommended for most cases.
- Available: Provides eventual consistency for both new objects and object updates. For S3 buckets, use only as required (for example, for a bucket that contains log values that are rarely read, or for HEAD or GET operations on keys that don't exist). Not supported for S3 FabricPool buckets.



Before selecting a consistency value, read the full description of consistency. You should understand the benefits and limitations before changing the default value.

Example of how consistency and ILM rules can interact

Suppose you have a two-site grid with the following ILM rule and the following consistency:

- **ILM rule**: Create two object copies, one at the local site and one at a remote site. Use Strict ingest behavior.
- consistency: Strong-global (object metadata is immediately distributed to all sites).

When a client stores an object to the grid, StorageGRID makes both object copies and distributes metadata to both sites before returning success to the client.

The object is fully protected against loss at the time of the ingest successful message. For example, if the local site is lost shortly after ingest, copies of both the object data and the object metadata still exist at the remote site. The object is fully retrievable.

If you instead used the same ILM rule and the strong-site consistency, the client might receive a success message after object data is replicated to the remote site but before object metadata is distributed there. In this case, the level of protection of object metadata does not match the level of protection for object data. If the local site is lost shortly after ingest, object metadata is lost. The object can't be retrieved.

The inter-relationship between consistency and ILM rules can be complex. Contact NetApp if you need assistance.

Related information

• Example 5: ILM rules and policy for Strict ingest behavior

How objects are stored (replication or erasure coding)

What is replication?

Replication is one of two methods used by StorageGRID to store object data. When objects match an ILM rule that uses replication, the system creates exact copies of object data and stores the copies on Storage Nodes or Archive Nodes.

When you configure an ILM rule to create replicated copies, you specify how many copies should be created, where those copies should be placed, and how long the copies should be stored at each location.

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



When StorageGRID matches objects to this rule, it creates two copies of the object, placing each copy on a different Storage Node in the storage pool. The two copies might be placed on any two of the three available Storage Nodes. In this case, the rule placed object copies on Storage Nodes 2 and 3. Because there are two copies, the object can be retrieved if any of the nodes in the storage pool fails.



StorageGRID can store only one replicated copy of an object on any given Storage Node. If your grid includes three Storage Nodes and you create a 4-copy ILM rule, only three copies will be made—one copy for each Storage Node. The **ILM placement unachievable** alert is triggered to indicate that the ILM rule could not be completely applied.

Related information

- · What is erasure coding
- What is a storage pool
- Enable site-loss protection using replication and erasure coding

Why you should not use single-copy replication

When creating an ILM rule to create replicated copies, you should always specify at least two copies for any time period in the placement instructions.



Don't use an ILM rule that creates only one replicated copy for any time period. If only one replicated copy of an object exists, that object is lost if a Storage Node fails or has a significant error. You also temporarily lose access to the object during maintenance procedures such as upgrades.

In the following example, the Make 1 Copy ILM rule specifies that one replicated copy of an object be placed in a storage pool that contains three Storage Nodes. When an object is ingested that matches this rule, StorageGRID places a single copy on only one Storage Node.



When an ILM rule creates only one replicated copy of an object, the object becomes inaccessible when the Storage Node is unavailable. In this example, you will temporarily lose access to object AAA whenever Storage Node 2 is offline, such as during an upgrade or other maintenance procedure. You will lose object AAA entirely if Storage Node 2 fails.



To avoid losing object data, you should always make at least two copies of all objects you want to protect with replication. If two or more copies exist, you can still access the object if one Storage Node fails or goes offline.



What is erasure coding?

Erasure coding is one of two methods StorageGRID uses to store object data. When objects match an ILM rule that uses erasure coding, those objects are sliced into data fragments, additional parity fragments are computed, and each fragment is stored on a different Storage Node.

When an object is accessed, it is reassembled using the stored fragments. If a data or a parity fragment becomes corrupt or lost, the erasure-coding algorithm can recreate that fragment using a subset of the remaining data and parity fragments.

As you create ILM rules, StorageGRID creates erasure-coding profiles that support those rules. You can view a list of erasure-coding profiles, rename an erasure-coding profile, or deactivate an erasure-coding profile if it is not currently used in any ILM rules.

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







The 4+2 erasure-coding scheme can be configured in various ways. For example, you can configure a singlesite storage pool that contains six Storage Nodes. For site-loss protection, you can use a storage pool containing three sites with three Storage Nodes at each site. An object can be retrieved as long as any four of the six fragments (data or parity) remain available. Up to two fragments can be lost without loss of the object data. If an entire site is lost, the object can still be retrieved or repaired, as long as all of the other fragments remain accessible.







If more than two Storage Nodes are lost, the object is not retrievable.







Related information

- What is replication
- What is a storage pool
- What are erasure-coding schemes

- Rename an erasure-coding profile
- Deactivate an erasure-coding profile

What are erasure-coding schemes?

Erasure-coding schemes control how many data fragments and how many parity fragments are created for each object.

When you 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.

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. To retrieve or repair an object, k fragments are needed.

When selecting the storage pool to use for a rule that will create an erasure-coded copy, use the following guidelines for storage pools:

• The storage pool must include three or more sites, or exactly one site.



You can't use erasure coding if the storage pool includes two sites.

- Erasure-coding schemes for storage pools containing three or more sites
- Erasure-coding schemes for one-site storage pools
- Don't use a storage pool that includes the default site, All Sites.
- The storage pool should include at least k+m+1 Storage Nodes that can store object data.



Storage Nodes can be configured during installation to contain only object metadata and not object data. For more information, see Types of Storage Nodes.

The minimum number of Storage Nodes required is k+m. However, having at least one additional Storage Node can help prevent ingest failures or ILM backlogs if a required Storage Node is temporarily unavailable.

The storage overhead of an erasure-coding scheme is calculated by dividing the number of parity fragments (m) by the number of data fragments (k). You can use the storage overhead to calculate how much disk space each erasure-coded object requires:

disk space = object size + (object size * storage overhead)

For example, if you store a 10 MB object using the 4+2 scheme (which has 50% storage overhead), the object consumes 15 MB of grid storage. If you store the same 10 MB object using the 6+2 scheme (which has 33% storage overhead), the object consumes approximately 13.3 MB.

Select the erasure-coding scheme with the lowest total value of k+m that meets your needs. Erasure-coding schemes with a lower number of fragments are overall more computationally efficient, as fewer fragments are created and distributed (or retrieved) per object, can show better performance due to the larger fragment size, and can require fewer nodes be added in an expansion when more storage is required. (For information about planning a storage expansion, see the instructions for expanding StorageGRID.)

Erasure-coding schemes for storage pools containing three or more sites

The following table describes the erasure-coding schemes currently supported by StorageGRID for storage pools that include three or more sites. All of these schemes provide site-loss protection. One site can be lost, and the object will still be accessible.

For erasure-coding schemes that provide site-loss protection, the recommended number of Storage Nodes in the storage pool exceeds k+m+1 because each site requires a minimum of three Storage Nodes.

Erasure-coding scheme (<i>k</i> + <i>m</i>)	Minimum number of deployed sites	Recommended number of Storage Nodes at each site	Total recommended number of Storage Nodes	Site loss protection?	Storage overhead
4+2	3	3	9	Yes	50%
6+2	4	3	12	Yes	33%
8+2	5	3	15	Yes	25%
6+3	3	4	12	Yes	50%
9+3	4	4	16	Yes	33%
2+1	3	3	9	Yes	50%
4+1	5	3	15	Yes	25%
6+1	7	3	21	Yes	17%
7+5	3	5	15	Yes	71%



StorageGRID requires a minimum of three Storage Nodes per site. To use the 7+5 scheme, each site requires a minimum of four Storage Nodes. Using five Storage Nodes per site is recommended.

When selecting an erasure-coding scheme that provides site protection, balance the relative importance of the following factors:

- **Number of fragments**: Performance and expansion flexibility are generally better when the total number of fragments is lower.
- **Fault tolerance**: Fault tolerance is increased by having more parity segments (that is, when m has a higher value.)
- Network traffic: When recovering from failures, using a scheme with more fragments (that is, a higher total for k+m) creates more network traffic.
- Storage overhead: Schemes with higher overhead require more storage space per object.

For example, when deciding between a 4+2 scheme and 6+3 scheme (which both have 50% storage overhead), select the 6+3 scheme if additional fault tolerance is required. Select the 4+2 scheme if network

resources are constrained. If all other factors are equal, select 4+2 because it has a lower total number of fragments.



If you are unsure of which scheme to use, select 4+2 or 6+3, or contact technical support.

Erasure-coding schemes for one-site storage pools

A one-site storage pool supports all of the erasure-coding schemes defined for three or more sites, provided that the site has enough Storage Nodes.

The minimum number of Storage Nodes required is k+m, but a storage pool with k+m +1 Storage Nodes is recommended. For example, the 2+1 erasure-coding scheme requires a storage pool with a minimum of three Storage Nodes, but four Storage Nodes is recommended.

Erasure-coding scheme (<i>k</i> + <i>m</i>)	Minimum number of Storage Nodes	Recommended number of Storage Nodes	Storage overhead
4+2	6	7	50%
6+2	8	9	33%
8+2	10	11	25%
6+3	9	10	50%
9+3	12	13	33%
2+1	3	4	50%
4+1	5	6	25%
6+1	7	8	17%
7+5	12	13	71%

Advantages, disadvantages, and requirements for erasure coding

Before deciding whether to use replication or erasure coding to protect object data from loss, you should understand the advantages, disadvantages, and the requirements for erasure coding.

Advantages of erasure coding

When compared to replication, erasure coding offers improved reliability, availability, and storage efficiency.

• **Reliability**: Reliability is gauged in terms of fault tolerance—that is, the number of simultaneous failures that can be sustained without loss of data. With replication, multiple identical copies are stored on different nodes and across sites. With erasure coding, an object is encoded into data and parity fragments and distributed across many nodes and sites. This dispersal provides both site and node failure protection.

When compared to replication, erasure coding provides improved reliability at comparable storage costs.

- Availability: Availability can be defined as the ability to retrieve objects if Storage Nodes fail or become inaccessible. When compared to replication, erasure coding provides increased availability at comparable storage costs.
- **Storage efficiency**: For similar levels of availability and reliability, objects protected through erasure coding consume less disk space than the same objects would if protected through replication. For example, a 10 MB object that is replicated to two sites consumes 20 MB of disk space (two copies), while an object that is erasure-coded across three sites with a 6+3 erasure-coding scheme only consumes 15 MB of disk space.



Disk space for erasure-coded objects is calculated as the object size plus the storage overhead. The storage overhead percentage is the number of parity fragments divided by the number of data fragments.

Disadvantages of erasure coding

When compared to replication, erasure coding has the following disadvantages:

- An increased number of Storage Nodes and sites is recommended, depending on the erasure-coding scheme. In contrast, if you replicate object data, you need only one Storage Node for each copy. See Erasure-coding schemes for storage pools containing three or more sites and Erasure-coding schemes for one-site storage pools.
- Increased cost and complexity of storage expansions. To expand a deployment that uses replication, you
 add storage capacity in every location where object copies are made. To expand a deployment that uses
 erasure coding, you must consider both the erasure-coding scheme in use and how full existing Storage
 Nodes are. For example, if you wait until existing nodes are 100% full, you must add at least k+m Storage
 Nodes, but if you expand when existing nodes are 70% full, you can add two nodes per site and still
 maximize usable storage capacity. For more information, see Add storage capacity for erasure-coded
 objects.
- There are increased retrieval latencies when you use erasure coding across geographically distributed sites. The object fragments for an object that is erasure-coded and distributed across remote sites take longer to retrieve over WAN connections than an object that is replicated and available locally (the same site to which the client connects).
- When you use erasure coding across geographically distributed sites, there is higher WAN network traffic usage for retrievals and repairs, especially for frequently retrieved objects or for object repairs over WAN network connections.
- When you use erasure coding across sites, the maximum object throughput declines sharply as network latency between sites increases. This decrease is due to the corresponding decrease in TCP network throughput, which affects how quickly the StorageGRID system can store and retrieve object fragments.
- Higher usage of compute resources.

When to use erasure coding

Erasure coding is best suited for the following requirements:

• Objects greater than 1 MB in size.



Erasure coding is best suited for objects greater than 1 MB. Don't use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.

- Long-term or cold storage for infrequently retrieved content.
- High data availability and reliability.
- Protection against complete site and node failures.
- Storage efficiency.
- Single-site deployments that require efficient data protection with only a single erasure-coded copy rather than multiple replicated copies.
- Multiple-site deployments where the inter-site latency is less than 100 ms.

How object retention is determined

StorageGRID provides options for both grid administrators and individual tenant users to specify how long to store objects. In general, any retention instructions provided by a tenant user take precedence over the retention instructions provided by the grid administrator.

How tenant users control object retention

Tenant users have three primary ways to control how long their objects are stored in StorageGRID:

- If the global S3 Object Lock setting is enabled for the grid, S3 tenant users can create buckets with S3 Object Lock enabled and then use the S3 REST API to specify retain-until-date and legal hold settings for each object version added to that bucket.
 - An object version that is under a legal hold can't be deleted by any method.
 - Before an object version's retain-until-date is reached, that version can't be deleted by any method.
 - Objects in buckets with S3 Object Lock enabled are retained by ILM "forever." However, after its retainuntil-date is reached, an object version can be deleted by a client request or the expiration of the bucket lifecycle. See Manage objects with S3 Object Lock.
- S3 tenant users can add a lifecycle configuration to their buckets that specifies an Expiration action. If a bucket lifecycle exists, StorageGRID stores an object until the date or number of days specified in the Expiration action are met, unless the client deletes the object first. See Create S3 lifecycle configuration.
- An S3 or Swift client can issue a delete object request. StorageGRID always prioritizes client delete requests over S3 bucket lifecycle or ILM when determining whether to delete or retain an object.

How grid administrators control object retention

Grid administrators use ILM placement instructions to control how long objects are stored. When objects are matched by an ILM rule, StorageGRID stores those objects until the last time period in the ILM rule has elapsed. Objects are retained indefinitely if "forever" is specified for the placement instructions.

Regardless of who controls how long objects are retained, ILM settings control what types of object copies (replicated or erasure-coded) are stored and where the copies are located (Storage Nodes, Cloud Storage Pools, or Archive Nodes).

How S3 bucket lifecycle and ILM interact

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

Examples for object retention

To better understand the interactions between S3 Object Lock, bucket lifecycle settings, client delete requests, and ILM, consider the following examples.

Example 1: S3 bucket lifecycle keeps objects longer than ILM

ILM

Store two copies for 1 year (365 days)

Bucket lifecycle

Expire objects in 2 years (730 days)

Result

StorageGRID stores the object for 730 days. StorageGRID uses the bucket lifecycle settings to determine whether to delete or retain an object.



If the bucket lifecycle specifies that objects should be kept longer than specified by ILM, StorageGRID continues to use the ILM placement instructions when determining the number and type of copies to store. In this example, two copies of the object will continue to be stored in StorageGRID from days 366 to 730.

Example 2: S3 bucket lifecycle expires objects before ILM

ILM

```
Store two copies for 2 years (730 days)
```

Bucket lifecycle

Expire objects in 1 year (365 days)

Result

StorageGRID deletes both copies of the object after day 365.

Example 3: Client delete overrides bucket lifecycle and ILM

ILM

Store two copies on Storage Nodes "forever"

Bucket lifecycle

Expire objects in 2 years (730 days)

Client delete request

Issued on day 400

Result

StorageGRID deletes both copies of the object on day 400 in response to the client delete request.

Example 4: S3 Object Lock overrides client delete request

S3 Object Lock

Retain-until-date for an object version is 2026-03-31. A legal hold is not in effect.

Compliant ILM rule

Store two copies on Storage Nodes "forever"

Client delete request

Issued on 2024-03-31

Result

StorageGRID will not delete the object version because the retain-until-date is still 2 years away.

How objects are deleted

StorageGRID can delete objects either in direct response to a client request or automatically as a result of the expiration of an S3 bucket lifecycle or the requirements of the ILM policy. Understanding the different ways that objects can be deleted and how StorageGRID handles delete requests can help you manage objects more effectively.

StorageGRID can use one of two methods to delete objects:

- Synchronous deletion: When StorageGRID receives a client delete request, all object copies are removed immediately. The client is informed that deletion was successful after the copies have been removed.
- Objects are queued for deletion: When StorageGRID receives a delete request, the object is queued for deletion and the client is informed immediately that deletion was successful. Object copies are removed later by background ILM processing.

When deleting objects, StorageGRID uses the method that optimizes delete performance, minimizes potential delete backlogs, and frees space most quickly.

The table summarizes when StorageGRID uses each method.

Method of performing deletion	When used
Objects are queued for deletion	When any of the following conditions are true:
	Automatic object deletion has been triggered by one of the following events:
	 The expiration date or number of days in the lifecycle configuration for an S3 bucket is reached.
	 The last time period specified in an ILM rule elapses.
	Note: Objects in a bucket that has S3 Object Lock enabled can't be deleted if they are under a legal hold or if a retain-until-date has been specified but not yet met.
	 An S3 or Swift client requests deletion and one or more of these conditions is true:
	 Copies can't be deleted within 30 seconds because, for example, an object location is temporarily unavailable.
	 Background deletion queues are idle.

Method of performing deletion	When used
Objects are removed immediately (synchronous deletion)	 When an S3 or Swift client makes a delete request and all of the following conditions are met: All copies can be removed within 30 seconds. Background deletion queues contain objects to process.

When S3 or Swift clients make delete requests, StorageGRID begins by adding objects to the delete queue. It then switches to performing synchronous deletion. Making sure that the background deletion queue has objects to process allows StorageGRID to process deletes more efficiently, especially for low concurrency clients, while helping to prevent client delete backlogs.

Time required to delete objects

The way that StorageGRID deletes objects can affect how the system appears to perform:

- When StorageGRID performs synchronous deletion, it can take StorageGRID up to 30 seconds to return a result to the client. This means that deletion can appear to be happening more slowly, even though copies are actually being removed more quickly than they are when StorageGRID queues objects for deletion.
- If you are closely monitoring delete performance during a bulk delete, you might notice that the deletion rate appears to be slow after a certain number of objects have been deleted. This change occurs when StorageGRID shifts from queuing objects for deletion to performing synchronous deletion. The apparent reduction in the deletion rate does not mean that object copies are being removed more slowly. On the contrary, it indicates that on average, space is now being freed more quickly.

If you are deleting large numbers of objects and your priority is to free space quickly, consider using a client request to delete objects rather than deleting them using ILM or other methods. In general, space is freed more quickly when deletion is performed by clients because StorageGRID can use synchronous deletion.

The amount of time required to free space after an object is deleted depends on several factors:

- Whether object copies are synchronously removed or are queued for removal later (for client delete requests).
- Other factors such as the number of objects in the grid or the availability of grid resources when object copies are queued for removal (for both client deletes and other methods).

How S3 versioned objects are deleted

When versioning is enabled for an S3 bucket, StorageGRID follows Amazon S3 behavior when responding to delete requests, whether those requests come from an S3 client, the expiration of an S3 bucket lifecycle, or the requirements of the ILM policy.

When objects are versioned, object delete requests don't delete the current version of the object and don't free space. Instead, an object delete request creates a zero-byte delete marker as the current version of the object, which makes the previous version of the object "noncurrent." An object delete marker becomes an expired object delete marker when it is the current version and there are no noncurrent versions.

Even though the object has not been removed, StorageGRID behaves as though the current version of the object is no longer available. Requests to that object return 404 NotFound. However, because noncurrent object data has not been removed, requests that specify a noncurrent version of the object can succeed.

To free space when deleting versioned objects, or to remove delete markers, use one of the following:

- **S3 client request**: Specify the object version ID in the S3 DELETE Object request (DELETE /object?versionId=ID). Keep in mind that this request only removes object copies for the specified version (the other versions are still taking up space).
- Bucket lifecycle: Use the NoncurrentVersionExpiration action in the bucket lifecycle configuration. When the number of NoncurrentDays specified is met, StorageGRID permanently removes all copies of noncurrent object versions. These object versions can't be recovered.

The NewerNoncurrentVersions action in the bucket lifecycle configuration specifies the number of noncurrent versions retained in a versioned S3 bucket. If there are more noncurrent versions than NewerNoncurrentVersions specifies, StorageGRID removes the older versions when the NoncurrentDays value has elapsed. The NewerNoncurrentVersions threshold overrides lifecycle rules provided by ILM, meaning that a noncurrent object with a version within the NewerNoncurrentVersions threshold is retained if ILM requests its deletion.

To remove expired object delete markers use the Expiration action with one of the following tags: ExpiredObjectDeleteMarker, Days, or Date.

- ILM: Clone an active policy and add two ILM rules to the new policy:
 - First rule: Use "Noncurrent time" as the Reference time to match the noncurrent versions of the object. In Step 1 (Enter details) of the Create an ILM rule wizard, select Yes for the question, "Apply this rule to older object versions only (in S3 buckets with versioning enabled)?"
 - Second rule: Use **Ingest time** to match the current version. The "Noncurrent time" rule must appear in the policy above the **Ingest time** rule.



ILM cannot be used to remove current object delete markers. Use an S3 client request or S3 Bucket Lifecycle to remove current object delete markers.

• **Delete objects in bucket**: Use the tenant manager to delete all object versions, including delete markers, from a bucket.

When a versioned object is deleted, StorageGRID creates a zero-byte delete marker as the current version of the object. All objects and delete markers must be removed before a versioned bucket can be deleted.

- Delete markers created in StorageGRID 11.7 or earlier can only be removed through S3 client requests, they are not removed by ILM, bucket lifecycle rules, or Delete objects in bucket operations.
- Delete markers from a bucket that was created in StorageGRID 11.8 or later can be removed by ILM, bucket lifecycle rules, Delete objects in bucket operations, or an explicit S3 client deletion. Expired delete markers in StorageGRID 11.8 or later must be removed by bucket lifecycle rules or by an explicit S3 client request with a version ID specified.

Related information

- Use S3 REST API
- Example 4: ILM rules and policy for S3 versioned objects

Create and assign storage grades

Storage grades identify the type of storage used by a Storage Node. You can create storage grades if you want ILM rules to place certain objects on certain Storage Nodes.

Before you begin

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

About this task

When you first install StorageGRID, the **Default** storage grade is automatically assigned to every Storage Node in your system. As required, you can optionally define custom storage grades and assign them to different Storage Nodes.

Using custom storage grades allows you to create ILM storage pools that contain only a specific type of Storage Node. For example, you might want certain objects to be stored on your fastest Storage Nodes, such as StorageGRID all-flash storage appliances.



Storage Nodes can be configured during installation to contain only object metadata and not object data. Metadata-only Storage Nodes can't be assigned a storage grade. For more information, see Types of Storage Nodes.

If storage grade is not a concern (for example, all Storage Nodes are identical), you can skip this procedure and use the **includes all storage grades** selection for the storage grade when you create storage pools. Using this selection ensures that the storage pool will include every Storage Node at the site, regardless of its storage grade.



Don't create more storage grades than necessary. For example, don't create a storage grade for each Storage Node. Instead, assign each storage grade to two or more nodes. Storage grades assigned to only one node can cause ILM backlogs if that node becomes unavailable.

Steps

- 1. Select ILM > Storage grades.
- 2. Define custom storage grades:
 - a. For each custom storage grade you want to add, select **Insert** 🚯 to add a row.
 - b. Enter a descriptive label.



Storage Grade Definitions			
Storage Grade	Label		Actions
0	Default		
1	disk		10
Storage Grades			1
LDR		Storage Grade	Actions
Data Center 1/DC1-S1/LDR		Default	0
Data Center 1/DC1-S2/LDR		Default	1

Data Center 1/DC1-S2/LDR	Default	0
Data Center 1/DC1-S3/LDR	Default	0
Data Center 2/DC2-S1/LDR	Default	0
Data Center 2/DC2-S2/LDR	Default	0
Data Center 2/DC2-S3/LDR	Default	1
Data Center 3/DC3-S1/LDR	Default	0
Data Center 3/DC3-S2/LDR	Default	1
Data Center 3/DC3-S3/LDR	Default	0



c. Select Apply Changes.

d. Optionally, if you need to modify a saved label, select **Edit** and select **Apply Changes**.



You can't delete storage grades.

3. Assign new storage grades to Storage Nodes:

- a. Locate the Storage Node in the LDR list, and select its Edit icon 🥢.
- b. Select the appropriate storage grade from the list.

Storage Grades

LDR	Storage Grade	Actions
Data Center 1/DC1-S1/LDR	Default	1
Data Center 1/DC1-S2/LDR	Default disk	Ø
Data Center 1/DC1-S3/LDR	Default 6	1
Data Center 2/DC2-S1/LDR	Default	Ø
Data Center 2/DC2-S2/LDR	Default	1
Data Center 2/DC2-S3/LDR	Default	Ø
Data Center 3/DC3-S1/LDR	Default	1
Data Center 3/DC3-S2/LDR	Default	1
Data Center 3/DC3-S3/LDR	Default	1





Assign a storage grade to a given Storage Node only once. A Storage Node recovered from failure maintains the previously assigned storage grade. Don't change this assignment after the ILM policy is activated. If the assignment is changed, data is stored based on the new storage grade.

c. Select Apply Changes.

Use storage pools

What is a storage pool?

A storage pool is a logical grouping of Storage Nodes or Archive Nodes.

When you install StorageGRID, one storage pool per site is automatically created. You can configure additional storage pools as needed for your storage requirements.



Storage Nodes can be configured during installation to contain object data and object metadata, or only object metadata. Metadata-only Storage Nodes can't be used in storage pools. For more information, see Types of Storage Nodes.



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

Storage pools have two attributes:

- Storage grade: For Storage Nodes, the relative performance of backing storage.
- Site: The data center where objects will be stored.

Storage pools are used in ILM rules to determine where object data is stored and the type of storage used. When you configure ILM rules for replication, you select one or more storage pools that include either Storage Nodes or Archive Nodes. When you create erasure-coding profiles, you select a storage pool that includes

Storage Nodes.

Guidelines for creating storage pools

Configure and use storage pools to protect against data loss by distributing data across multiple sites. Replicated copies and erasure-coded copies require different storage pool configurations.

See Examples of enabling site-loss protection using replication and erasure coding.

Guidelines for all storage pools

- Keep storage pool configurations as simple as possible. Don't create more storage pools than necessary.
- Create storage pools with as many nodes as possible. Each storage pool should contain two or more nodes. A storage pool with insufficient nodes can cause ILM backlogs if a node becomes unavailable.
- Avoid creating or using storage pools that overlap (contain one or more of the same nodes). If storage pools overlap, more than one copy of object data might be saved on the same node.
- In general, don't use the All Storage Nodes storage pool (StorageGRID 11.6 and earlier) or the All Sites site. These items are automatically updated to include any new sites you add in an expansion, which might not be the behavior you want.

Guidelines for storage pools used for replicated copies

• For site-loss protection using replication, specify one or more site-specific storage pools in the placement instructions for each ILM rule.

One storage pool is automatically created for each site during StorageGRID installation.

Using a storage pool for each site ensures that replicated object copies are placed exactly where you expect (for example, one copy of every object at each site for site-loss protection).

- If you add a site in an expansion, create a new storage pool that contains only the new site. Then, update ILM rules to control which objects are stored on the new site.
- If the number of copies is less than the number of storage pools, the system distributes the copies to balance disk usage among the pools.
- If the storage pools overlap (contain the same Storage Nodes), all copies of the object might be saved at only one site. You must ensure that the selected storage pools don't contain the same Storage Nodes.

Guidelines for storage pools used for erasure-coded copies

- For site-loss protection using erasure coding, create storage pools that consist of at least three sites. If a storage pool includes only two sites, you can't use that storage pool for erasure coding. No erasure-coding schemes are available for a storage pool that has two sites.
- The number of Storage Nodes and sites contained in the storage pool determine which erasure-coding schemes are available.
- If possible, a storage pool should include more than the minimum number of Storage Nodes required for the erasure-coding scheme you select. For example, if you use a 6+3 erasure-coding scheme, you must have at least nine Storage Nodes. However, having at least one additional Storage Node per site is recommended.
- Distribute Storage Nodes across sites as evenly as possible. For example, to support a 6+3 erasure-coding scheme, configure a storage pool that includes at least three Storage Nodes at three sites.

• If you have high throughput requirements, using a storage pool that includes multiple sites is not recommended if the network latency between sites is greater than 100 ms. As latency increases, the rate at which StorageGRID can create, place, and retrieve object fragments decreases sharply due to the decrease in TCP network throughput.

The decrease in throughput affects the maximum achievable rates of object ingest and retrieval (when Balanced or Strict are selected as the ingest behavior) or could lead to ILM queue backlogs (when Dual commit is selected as the ingest behavior). See ILM rule ingest behavior.



If your grid includes only one site, you are prevented from using the All Storage Nodes storage pool (StorageGRID 11.6 and earlier) or the All Sites default site in an erasure-coding profile. This behavior prevents the profile from becoming invalid if a second site is added.

• You can't use Archive Nodes for erasure-coded data.

Guidelines for storage pools used for archived copies

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

The Cloud Tiering - Simple Storage Service (S3) option is also deprecated. If you are currently using an Archive Node with this option, migrate your objects to a Cloud Storage Pool instead.

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

- You can't create a storage pool that includes both Storage Nodes and Archive Nodes. Archived copies require a storage pool that only includes Archive Nodes.
- When using a storage pool that includes Archive Nodes, you should also maintain at least one replicated or erasure-coded copy on a storage pool that includes Storage Nodes.
- If the global S3 Object Lock setting is enabled and you are creating a compliant ILM rule, you can't use a storage pool that includes Archive Nodes. See the instructions for managing objects with S3 Object Lock.
- If an Archive Node's Target Type is Cloud Tiering Simple Storage Service (S3), the Archive Node must be in its own storage pool.

Enable site-loss protection

If your StorageGRID deployment includes more than one site, you can use replication and erasure coding with appropriately configured storage pools to enable site-loss protection.

Replication and erasure coding require different storage pool configurations:

- To use replication for site-loss protection, use the site-specific storage pools that are automatically created during StorageGRID installation. Then create ILM rules with placement instructions that specify multiple storage pools so that one copy of each object will be placed at each site.
- To use erasure coding for site-loss protection, create storage pools that consist of multiple sites. Then create ILM rules that use one storage pool consisting of multiple sites and any available erasure-coding schema.



When configuring your StorageGRID deployment for site-loss protection, you must also take into account the effects of ingest options and consistency.

Replication example

By default, one storage pool is created for each site during StorageGRID installation. Having storage pools that consist of only one site enables you to configure ILM rules that use replication for site-loss protection. In this example:

- Storage pool 1 contains Site 1
- Storage pool 2 contains Site 2
- The ILM rule contains two placements:
 - Store objects by replicating 1 copy at Site 1
 - Store objects by replicating 1 copy at Site 2

ILM rule placements:



If one site is lost, copies of the objects are available at the other site.

Erasure coding example

Having storage pools that consist of more than one site per storage pool enables you to configure ILM rules that use erasure coding for site-loss protection. In this example:

- Storage pool 1 contains Sites 1 through 3
- The ILM rule contains one placement: Store objects by erasure coding using a 4+2 EC scheme at Storage pool 1, which contains three sites

ILM rule placements:

erasure coding

V



In this example:

- The ILM rule uses a 4+2 erasure-coding scheme.
- Each object is sliced into four equal data fragments, and two parity fragments are computed from the object data.
- Each of the six fragments is stored on a different node across three data center sites to provide data protection for node failures or site loss.



Erasure coding is allowed in storage pools containing any number of sites *except* two sites.

ILM rule using 4+2 erasure-coding scheme:



If one site is lost, data can still be recovered:



Create a storage pool

You create storage pools to determine where the StorageGRID system stores object data and the type of storage used. Each storage pool includes one or more sites and one or more storage grades.



When you install StorageGRID 11.8 on a new grid, storage pools are automatically created for each site. However, if you initially installed StorageGRID 11.6 or earlier, storage pools aren't automatically created for each site.

If you want to create Cloud Storage Pools to store object data outside of your StorageGRID system, see the information about using Cloud Storage Pools.

Before you begin

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

About this task

Storage pools determine where object data is stored. The number of storage pools you need depends on the number of sites in your grid and on the types of copies you want: replicated or erasure-coded.

- For replication and single-site erasure coding, create a storage pool for each site. For example, if you want to store replicated object copies at three sites, create three storage pools.
- For erasure coding at three or more sites, create one storage pool that includes an entry for each site. For example, if you want to erasure code objects across three sites, create one storage pool.



Don't include the All Sites site in a storage pool that will be used in an erasure-coding profile. Instead, add a separate entry to the storage pool for each site that will store erasure-coded data. See this step for an example.

• If you have more than one storage grade, don't create a storage pool that includes different storage grades

at a single site. See the Guidelines for creating storage pools.

Steps

1. Select ILM > Storage pools.

The Storage pools tab lists all defined storage pools.



For new installations of StorageGRID 11.6 or earlier, the All Storage Nodes storage pool is automatically updated whenever you add new data center sites. Don't use this pool in ILM rules.

- 2. To create a new storage pool, select Create.
- 3. Enter a unique name for the storage pool. Use a name that will be easy to identify when you configure erasure-coding profiles and ILM rules.
- 4. From the Site drop-down list, select a site for this storage pool.

When you select a site, the number of Storage Nodes and Archive Nodes in the table are automatically updated.

In general, don't use the All Sites site in any storage pool. ILM rules that use an All Sites storage pool place objects at any available site, giving you less control of object placement. Also, an All Sites storage pool uses the Storage Nodes at a new site immediately, which might not be the behavior you expect.

5. From the **Storage grade** drop-down list, select the type of storage that will be used if an ILM rule uses this storage pool.

The storage grade, *includes all storage grades*, includes all Storage Nodes at the selected site. The default Archive Nodes storage grade includes all Archive Nodes at the selected site. If you created additional storage grades for the Storage Nodes in your grid, they are listed in the drop-down.

6. If you want to use the storage pool in a multi-site erasure-coding profile, select **Add more nodes** to add an entry for each site to the storage pool.



You are prevented from creating duplicate entries or from creating a storage pool that includes both the Archive Nodes storage grade and any storage grade that contains Storage Nodes.

You are warned if you add more than one entry with different storage grades for a site.

To remove an entry, select the delete icon \mathbf{X} .

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

The new storage pool is added to the list.

View storage pool details

You can view the details of a storage pool to determine where the storage pool is used and to see which nodes and storage grades are included.

Before you begin

• You are signed in to the Grid Manager using a supported web browser.

· You have specific access permissions.

Steps

1. Select ILM > Storage pools.

The Storage pools table includes the following information for each storage pool that includes Storage Nodes:

- Name: The unique display name of the storage pool.
- Node count: The number of nodes in the storage pool.
- **Storage usage**: The percentage of the total usable space that has been used for object data on this node. This value does not include object metadata.
- **Total capacity**: The size of the storage pool, which equals the total amount of usable space for object data for all nodes in the storage pool.
- **ILM usage**: How the storage pool is currently being used. A storage pool might be unused or it might be used in one or more ILM rules, erasure-coding profiles, or both.



You can't remove a storage pool if it is being used.

2. To view details about a specific storage pool, select its name.

The details page for the storage pool appears.

3. View the **Nodes** tab to learn about the Storage Nodes or Archive Nodes included in the storage pool.

The table includes the following information for each node:

- Node name
- Site name
- · Storage grade
- Storage usage: The percentage of the total usable space for object data that has been used for the Storage Node. This field is not visible for Archive Node pools.



The same Storage usage (%) value is also shown in the Storage Used - Object Data chart for each Storage Node (select **NODES** > *Storage Node* > *Storage*).

- 4. Select the **ILM usage** tab to determine if the storage pool is currently being used in any ILM rules or erasure-coding profiles.
- 5. Optionally, go to the **ILM rules page** to learn about and manage any rules that use the storage pool.

See the instructions for working with ILM rules.

Edit storage pool

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

Before you begin

• You are signed in to the Grid Manager using a supported web browser.

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

About this task

If you are adding a new site or storage grade to a storage pool that is used in the active ILM policy, be aware that the Storage Nodes in the new site or storage grade will not be used automatically. To force StorageGRID to use a new site or storage grade, you must activate a new ILM policy after saving the edited storage pool.

Steps

- 1. Select **ILM > Storage pools**.
- 2. Select the checkbox for the storage pool you want to edit.

You can't edit the All Storage Nodes storage pool (StorageGRID 11.6 and earlier).

3. Select Edit.

- 4. As required, change the storage pool name.
- 5. As required, select other sites and storage grades.



You are prevented from changing the site or storage grade if the storage pool is used in an erasure-coding profile and the change would cause the erasure-coding scheme to become invalid. For example, if a storage pool used in a erasure-coding profile currently includes a storage grade with only one site, you are prevented from using a storage grade with two sites because the change would make the erasure-coding scheme invalid.

6. Select Save.

After you finish

If you added a new site or storage grade to a storage pool used in the active ILM policy, activate a new ILM policy to force StorageGRID to use the new site or storage grade. For example, clone your existing ILM policy and then activate the clone. See Work with ILM rules and ILM policies.

Remove a storage pool

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

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the required access permissions.

Steps

- 1. Select ILM > Storage pools.
- 2. Look at the ILM usage column in the table to determine whether you can remove the storage pool.

You can't remove a storage pool if it is being used in an ILM rule or in an erasure-coding profile. As required, select *storage pool name* > ILM usage to determine where the storage pool is used.

3. If the storage pool you want to remove is not being used, select the checkbox.

- 4. Select Remove.
- 5. Select OK.

Use Cloud Storage Pools

What is a Cloud Storage Pool?

A Cloud Storage Pool lets you use ILM to move object data outside of your StorageGRID system. For example, you might want to move infrequently accessed objects to lower-cost cloud storage, such as Amazon S3 Glacier, S3 Glacier Deep Archive, Google Cloud, or the Archive access tier in Microsoft Azure Blob storage. Or, you might want to maintain a cloud backup of StorageGRID objects to enhance disaster recovery.

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



Moving objects from an Archive Node to an external archival storage system through the S3 API is deprecated and has been replaced by ILM Cloud Storage Pools, which offer more functionality. If you are currently using an Archive Node with the Cloud Tiering - Simple Storage Service (S3) option, migrate your objects to a Cloud Storage Pool instead.

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

	Storage pool	Cloud Storage Pool
How is it created?	Using the ILM > Storage pools option in Grid Manager.	Using the ILM > Storage pools > Cloud Storage Pools option in Grid Manager. You must set up the external bucket or container before you can create the Cloud Storage Pool.
How many pools can you create?	Unlimited.	Up to 10.

	Storage pool	Cloud Storage Pool
Where are objects Nodes or Archive Nodes stored? Within StorageGRID.	On one or more Storage Nodes or Archive Nodes	In an Amazon S3 bucket, Azure Blob storage container, or Google Cloud that is external to the StorageGRID system.
	within Storage GRID.	If the Cloud Storage Pool is an Amazon S3 bucket:
	 You can optionally configure a bucket lifecycle to transition objects to low-cost, long-term storage, such as Amazon S3 Glacier or S3 Glacier Deep Archive. The external storage system must support the Glacier storage class and the S3 RestoreObject API. 	
		 You can create Cloud Storage Pools for use with AWS Commercial Cloud Services (C2S), which supports the AWS Secret Region.
		If the Cloud Storage Pool is an Azure Blob storage container, StorageGRID transitions the object to the Archive tier.
		Note: In general, don't configure Azure Blob storage lifecycle management for the container used for a Cloud Storage Pool. RestoreObject operations on objects in the Cloud Storage Pool can be affected by the configured lifecycle.
What controls object placement?	An ILM rule in the active ILM policies.	An ILM rule in the active ILM policies.
Which data protection method is used?	Replication or erasure coding.	Replication.
How many copies of each object are	Multiple.	One copy in the Cloud Storage Pool and, optionally, one or more copies in StorageGRID. Note: You can't store an object in more than one Cloud Storage
allowed?		Fool at any given time.
What are the advantages ?	Objects are quickly accessible at any time.	Low-cost storage.
		Note : FabricPool data can't be tiered to Cloud Storage Pools. Objects with S3 Object Lock enabled can't be placed in Cloud Storage Pools.

Lifecycle of a Cloud Storage Pool object

Before implementing Cloud Storage Pools, review the lifecycle of objects that are stored in each type of Cloud Storage Pool.

S3: Lifecycle of a Cloud Storage Pool object

The steps describe the lifecycle stages of an object that is stored in an S3 Cloud Storage Pool.



"Glacier" refers to both the Glacier storage class and the Glacier Deep Archive storage class, with one exception: the Glacier Deep Archive storage class does not support the Expedited restore tier. Only Bulk or Standard retrieval is supported.



The Google Cloud Platform (GCP) supports object retrieval from long-term storage without requiring a POST Restore operation.

1. Object stored in StorageGRID

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

2. Object moved to S3 Cloud Storage Pool

- When the object is matched by an ILM rule that uses an S3 Cloud Storage Pool as its placement location, StorageGRID moves the object to the external S3 bucket specified by the Cloud Storage Pool.
- When the object has been moved to the S3 Cloud Storage Pool, the client application can retrieve it using an S3 GetObject request from StorageGRID, unless the object has been transitioned to Glacier storage.

3. Object transitioned to Glacier (non-retrievable state)

 Optionally, the object can be transitioned to Glacier storage. For example, the external S3 bucket might use lifecycle configuration to transition an object to Glacier storage immediately or after some number of days.



If you want to transition objects, you must create a lifecycle configuration for the external S3 bucket, and you must use a storage solution that implements the Glacier storage class and supports the S3 RestoreObject API.



Don't use Cloud Storage Pools for objects that have been ingested by Swift clients. Swift does not support RestoreObject requests, so StorageGRID will not be able to retrieve any Swift objects that have been transitioned to S3 Glacier storage. Issuing a Swift GET object request to retrieve these objects will fail (403 Forbidden).

 During the transition, the client application can use an S3 HeadObject request to monitor the object's status.

4. Object restored from Glacier storage

If an object has been transitioned to Glacier storage, the client application can issue an S3 RestoreObject request to restore a retrievable copy to the S3 Cloud Storage Pool. The request specifies how many days the copy should be available in the Cloud Storage Pool and the data-access tier to use for the restore operation (Expedited, Standard, or Bulk). When the expiration date of the retrievable copy is reached, the copy is automatically returned to a non-retrievable state.



If one or more copies of the object also exist on Storage Nodes within StorageGRID, there is no need to restore the object from Glacier by issuing a RestoreObject request. Instead, the local copy can be retrieved directly, using a GetObject request.

5. Object retrieved

Once an object has been restored, the client application can issue a GetObject request to retrieve the restored object.

Azure: Lifecycle of a Cloud Storage Pool object

The steps describe the lifecycle stages of an object that is stored in an Azure Cloud Storage Pool.

1. Object stored in StorageGRID

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

2. Object moved to Azure Cloud Storage Pool

When the object is matched by an ILM rule that uses an Azure Cloud Storage Pool as its placement location, StorageGRID moves the object to the external Azure Blob storage container specified by the Cloud Storage Pool.



Don't use Cloud Storage Pools for objects that have been ingested by Swift clients. Swift does not support RestoreObject requests, so StorageGRID will not be able to retrieve any Swift objects that have been transitioned to the Azure Blob storage Archive tier. Issuing a Swift GET object request to retrieve these objects will fail (403 Forbidden).

3. Object transitioned to Archive tier (non-retrievable state)

Immediately after moving the object to the Azure Cloud Storage Pool, StorageGRID automatically transitions the object to the Azure Blob storage Archive tier.

4. Object restored from Archive tier

If an object has been transitioned to the Archive tier, the client application can issue an S3 RestoreObject request to restore a retrievable copy to the Azure Cloud Storage Pool.

When StorageGRID receives the RestoreObject, it temporarily transitions the object to the Azure Blob storage Cool tier. As soon as the expiration date in the RestoreObject request is reached, StorageGRID transitions the object back to the Archive tier.



If one or more copies of the object also exist on Storage Nodes within StorageGRID, there is no need to restore the object from the Archive access tier by issuing a RestoreObject request. Instead, the local copy can be retrieved directly, using a GetObject request.

5. Object retrieved

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

Related information

Use S3 REST API

When to use Cloud Storage Pools

Using Cloud Storage Pools, you can back up or tier data to an external location. Additionally, you can back up or tier data to more than one cloud.

Back up StorageGRID data to external location

You can use a Cloud Storage Pool to back up StorageGRID objects to an external location.

If the copies in StorageGRID are inaccessible, the object data in the Cloud Storage Pool can be used to serve client requests. However, you might need to issue S3 RestoreObject request to access the backup object copy in the Cloud Storage Pool.

The object data in a Cloud Storage Pool can also be used to recover data lost from StorageGRID because of a storage volume or Storage Node failure. If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID temporarily restores the object and creates a new copy on the recovered Storage Node.

To implement a backup solution:

- 1. Create a single Cloud Storage Pool.
- 2. Configure an ILM rule that simultaneously stores object copies on Storage Nodes (as replicated or erasurecoded copies) and a single object copy in the Cloud Storage Pool.
- 3. Add the rule to your ILM policy. Then, simulate and activate the policy.

Tier data from StorageGRID to external location

You can use a Cloud Storage Pool to store objects outside of the StorageGRID system. For example, suppose you have a large number of objects that you need to retain, but you expect to access those objects rarely, if ever. You can use a Cloud Storage Pool to tier the objects to lower-cost storage and to free up space in StorageGRID.

To implement a tiering solution:

- 1. Create a single Cloud Storage Pool.
- 2. Configure an ILM rule that moves rarely used objects from Storage Nodes to the Cloud Storage Pool.
- 3. Add the rule to your ILM policy. Then, simulate and activate the policy.

Maintain multiple cloud endpoints

You can configure multiple Cloud Storage Pool endpoints if you want to tier or back up object data to more than one cloud. The filters in your ILM rules let you specify which objects are stored in each Cloud Storage Pool. For example, you might want to store objects from some tenants or buckets in Amazon S3 Glacier and objects from other tenants or buckets in Azure Blob storage. Or, you might want to move data between Amazon S3 Glacier and Azure Blob storage.



When using multiple Cloud Storage Pool endpoints, keep in mind that an object can be stored in only one Cloud Storage Pool at a time.

To implement multiple cloud endpoints:

- 1. Create up to 10 Cloud Storage Pools.
- 2. Configure ILM rules to store the appropriate object data at the appropriate time in each Cloud Storage

Pool. For example, store objects from bucket A in Cloud Storage Pool A, and store objects from bucket B in Cloud Storage Pool B. Or, store objects in Cloud Storage Pool A for some amount of time and then move them to Cloud Storage Pool B.

3. Add the rules to your ILM policy. Then, simulate and activate the policy.

Considerations for Cloud Storage Pools

If you plan to use a Cloud Storage Pool to move objects out of the StorageGRID system, you must review the considerations for configuring and using Cloud Storage Pools.

General considerations

- In general, cloud archival storage, such as Amazon S3 Glacier or Azure Blob storage, is an inexpensive place to store object data. However, the costs to retrieve data from cloud archival storage are relatively high. To achieve the lowest overall cost, you must consider when and how often you will access the objects in the Cloud Storage Pool. Using a Cloud Storage Pool is recommended only for content that you expect to access infrequently.
- Don't use Cloud Storage Pools for objects that have been ingested by Swift clients. Swift does not support RestoreObject requests, so StorageGRID will not be able to retrieve any Swift objects that have been transitioned to S3 Glacier storage or the Azure Blob storage Archive tier. Issuing a Swift GET object request to retrieve these objects will fail (403 Forbidden).
- Using Cloud Storage Pools with FabricPool is not supported because of the added latency to retrieve an object from the Cloud Storage Pool target.
- Objects with S3 Object Lock enabled can't be placed in Cloud Storage Pools.
- If the destination S3 bucket for a Cloud Storage Pool has S3 Object Lock enabled, the attempt to configure bucket replication (PutBucketReplication) will fail with an AccessDenied error.

Considerations for the ports used for Cloud Storage Pools

To ensure that the ILM rules can move objects to and from the specified Cloud Storage Pool, you must configure the network or networks that contain your system's Storage Nodes. You must ensure that the following ports can communicate with the Cloud Storage Pool.

By default, Cloud Storage Pools use the following ports:

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

You can specify a different port when you create or edit a Cloud Storage Pool.

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

Considerations for costs

Access to storage in the cloud using a Cloud Storage Pool requires network connectivity to the cloud. You must consider the cost of the network infrastructure you will use to access the cloud and provision it appropriately, based on the amount of data you expect to move between StorageGRID and the cloud using the Cloud Storage Pool.

When StorageGRID connects to the external Cloud Storage Pool endpoint, it issues various requests to monitor connectivity and to ensure it can perform the required operations. While some additional costs will be

associated with these requests, the cost of monitoring a Cloud Storage Pool should only be a small fraction of the overall cost of storing objects in S3 or Azure.

More significant costs might be incurred if you need to move objects from an external Cloud Storage Pool endpoint back to StorageGRID. Objects might be moved back to StorageGRID in either of these cases:

- The only copy of the object is in a Cloud Storage Pool and you decide to store the object in StorageGRID instead. In this case, you reconfigure your ILM rules and policy. When ILM evaluation occurs, StorageGRID issues multiple requests to retrieve the object from the Cloud Storage Pool. StorageGRID then creates the specified number of replicated or erasure-coded copies locally. After the object is moved back to StorageGRID, the copy in the Cloud Storage Pool is deleted.
- Objects are lost because of Storage Node failure. If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID temporarily restores the object and creates a new copy on the recovered Storage Node.



When objects are moved back to StorageGRID from a Cloud Storage Pool, StorageGRID issues multiple requests to the Cloud Storage Pool endpoint for each object. Before moving large numbers of objects, contact technical support for help in estimating the time frame and associated costs.

S3: Permissions required for the Cloud Storage Pool bucket

The bucket policy for the external S3 bucket used for a Cloud Storage Pool must grant StorageGRID permission to move an object to the bucket, get an object's status, restore an object from Glacier storage when required, and more. Ideally, StorageGRID should have full-control access to the bucket (s3:*); however, if this is not possible, the bucket policy must grant the following S3 permissions to StorageGRID:

- s3:AbortMultipartUpload
- s3:DeleteObject
- s3:GetObject
- s3:ListBucket
- s3:ListBucketMultipartUploads
- s3:ListMultipartUploadParts
- s3:PutObject
- s3:RestoreObject

S3: Considerations for the external bucket's lifecycle

The movement of objects between StorageGRID and the external S3 bucket specified in the Cloud Storage Pool is controlled by ILM rules and the active ILM policies in StorageGRID. In contrast, the transition of objects from the external S3 bucket specified in the Cloud Storage Pool to Amazon S3 Glacier or S3 Glacier Deep Archive (or to a storage solution that implements the Glacier storage class) is controlled by that bucket's lifecycle configuration.

If you want to transition objects from the Cloud Storage Pool, you must create the appropriate lifecycle configuration on the external S3 bucket, and you must use a storage solution that implements the Glacier storage class and supports the S3 RestoreObject API.

For example, suppose you want all objects that are moved from StorageGRID to the Cloud Storage Pool to be transitioned to Amazon S3 Glacier storage immediately. You would create a lifecycle configuration on the
external S3 bucket that specifies a single action (Transition) as follows:

```
<LifecycleConfiguration>
<Rule>
<ID>Transition Rule</ID>
<Filter>
<Prefix></Prefix>
</Filter>
<Status>Enabled</Status>
<Transition>
<Days>0</Days>
<StorageClass>GLACIER</StorageClass>
</Transition>
</Rule>
</LifecycleConfiguration>
```

This rule would transition all bucket objects to Amazon S3 Glacier on the day they were created (that is, on the day they were moved from StorageGRID to the Cloud Storage Pool).



When configuring the external bucket's lifecycle, never use **Expiration** actions to define when objects expire. Expiration actions cause the external storage system to delete expired objects. If you later attempt to access an expired object from StorageGRID, the deleted object will not be found.

If you want to transition objects in the Cloud Storage Pool to S3 Glacier Deep Archive (instead of to Amazon S3 Glacier), specify <StorageClass>DEEP_ARCHIVE</StorageClass> in the bucket lifecycle. However, be aware that you can't use the Expedited tier to restore objects from S3 Glacier Deep Archive.

Azure: Considerations for Access tier

When you configure an Azure storage account, you can set the default Access tier to Hot or Cool. When creating a storage account for use with a Cloud Storage Pool, you should use the Hot tier as the default tier. Even though StorageGRID immediately sets the tier to Archive when it moves objects to the Cloud Storage Pool, using a default setting of Hot ensures that you will not be charged an early deletion fee for objects removed from the Cool tier before the 30-day minimum.

Azure: Lifecycle management not supported

Don't use Azure Blob storage lifecycle management for the container used with a Cloud Storage Pool. The lifecycle operations might interfere with Cloud Storage Pool operations.

Related information

Create a Cloud Storage Pool

Compare Cloud Storage Pools and CloudMirror replication

As you begin using Cloud Storage Pools, it might be helpful to understand the similarities and differences between Cloud Storage Pools and the StorageGRID CloudMirror replication service.

	Cloud Storage Pool	CloudMirror replication service
What is the primary purpose?	Acts as an archive target. The object copy in the Cloud Storage Pool can be the only copy of the object, or it can be an additional copy. That is, instead of keeping two copies onsite, you can keep one copy within StorageGRID and send a copy to the Cloud Storage Pool.	Enables a tenant to automatically replicate objects from a bucket in StorageGRID (source) to an external S3 bucket (destination). Creates an independent copy of an object in an independent S3 infrastructure.
How is it set up?	Defined in the same way as storage pools, using the Grid Manager or the Grid Management API. Can be selected as the placement location in an ILM rule. While a storage pool consists of a group of Storage Nodes, a Cloud Storage Pool is defined using a remote S3 or Azure endpoint (IP address, credentials, and so on).	A tenant user configures CloudMirror replication by defining a CloudMirror endpoint (IP address, credentials, and so on) using the Tenant Manager or the S3 API. After the CloudMirror endpoint is set up, any bucket owned by that tenant account can be configured to point to the CloudMirror endpoint.
Who is responsible for setting it up?	Typically, a grid administrator	Typically, a tenant user
What is the destination?	 Any compatible S3 infrastructure (including Amazon S3) Azure Blob Archive tier Google Cloud Platform (GCP) 	 Any compatible S3 infrastructure (including Amazon S3) Google Cloud Platform (GCP)
What causes objects to be moved to the destination?	One or more ILM rules in the active ILM policies. The ILM rules define which objects StorageGRID moves to the Cloud Storage Pool and when the objects are moved.	The act of ingesting a new object into a source bucket that has been configured with a CloudMirror endpoint. Objects that existed in the source bucket before the bucket was configured with the CloudMirror endpoint aren't replicated, unless they are modified.
How are objects retrieved?	Applications must make requests to StorageGRID to retrieve objects that have been moved to a Cloud Storage Pool. If the only copy of an object has been transitioned to archival storage, StorageGRID manages the process of restoring the object so it can be retrieved.	Because the mirrored copy in the destination bucket is an independent copy, applications can retrieve the object by making requests either to StorageGRID or to the S3 destination. For example, suppose you use CloudMirror replication to mirror objects to a partner organization. The partner can use its own applications to read or update objects directly from the S3 destination. Using StorageGRID is not required.

	Cloud Storage Pool	CloudMirror replication service
Can you read from the destination directly?	No. Objects moved to a Cloud Storage Pool are managed by StorageGRID. Read requests must be directed to StorageGRID (and StorageGRID will be responsible for retrieval from Cloud Storage Pool).	Yes, because the mirrored copy is an independent copy.
What happens if an object is deleted from the source?	The object is also deleted from the Cloud Storage Pool.	The delete action is not replicated. A deleted object no longer exists in the StorageGRID bucket, but it continues to exist in the destination bucket. Similarly, objects in the destination bucket can be deleted without affecting the source.
How do you access objects after a disaster (StorageGRID system not operational)?	Failed StorageGRID nodes must be recovered. During this process, copies of replicated objects might be restored using the copies in the Cloud Storage Pool.	The object copies in the CloudMirror destination are independent of StorageGRID, so they can be accessed directly before the StorageGRID nodes are recovered.

Create a Cloud Storage Pool

A Cloud Storage Pool specifies a single external Amazon S3 bucket or other S3compatible provider, or Azure Blob storage container.

When you create a Cloud Storage Pool, you specify the name and location of the external bucket or container that StorageGRID will use to store objects, the cloud provider type (Amazon S3/GCP or Azure Blob storage), and the information StorageGRID needs to access the external bucket or container.

StorageGRID validates the Cloud Storage Pool as soon as you save it, so you must ensure that the bucket or container specified in the Cloud Storage Pool exists and is reachable.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the required access permissions.
- You have reviewed the considerations for Cloud Storage Pools.
- The external bucket or container referenced by the Cloud Storage Pool already exists, and you know its name and location.
- To access the bucket or container, you have the following information for the authentication type you will choose:

S3 access key

For the external S3 bucket

- $\circ\,$ The access key ID for the account that owns the external bucket.
- The associated secret access key.

Alternatively, you can specify Anonymous for the authentication type.

C2S access portal

For Commercial Cloud Services (C2S) S3 service

You have the following:

- Complete URL that StorageGRID will use to obtain temporary credentials from the C2S access portal (CAP) server, including all the required and optional API parameters assigned to your C2S account.
- Server CA certificate issued by an appropriate Government Certificate Authority (CA).
 StorageGRID uses this certificate to verify the identity of the CAP server. The server CA certificate must use PEM encoding.
- Client certificate issued by an appropriate Government Certificate Authority (CA). StorageGRID uses this certificate to identity itself to the CAP server. The client certificate must use PEM encoding and must have been granted access to your C2S account.
- PEM-encoded private key for the client certificate.
- Passphrase for decrypting the private key for the client certificate, if it is encrypted.



If the client certificate will be encrypted, use the traditional format for the encryption. PKCS #8 encrypted format is not supported.

Azure Blob storage

For the external container

- Uniform Resource Identifier (URI) used to access the Blob Storage container.
- $\,\circ\,$ Name of the storage account and the account key. You can use the Azure portal to find these values.

Steps

1. Select ILM > Storage pools > Cloud Storage Pools.

2. Select **Create**, then enter the following information:

Field	Description
Cloud Storage Pool name	A name that briefly describes the Cloud Storage Pool and its purpose. Use a name that will be easy to identify when you configure ILM rules.

Field	Description
Provider type	 Which cloud provider you will use for this Cloud Storage Pool: Amazon S3/GCP: Select this option for an Amazon S3, Commercial Cloud Services (C2S) S3, Google Cloud Platform (GCP), or other S3-compatible provider. Azure Blob Storage
Bucket or container	The name of the external S3 bucket or Azure container. You can't change this value after the Cloud Storage Pool is saved.

3. Based on your Provider type selection, enter the Service endpoint information.

Amazon S3/GCP

a. For the protocol, select either HTTPS or HTTP.



Don't use HTTP connections for sensitive data.

b. Enter the hostname. Example:

s3-aws-region.amazonaws.com

c. Select the URL style:

Option	Description
Auto-detect	Attempt to automatically detect which URL style to use, based on the information provided. For example, if you specify an IP address, StorageGRID will use a path-style URL. Select this option only if you don't know which specific style to use.
Virtual-hosted-style	Use a virtual-hosted-style URL to access the bucket. Virtual-hosted- style URLs include the bucket name as part of the domain name. Example: https://bucket-name.s3.company.com/key-name
Path-style	Use a path-style URL to access the bucket. Path-style URLs include the bucket name at the end. Example: https://s3.company.com/bucket-name/key-name Note: The path-style URL option is not recommended and will be deprecated in a future release of StorageGRID.

d. Optionally, enter the port number, or use the default port: 443 for HTTPS or 80 for HTTP.

Azure Blob Storage

a. Using one of the following formats, enter the URI for the service endpoint.

- https://host:port
- http://host:port

Example: https://myaccount.blob.core.windows.net:443

If you don't specify a port, by default port 443 is used for HTTPS and port 80 is used for HTTP.

4. Select **Continue**. Then select the authentication type and enter the required information for the Cloud Storage Pool endpoint:

Access key

For Amazon S3/GCP provider type only

- a. For Access key ID, enter the access key ID for the account that owns the external bucket.
- b. For Secret access key, enter the secret access key.

CAP (C2S access portal)

For Commercial Cloud Services (C2S) S3 service

- a. For **Temporary credentials URL**, enter the complete URL that StorageGRID will use to obtain temporary credentials from the CAP server, including all the required and optional API parameters assigned to your C2S account.
- b. For **Server CA certificate**, select **Browse**, and upload the PEM-encoded CA certificate that StorageGRID will use to verify the CAP server.
- c. For **Client certificate**, select **Browse**, and upload the PEM-encoded certificate that StorageGRID will use to identify itself to the CAP server.
- d. For **Client private key**, select **Browse**, and upload the PEM-encoded private key for the client certificate.
- e. If the client private key is encrypted, enter the passphrase for decrypting the client private key. Otherwise, leave the **Client private key passphrase** field blank.

Azure Blob Storage

- a. For **Account name**, enter the name of the Blob storage account that owns the external service container.
- b. For Account key, enter the secret key for the Blob storage account.

Anonymous

No additional information is required.

5. Select **Continue**. Then choose the type of server verification you want to use:

Option	Description
Use root CA certificates in Storage Node OS	Use the Grid CA certificates installed on the operating system to secure connections.
Use custom CA certificate	Use a custom CA certificate. Select Browse , and upload the PEM- encoded certificate.
Do not verify certificate	The certificate used for the TLS connection is not verified.

6. Select Save.

When you save a Cloud Storage Pool, StorageGRID does the following:

• Validates that the bucket or container and the service endpoint exist and that they can be reached using the credentials that you specified.

• Writes a marker file to the bucket or container to identify it as a Cloud Storage Pool. Never remove this file, which is named x-ntap-sgws-cloud-pool-uuid.

If Cloud Storage Pool validation fails, you receive an error message that explains why validation failed. For example, an error might be reported if there is a certificate error or if the bucket or container you specified does not already exist.

7. If an error occurs, see the instructions for troubleshooting Cloud Storage Pools, resolve any issues, and then try saving the Cloud Storage Pool again.

Edit a Cloud Storage Pool

You can edit a Cloud Storage Pool to change its name, service endpoint, or other details; however, you can't change the S3 bucket or Azure container for a Cloud Storage Pool.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.
- You have reviewed the considerations for Cloud Storage Pools.

Steps

1. Select ILM > Storage pools > Cloud Storage Pools.

The Cloud Storage Pools table lists the existing Cloud Storage Pools.

- 2. Select the checkbox for the Cloud Storage Pool you want to edit.
- 3. Select Actions > Edit.
- 4. As required, change the display name, service endpoint, authentication credentials, or certificate validation method.



You can't change the provider type or the S3 bucket or Azure container for a Cloud Storage Pool.

If you previously uploaded a server or client certificate, you can select **Certificate details** to review the certificate that is currently in use.

5. Select Save.

When you save a Cloud Storage Pool, StorageGRID validates that the bucket or container and the service endpoint exist, and that they can be reached using the credentials that you specified.

If Cloud Storage Pool validation fails, an error message is displayed. For example, an error might be reported if there is a certificate error.

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

Remove a Cloud Storage Pool

You can remove a Cloud Storage Pool if it not used in an ILM rule and it does not contain object data.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the required access permissions.

If needed, use ILM to move object data

If the Cloud Storage Pool you want to remove contains object data, you must use ILM to move the data to a different location. For example, you can move the data to Storage Nodes on your grid or to a different Cloud Storage Pool.

Steps

- 1. Select ILM > Storage pools > Cloud Storage Pools.
- 2. Look at the ILM usage column in the table to determine whether you can remove the Cloud Storage Pool.

You can't remove a Cloud Storage Pool if it is being used in an ILM rule or in an erasure-coding profile.

- 3. If the Cloud Storage Pool is being used, select *cloud storage pool name* > ILM usage.
- 4. Clone each ILM rule that currently places objects in the Cloud Storage Pool you want to remove.
- 5. Determine where you want to move the existing objects managed by each rule you cloned.

You can use one or more storage pools or a different Cloud Storage Pool.

6. Edit each of the rules you cloned.

For Step 2 of the Create ILM rule wizard, select the new location from the copies at field.

- 7. Create a new ILM policy and replace each of the old rules with a cloned rule.
- 8. Activate the new policy.
- 9. Wait for ILM to remove objects from the Cloud Storage Pool and place them in the new location.

Delete Cloud Storage Pool

When the Cloud Storage Pool is empty and not used in any ILM rules, you can delete it.

Before you begin

- · You have removed any ILM rules that might have used the pool.
- You have confirmed that the S3 bucket or Azure container does not contain any objects.

An error occurs if you attempt to remove a Cloud Storage Pool if it contains objects. See Troubleshoot Cloud Storage Pools.



When you create a Cloud Storage Pool, StorageGRID writes a marker file to the bucket or container to identify it as a Cloud Storage Pool. Don't remove this file, which is named x-ntap-sgws-cloud-pool-uuid.

Steps

- 1. Select ILM > Storage pools > Cloud Storage Pools.
- 2. If the ILM usage column indicates that Cloud Storage Pool is not being used, select the checkbox.
- 3. Select Actions > Remove.

4. Select OK.

Troubleshoot Cloud Storage Pools

Use these troubleshooting steps to help resolve errors you might encounter when creating, editing, or deleting a Cloud Storage Pool.

Determine if an error has occurred

StorageGRID performs a simple health check on every Cloud Storage Pool once a minute to ensure that the Cloud Storage Pool can be accessed and that it is functioning correctly. If the health check detects an issue, a message is shown in the Last error column of the Cloud Storage Pools table on the Storage pools page.

The table shows the most recent error detected for each Cloud Storage Pool and indicates how long ago the error occurred.

In addition, a **Cloud Storage Pool connectivity error** alert is triggered if the health check detects that one or more new Cloud Storage Pool errors have occurred within the past 5 minutes. If you receive an email notification for this alert, go to the Storage pools page (select **ILM > Storage pools**), review the error messages in the Last error column, and refer to the troubleshooting guidelines below.

Check if an error has been resolved

After resolving any underlying issues, you can determine if the error has been resolved. From the Cloud Storage Pool page, select the endpoint, and select **Clear error**. A confirmation message indicates that StorageGRID has cleared the error for the Cloud Storage Pool.

If the underlying problem has been resolved, the error message is no longer displayed. However, if the underlying problem has not been fixed (or if a different error is encountered), the error message will be shown in the Last error column within a few minutes.

Error: This Cloud Storage Pool contains unexpected content

You might encounter this error when you try to create, edit, or delete a Cloud Storage Pool. This error occurs if the bucket or container includes the x-ntap-sgws-cloud-pool-uuid marker file, but that file does not have the expected UUID.

Typically, you will only see this error if you are creating a new Cloud Storage Pool and another instance of StorageGRID is already using the same Cloud Storage Pool.

Try these steps to correct the issue:

- Check to make sure that no one in your organization is also using this Cloud Storage Pool.
- Delete the x-ntap-sgws-cloud-pool-uuid file and try configuring the Cloud Storage Pool again.

Error: Could not create or update Cloud Storage Pool. Error from endpoint

You might encounter this error when you try to create or edit a Cloud Storage Pool. This error indicates that some kind of connectivity or configuration issue is preventing StorageGRID from writing to the Cloud Storage Pool.

To correct the issue, review the error message from the endpoint.

• If the error message contains Get url: EOF, check that the service endpoint used for the Cloud Storage

Pool does not use HTTP for a container or bucket that requires HTTPS.

- If the error message contains Get *url*: net/http: request canceled while waiting for connection, verify that the network configuration allows Storage Nodes to access the service endpoint used for the Cloud Storage Pool.
- For all other endpoint error messages, try one or more of the following:
 - Create an external container or bucket with the same name you entered for the Cloud Storage Pool, and try to save the new Cloud Storage Pool again.
 - Correct the container or bucket name you specified for the Cloud Storage Pool, and try to save the new Cloud Storage Pool again.

Error: Failed to parse CA certificate

You might encounter this error when you try to create or edit a Cloud Storage Pool. The error occurs if StorageGRID could not parse the certificate you entered when configuring the Cloud Storage Pool.

To correct the issue, check the CA certificate you provided for issues.

Error: A Cloud Storage Pool with this ID was not found

You might encounter this error when you try to edit or delete a Cloud Storage Pool. This error occurs if the endpoint returns a 404 response, which can mean either of the following:

- The credentials used for the Cloud Storage Pool don't have read permission for the bucket.
- The bucket used for the Cloud Storage Pool does not include the x-ntap-sgws-cloud-pool-uuid marker file.

Try one or more of these steps to correct the issue:

- Check that the user associated with the configured Access Key has the requisite permissions.
- Edit the Cloud Storage Pool with credentials that have the requisite permissions.
- If the permissions are correct, contact support.

Error: Could not check the content of the Cloud Storage Pool. Error from endpoint

You might encounter this error when you try to delete a Cloud Storage Pool. This error indicates that some kind of connectivity or configuration issue is preventing StorageGRID from reading the contents of the Cloud Storage Pool bucket.

To correct the issue, review the error message from the endpoint.

Error: Objects have already been placed in this bucket

You might encounter this error when you try to delete a Cloud Storage Pool. You can't delete a Cloud Storage Pool if it contains data that was moved there by ILM, data that was in the bucket before you configured the Cloud Storage Pool, or data that was put in the bucket by some other source after the Cloud Storage Pool was created.

Try one or more of these steps to correct the issue:

 Follow the instructions for moving objects back to StorageGRID in "Lifecycle of a Cloud Storage Pool object." • If you are certain the remaining objects were not placed in the Cloud Storage Pool by ILM, manually delete the objects from the bucket.



Never manually delete objects from a Cloud Storage Pool that might have been placed there by ILM. If you later attempt to access a manually deleted object from StorageGRID, the deleted object will not be found.

Error: Proxy encountered an external error while trying to reach the Cloud Storage Pool

You might encounter this error if you have configured a non-transparent storage proxy between Storage Nodes and the external S3 endpoint used for the Cloud Storage Pool. This error occurs if the external proxy server can't reach the Cloud Storage Pool endpoint. For example, the DNS server might not be able to resolve the hostname or there might be an external networking issue.

Try one or more of these steps to correct the issue:

- Check the settings for the Cloud Storage Pool (ILM > Storage pools).
- Check the networking configuration of the storage proxy server.

Related information

Lifecycle of a Cloud Storage Pool object

Manage erasure-coding profiles

You can view the details for an erasure-coding profile and rename a profile if needed. You can deactivate an erasure-coding profile if it is not currently used in any ILM rules.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the required access permissions.

View erasure-coding profile details

You can view the details of an erasure-coding profile to determine its status, the erasure-coding scheme used, and other information.

Steps

- 1. Select ILM > Erasure coding.
- 2. Select the profile. The detail page for the profile appears.
- 3. Optionally, view the ILM rules tab for a list of ILM rules that use the profile, and the ILM policies that use those rules.
- 4. Optionally, view the Storage Nodes tab for details about each Storage Node in the profile's storage pool, such as the site where it's located and the storage usage.

Rename an erasure-coding profile

You might want to rename an erasure-coding profile to make it more obvious what the profile does.

Steps

- 1. Select ILM > Erasure coding.
- 2. Select the profile you want to rename.
- 3. Select Rename.
- 4. Enter a unique name for the erasure-coding profile.

The erasure-coding profile name is appended to the storage pool name in the placement instruction for an ILM rule.



Erasure-coding profile names must be unique. A validation error occurs if you use the name of an existing profile, even if that profile has been deactivated.

5. Select Save.

Deactivate an erasure-coding profile

You can deactivate an erasure-coding profile if you no longer plan to use it and if the profile is not currently used in any ILM rules.



Confirm that no erasure-coded data repair operations or decommission procedures are in process. An error message is returned if you attempt to deactivate an erasure-coding profile while either of these operations are in progress.

About this task

StorageGRID prevents you from deactivating an erasure-coding profile if either of the following is true:

- The erasure-coding profile is currently used in an ILM rule.
- The erasure-coding profile is no longer used in any ILM rules, but object data and parity fragments for the profile still exist.

Steps

- 1. Select ILM > Erasure coding.
- 2. On the Active tab, review the **Status** column to confirm that the erasure-coding profile you want to deactivate is not used in any ILM rules.

You can't deactivate an erasure-coding profile if it is used in any ILM rule. In the example, the 2+1 Data Center 1 profile is used in at least one ILM rule.

	Profile name 🌍 💠	Status 💡 ≑	Storage pool 🚷 🌩	Erasure-coding scheme 🚷 🗢
	2+1 Data Center 1	Used in <u>5 rules</u>	Data Center 1	2+1
	New profile	Deactivated	Data Center 1	2+1

- 3. If the profile is used in an ILM rule, follow these steps:
 - a. Select ILM > Rules.
 - b. Select each rule and review the retention diagram to determine if the rule uses the erasure-coding profile you want to deactivate.

- c. If the ILM rule uses the erasure-coding profile you want to deactivate, determine if the rule is used in any ILM policy.
- d. Complete the additional steps in the table, based on where the erasure-coding profile is used.

Where has the profile been used?	Additional steps to perform before deactivating the profile	Refer to these additional instructions
Never used in any ILM rule	No additional steps required. Continue with this procedure.	None
In an ILM rule that has never been used in any ILM policy	 Edit or delete all affected ILM rules. If you edit the rule, remove all placements that use the erasure- coding profile. Continue with this procedure. 	Work with ILM rules and ILM policies
In an ILM rule that is currently in an active ILM policy	 Clone the policy. Remove the ILM rule that uses the erasure-coding profile. Add one or more new ILM rules to ensure objects are protected. Save, simulate, and activate the new policy. Wait for the new policy to be applied and for existing objects to be moved to new locations based on the new rules you added. Note: Depending on the number of objects and the size of your StorageGRID system, it might take weeks or even months for ILM operations to move the objects to new locations, based on the new ILM rules. While you can safely attempt to deactivate an erasure-coding profile while it is still associated with data, the deactivation operation will fail. An error message will inform you if the profile is not yet ready to be deactivated. Edit or delete the rule you removed from the policy. If 	Create an ILM policy Work with ILM rules and ILM policies
	erasure-coding profile. 7. Continue with this procedure.	

Where has the profile been used?	Additional steps to perform before deactivating the profile	Refer to these additional instructions
In an ILM rule that is	1. Edit the policy.	Create an
currently in an ILM policy	2. Remove the ILM rule that uses the erasure-coding profile.	ILM policy Work with
	3. Add one or more new ILM rules to ensure all objects are protected.	ILM rules and ILM
	4. Save the policy.	policies
	5. Edit or delete the rule you removed from the policy. If you edit the rule, remove all placements that use the erasure-coding profile.	
	6. Continue with this procedure.	

- e. Refresh the Erasure-Coding Profiles page to ensure that the profile is not used in an ILM rule.
- 4. If the profile is not used in an ILM rule, select the radio button and select **Deactivate**. The Deactivate erasure-coding profile dialog box appears.



You can select multiple profiles to deactivate at the same time, as long as each profile is not used in any rule.

5. If you are sure you want to deactivate the profile, select Deactivate.

Results

- If StorageGRID is able to deactivate the erasure-coding profile, its status is Deactivated. You can no longer select this profile for any ILM rule. You can't reactivate a deactivated profile.
- If StorageGRID is not able to deactivate the profile, an error message appears. For example, an error message appears if object data is still associated with this profile. You might need to wait several weeks before trying the deactivation process again.

Configure regions (optional and S3 only)

ILM rules can filter objects based on the regions where S3 buckets are created, allowing you to store objects from different regions in different storage locations.

If you want to use an S3 bucket region as a filter in a rule, you must first create the regions that can be used by the buckets in your system.



You can't change the region for a bucket after the bucket has been created.

Before you begin

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

About this task

When creating an S3 bucket, you can specify that the bucket be created in a specific region. Specifying a

region allows the bucket to be geographically close to its users, which can help optimize latency, minimize costs, and address regulatory requirements.

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

When configuring regions, follow these guidelines:

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



EU is not considered to be an alias for eu-west-1. If you want to use the EU or eu-west-1 region, you must use the exact name.

- You can't delete or modify a region if it's used in a rule that is assigned to any policy (active or inactive).
- If you use an invalid region as the advanced filter in an ILM rule, you can't add that rule to a policy.

An invalid region can result if you use a region as an advanced filter in an ILM rule but you later delete that region, or if you use the Grid Management API to create a rule and specify a region that you have not defined.

• If you delete a region after using it to create an S3 bucket, you will need to re-add the region if you ever want to use the Location Constraint advanced filter to find objects in that bucket.

Steps

1. Select **ILM > Regions**.

The Regions page appears, with the currently defined regions listed. **Region 1** shows the default region, us-east-1, which can't be modified or removed.

2. To add a region:

a. Select Add another region.

b. Enter the name of a region that you want to use when creating S3 buckets.

You must use this exact region name as the LocationConstraint request element when you create the corresponding S3 bucket.

To remove an unused region, select the delete icon X.

An error message appears if you attempt to remove a region that is currently used in any policy (active or inactive).

4. When you are done making changes, select **Save**.

You can now select these regions from the Advanced filters section in step 1 of the Create ILM rule wizard.

Create ILM rule

Create an ILM rule: Overview

To manage objects, you create a set of information lifecycle management (ILM) rules and organize them into an ILM policy.

Every object ingested into the system is evaluated against the active policy. When a rule in the policy matches an object's metadata, the instructions in the rule determine what actions StorageGRID takes to copy and store that object.



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

Elements of an ILM rule

An ILM rule has three elements:

- Filtering criteria: A rule's basic and advanced filters define which objects the rule applies to. If an object matches all filters, StorageGRID applies the rule and creates the object copies specified in the rule's placement instructions.
- **Placement instructions**: A rule's placement instructions define the number, type, and location of object copies. Each rule can include a sequence of placement instructions to change the number, type, and location of object copies over time. When the time period for one placement expires, the instructions in the next placement are automatically applied by the next ILM evaluation.
- **Ingest behavior**: A rule's ingest behavior allows you to choose how the objects filtered by the rule are protected as they are ingested (when an S3 or Swift client saves an object to the grid).

ILM rule filtering

When you create an ILM rule, you specify filters to identify which objects the rule applies to.

In the simplest case, a rule might not use any filters. Any rule that does not use filters applies to all objects, so it must be the last (default) rule in an ILM policy. The default rule provides storage instructions for objects that don't match the filters in another rule.

• Basic filters allow you to apply different rules to large, distinct groups of objects. These filters allow you to apply a rule to specific tenant accounts, specific S3 buckets or Swift containers, or both.

Basic filters give you a simple way to apply different rules to large numbers of objects. For example, your company's financial records might need to be stored to meet regulatory requirements, while data from the marketing department might need to be stored to facilitate daily operations. After creating separate tenant accounts for each department or after segregating data from the different departments into separate S3 buckets, you can easily create one rule that applies to all financial records and a second rule that applies to all marketing data.

- Advanced filters give you granular control. You can create filters to select objects based on the following object properties:
 - Ingest time

- Last access time
- All or part of the object name (Key)
- Location constraint (S3 only)
- Object size
- User metadata
- Object tag (S3 only)

You can filter objects on very specific criteria. For example, objects stored by a hospital's imaging department might be used frequently when they are less than 30 days old and infrequently afterwards, while objects that contain patient visit information might need to be copied to the billing department at the health network's headquarters. You can create filters that identify each type of object based on object name, size, S3 object tags, or any other relevant criteria, and then create separate rules to store each set of objects appropriately.

You can combine filters as needed in a single rule. For example, the marketing department might want to store large image files differently than their vendor records, while the Human Resources department might need to store personnel records in a specific geography and policy information centrally. In this case you can create rules that filter by tenant account to segregate the records from each department, while using filters in each rule to identify the specific type of objects that the rule applies to.

ILM rule placement instructions

Placement instructions determine where, when, and how object data is stored. An ILM rule can include one or more placement instructions. Each placement instruction applies to a single period of time.

When you create placement instructions:

- You start by specifying the reference time, which determines when the placement instructions start. The reference time might be when an object is ingested, when an object is accessed, when a versioned object becomes noncurrent, or a user-defined time.
- Next, you specify when the placement will apply, relative to the reference time. For example, a placement might start on day 0 and continue for 365 days, relative to when the object was ingested.
- Finally, you specify the type of copies (replication or erasure coding) and the location where the copies are stored. For example, you might want to store two replicated copies at two different sites.

Each rule can define multiple placements for a single time period and different placements for different time periods.

- To place objects in multiple locations during a single time period, select **Add other type or location** to add more than one line for that time period.
- To place objects in different locations in different time periods, select **Add another time period** to add the next time period. Then, specify one or more lines within the time period.

The example shows two placement instructions on the Define placements page of the Create ILM rule wizard.

Time period 1	From Day	0 0	store	for 💙	365	days		
Store objects by	replicating	~	2	copies at	Data Center 1 🗙	, Data Center 2 ×	1	×
and store object	s by erasur	e coding 🛛 🗸	using	6+3 EC schem	e at all sites 🏾 🧪	×	\bigcirc	
Add other type o	r location						1	
Add other type o	r location From Day	365 0	store	forever 🗸				

The first placement instruction ① has two lines for the first year:

- The first line creates two replicated object copies at two data center sites.
- The second line creates a 6+3 erasure-coded copy using all data center sites.

The second placement instruction (2) creates two copies after one year and keeps those copies forever.

When you define the set of placement instructions for a rule, you must ensure that at least one placement instruction begins at day 0, that there are no gaps between the time periods you have defined, and that the final placement instruction continues either forever or until you no longer require any object copies.

As each time period in the rule expires, the content placement instructions for the next time period are applied. New object copies are created and any unneeded copies are deleted.

ILM rule ingest behavior

Ingest behavior controls whether object copies are immediately placed according to the instructions in the rule, or if interim copies are made and the placement instructions are applied later. The following ingest behaviors are available for ILM rules:

- **Balanced**: StorageGRID attempts to make all copies specified in the ILM rule at ingest; if this is not possible, interim copies are made and success is returned to the client. The copies specified in the ILM rule are made when possible.
- Strict: All copies specified in the ILM rule must be made before success is returned to the client.
- **Dual commit**: StorageGRID immediately makes interim copies of the object and returns success to the client. Copies specified in the ILM rule are made when possible.

Related information

Ingest options

- · Advantages, disadvantages, and limitations of the ingest options
- How consistency and ILM rules interact to affect data protection

Example ILM rule

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

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

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

Related information

- What is a storage pool
- What is a Cloud Storage Pool

Access the Create an ILM rule wizard

ILM rules allow you to manage the placement of object data over time. To create an ILM rule, you use the Create an ILM rule wizard.



If you want to create the default ILM rule for a policy, follow the instructions for creating a default ILM rule instead.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.
- If you want to specify which tenant accounts this rule applies to, you have the Tenant accounts permission or you know the account ID for each account.
- If you want the rule to filter objects on last access time metadata, Last access time updates must be enabled by bucket for S3 or by container for Swift.
- You have configured any Cloud Storage Pools you plan to use. See Create Cloud Storage Pool.
- You are familiar with the ingest options.
- If you need to create a compliant rule for use with S3 Object Lock, you are familiar with the requirements for S3 Object Lock.
- Optionally, you have watched the video: Video: Information lifecycle management rules in StorageGRID 11.8.



About this task

When creating ILM rules:

- Consider the StorageGRID system's topology and storage configurations.
- Consider what types of object copies you want to make (replicated or erasure-coded) and the number of copies of each object that are required.
- Determine what types of object metadata are used in the applications that connect to the StorageGRID system. ILM rules filter objects based on their metadata.
- · Consider where you want object copies to be placed over time.
- Decide which ingest option to use (Balanced, Strict, or Dual commit).

Steps

- 1. Select ILM > Rules.
- 2. Select Create. Step 1 (Enter details) of the Create an ILM rule wizard appears.

Step 1 of 3: Enter details

The **Enter details** step of the Create an ILM rule wizard allows you to enter a name and description for the rule and to define filters for the rule.

Entering a description and defining filters for the rule are optional.

About this task

When evaluating an object against an ILM rule, StorageGRID compares the object metadata to the rule's filters. If the object metadata matches all filters, StorageGRID uses the rule to place the object. You can design a rule to apply to all objects, or you can specify basic filters, such as one or more tenant accounts or bucket names, or advanced filters, such as the object's size or user metadata.

Steps

- 1. Enter a unique name for the rule in the Name field.
- 2. Optionally, enter a short description for the rule in the **Description** field.

You should describe the rule's purpose or function so you can recognize the rule later.

3. Optionally, select one or more S3 or Swift tenant accounts to which this rule applies. If this rule applies to all tenants, leave this field blank.

If you don't have either the Root access permission or the Tenant accounts permission, you can't select tenants from the list. Instead, enter the tenant ID or enter multiple IDs as a comma-delimited string.

4. Optionally, specify the S3 buckets or Swift containers to which this rule applies.

If applies to all buckets is selected (default), the rule applies to all S3 buckets or Swift containers.

5. For S3 tenants, optionally select **Yes** to apply the rule only to older object versions in S3 buckets that have versioning enabled.

If you select **Yes**, "Noncurrent time" will be automatically selected for Reference time in Step 2 of the Create an ILM rule wizard.



Noncurrent time applies only to S3 objects in versioning-enabled buckets. See Operations on buckets, PutBucketVersioning and Manage objects with S3 Object Lock.

You can use this option to reduce the storage impact of versioned objects by filtering for noncurrent object versions. See Example 4: ILM rules and policy for S3 versioned objects.

6. Optionally, select Add an advanced filter to specify additional filters.

If you don't configure advanced filtering, the rule applies to all objects that match the basic filters. For more information about advanced filtering, see Use advanced filters in ILM rules and Specify multiple metadata types and values.

7. Select Continue. Step 2 (Define placements) of the Create an ILM rule wizard appears.

Use advanced filters in ILM rules

Advanced filtering allows you to create ILM rules that apply only to specific objects based on their metadata. When you set up advanced filtering for a rule, you select the type of metadata you want to match, select an operator, and specify a metadata value. When objects are evaluated, the ILM rule is applied only to those objects that have metadata matching the advanced filter.

The table shows the types of metadata you can specify in advanced filters, the operators you can use for each type of metadata, and the metadata values expected.

Metadata type	Supported operators	Metadata value
Ingest time	 is is not is before is on or before is after is on or after 	Time and date the object was ingested. Note: To avoid resource issues when activating an new ILM policy, you can use the Ingest time advanced filter in any rule that might change the location of large numbers of existing objects. Set Ingest time to be greater than or equal to the approximate time when the new policy will go into effect to ensure that existing objects aren't moved unnecessarily.

Metadata type	Supported operators	Metadata value
Key	 equals does not equal contains does not contain starts with does not start with ends with does not end with 	All or part of a unique S3 or Swift object key. For example, you might want to match objects that end with .txt or start with test-object/.
Last access time	 is is not is before is on or before is after is on or after 	Time and date the object was last retrieved (read or viewed). Note: If you plan to use last access time as an advanced filter, Last access time updates must be enabled for the S3 bucket or Swift container.
Location constraint (S3 only)	equalsdoes not equal	The region where an S3 bucket was created. Use ILM > Regions to define the regions that are shown. Note: A value of us-east-1 will match objects in buckets created in the us-east-1 region as well as objects in buckets that have no region specified. See Configure regions (optional and S3 only).
Object size	 equals does not equal less than less than or equal to greater than greater than or equal to 	The object's size. Erasure coding is best suited for objects greater than 1 MB. Don't use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.

Metadata type	Supported operators	Metadata value
User metadata	 contains ends with	Key-value pair, where User metadata name is the key and Metadata value is the value.
	 equals exists starts with does not contain does not end with does not equal does not exist does not start with 	For example, to filter on objects that have user metadata of color=blue, specify color for User metadata name, equals for the operator, and blue for Metadata value. Note: User-metadata names aren't case sensitive; user-metadata values are case sensitive.
Object tag (S3 only)	 contains ends with equals exists starts with does not contain does not end with does not equal does not exist does not start with 	 Key-value pair, where Object tag name is the key and Object tag value is the value. For example, to filter on objects that have an object tag of 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.

Specify multiple metadata types and values

When you define advanced filtering, you can specify multiple types of metadata and multiple metadata values. For example, if you want a rule to match objects between 10 MB and 100 MB in size, you would select the **Object size** metadata type and specify two metadata values.

- The first metadata value specifies objects greater than or equal to 10 MB.
- The second metadata value specifies objects less than or equal to 100 MB.

ilter grou	up 1 Objects wi	th all of followir	ng metadata will be evaluated by this	s rule:					×
Obje	ct size	~	greater than or equal to $$	10	٢	MB	~	×	
and	Object size	~	less than or equal to 🗸	100	0	0	МВ	~	×

Using multiple entries allows you to have precise control over which objects are matched. In the following example, the rule applies to objects that have Brand A or Brand B as the value of the camera_type user metadata. However, the rule only applies to those Brand B objects that are smaller than 10 MB.

User metadata	~	camera_type	equals	~	Brand A	×
Add another advanced filt	ter					
Filter group 2 Object	cts with all of follo	owing metadata will be evalu	uated by this rule:			
Filter group 2 Object	cts with all of follo	owing metadata will be evalu camera_type	uated by this rule:	~	Brand B	×

Step 2 of 3: Define placements

The **Define placements** step of the Create ILM Rule wizard allows you to define the placement instructions that determine how long objects are stored, the type of copies (replicated or erasure-coded), the storage location, and the number of copies.

About this task

An ILM rule can include one or more placement instructions. Each placement instruction applies to a single period of time. When you use more than one instruction, the time periods must be contiguous, and at least one instruction must start on day 0. The instructions can continue either forever, or until you no longer require any object copies.

Each placement instruction can have multiple lines if you want to create different types of copies or use different locations during that time period.

In this example, the ILM rule stores one replicated copy in Site 1 and one replicated copy in Site 2 for the first year. After one year, a 2+1 erasure-coded copy is made and saved at only one site.

Time period 1	From Day	0	store	for 🗸	365	days	×
Store objects by	re <mark>pl</mark> icatin	g 🗸	1	copies at	Site 1 X	/ ×	
and store objects	by replic	ating	✓ 1	copies at	Site 2	× / ×	
Add other type or	location						
Time period 2	From Day	365	© store	forever 🗸			×
Store objects by	erasure co	oding 🗸	using 2+	1 EC scheme at Sit	e 3 🧪	×	

1. For **Reference time**, select the type of time to use when calculating the start time for a placement instruction.

Option	Description
Ingest time	The time when the object was ingested.
Last access time	The time when the object was last retrieved (read or viewed). Note: To use this option, updates to Last access time must be enabled for the S3 bucket or Swift container. See Use Last access time in ILM rules.
User defined creation time	A time specified in user-defined metadata.
Noncurrent time	"Noncurrent time" is automatically selected if you selected Yes for the question, "Apply this rule to older object versions only (in S3 buckets with versioning enabled)?" in Step 1 of the Create an ILM rule wizard.



If you want to create a compliant rule, you must select **Ingest time**. See Manage objects with S3 Object Lock.

2. In the **Time period and placements** section, enter a starting time and a duration for the first time period.

For example, you might want to specify where to store objects for the first year (*From day 0 store for 365 days*). At least one instruction must start at day 0.

- 3. If you want to create replicated copies:
 - a. From the Store objects by drop-down list, select replicating.
 - b. Select the number of copies you want to make.

A warning appears if you change the number of copies to 1. An ILM rule that creates only one replicated copy for any time period puts data at risk of permanent loss. See Why you should not use single-copy replication.

To avoid the risk, do one or more of the following:

- Increase the number of copies for the time period.
- Add copies to other storage pools or to a Cloud Storage Pool.
- Select erasure coding instead of replicating.

You can safely ignore this warning if this rule already creates multiple copies for all time periods.

c. In the **copies at** field, select the storage pools you want to add.

If you specify only one storage pool, be aware that StorageGRID can store only one replicated copy of an object on any given Storage Node. If your grid includes three Storage Nodes and you select 4 as the number of copies, only three copies will be made—one copy for each Storage Node.



The **ILM placement unachievable** alert is triggered to indicate that the ILM rule could not be completely applied.

If you specify more than one storage pool, keep these rules in mind:

- The number of copies can't be greater than the number of storage pools.
- If the number of copies equals the number of storage pools, one copy of the object is stored in each storage pool.
- If the number of copies is less than the number of storage pools, one copy is stored at the ingest site, and then the system distributes the remaining copies to keep disk usage among the pools balanced, while ensuring that no site gets more than one copy of an object.
- If the storage pools overlap (contain the same Storage Nodes), all copies of the object might be saved at only one site. For this reason, don't specify the All Storage Nodes storage pool (StorageGRID 11.6 and earlier) and another storage pool.
- 4. If you want to create an erasure-coded copy:
 - a. From the Store objects by drop-down list, select erasure coding.



Erasure coding is best suited for objects greater than 1 MB. Don't use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.

- b. If you didn't add an Object size filter for a value greater than 200 KB, select **Previous** to return to Step 1. Then, select **Add an advanced filter** and set an **Object size** filter to any value greater than 200 KB.
- c. Select the storage pool you want to add and the erasure-coding scheme you want to use.

The storage location for an erasure-coded copy includes the name of the erasure-coding scheme, followed by the name of the storage pool.

- 5. Optionally:
 - a. Select Add other type or location to create additional copies at different locations.
 - b. Select Add another time period to add different time periods.



Objects are automatically deleted at the end of the final time period unless another time period ends with **forever**.

- 6. If you want to store objects in a Cloud Storage Pool:
 - a. In the Store objects by drop-down list, select replicating.
 - b. Select the copies at field, then select a Cloud Storage Pool.

When using Cloud Storage Pools, keep these rules in mind:

- You can't select more than one Cloud Storage Pool in a single placement instruction. Similarly, you can't select a Cloud Storage Pool and a storage pool in the same placement instruction.
- You can store only one copy of an object in any given Cloud Storage Pool. An error message appears if you set **Copies** to 2 or more.
- You can't store more than one object copy in any Cloud Storage Pool at the same time. An error
 message appears if multiple placements that use a Cloud Storage Pool have overlapping dates or
 if multiple lines in the same placement use a Cloud Storage Pool.

- You can store an object in a Cloud Storage Pool at the same time that object is being stored as replicated or erasure-coded copies in StorageGRID. However, you must include more than one line in the placement instruction for the time period, so you can specify the number and types of copies for each location.
- 7. In the Retention diagram, confirm your placement instructions.

In this example, the ILM rule stores one replicated copy in Site 1 and one replicated copy in Site 2 for the first year. After one year and for an additional 10 years, a 6+3 erasure-coded copy will be saved at three sites. After 11 years total, the objects will be deleted from StorageGRID.

The Rule analysis section of the Retention diagram states:

- StorageGRID site-loss protection will apply for the duration of this rule.
- Objects processed by this rule will be deleted after Day 4015.

Day 401

8. Select Continue. Step 3 (Select ingest behavior) of the Create an ILM rule wizard appears.

Use Last access time in ILM rules

You can use Last access time as the reference time in an ILM rule. For example, you might want to leave objects that have been viewed in the last three months on local Storage Nodes, while moving objects that have not been viewed as recently to an off-site location. You can also use Last access time as an advanced filter if you want an ILM rule to apply only to objects that were last accessed on a specific date.

About this task

Before using Last access time in an ILM rule, review the following considerations:

• When using Last access time as a reference time, be aware that changing the Last access time for an object does not trigger an immediate ILM evaluation. Instead, the object's placements are assessed and the object is moved as required when background ILM evaluates the object. This could take two weeks or more after the object is accessed.

Take this latency into account when creating ILM rules based on Last access time and avoid placements that use short time periods (less than one month).

• When using Last access time as an advanced filter or as a reference time, you must enable last access time updates for S3 buckets. You can use the Tenant Manager or the Tenant Management API.



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

()

Be aware that enabling last access time updates can reduce performance, especially in systems with small objects. The performance impact occurs because StorageGRID must update the objects with new timestamps every time the objects are retrieved.

The following table summarizes whether the Last access time is updated for all objects in the bucket for different types of requests.

Type of request	Whether Last access time is updated when last access time updates are disabled	Whether Last access time is updated when last access time updates are enabled
Request to retrieve an object, its access control list, or its metadata	No	Yes
Request to update an object's metadata	Yes	Yes
Request to copy an object from one bucket to another	No, for the source copyYes, for the destination copy	Yes, for the source copyYes, for the destination copy
Request to complete a multipart upload	Yes, for the assembled object	Yes, for the assembled object

Step 3 of 3: Select ingest behavior

The **Select ingest behavior** step of the Create ILM Rule wizard allows you to choose how the objects filtered by this rule are protected as they are ingested.

About this task

StorageGRID can make interim copies and queue the objects for ILM evaluation later, or it can make copies to meet the rule's placement instructions immediately.

Steps

1. Select the ingest behavior to use.

For more information, see Advantages, disadvantages, and limitations of the ingest options.

You can't use the Balanced or Strict option if the rule uses one of these placements:

- A Cloud Storage Pool at day 0
- (\mathbf{i})
- An Archive Node at day 0
- A Cloud Storage Pool or an Archive Node when the rule uses a User defined creation time as a Reference time

See Example 5: ILM rules and policy for Strict ingest behavior.

2. Select Create.

The ILM rule is created. The rule does not become active until it is added to an ILM policy and that policy is activated.

To view the details of the rule, select the rule's name on the ILM rules page.

Create a default ILM rule

Before creating an ILM policy, you must create a default rule to place any objects not matched by another rule in the policy. The default rule can't use any filters. It must apply to all tenants, all buckets, and all object versions.

Before you begin

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

About this task

The default rule is the last rule to be evaluated in an ILM policy, so it can't use any filters. The placement instructions for the default rule are applied to any objects that aren't matched by another rule in the policy.

In this example policy, the first rule applies only to objects belonging to test-tenant-1. The default rule, which is last, applies to objects belonging to all other tenant accounts.

Proposed policy name			-			
Example ILM policy						
Reason for change						
Example						
Manage rules 1. Select the rules you want to add to the policy. 2. Determine the order in which the rules will be evaluated by dragging and dropping the rows. The default rule will be automatically placed at the end of the policy and cannot be moved. Select rules						
Rule order	Rule name	Filters				
1 EC for test-tenant-1 Tenant is test-tenant-1						
Default	Default rule	-				

When you create the default rule, keep these requirements in mind:

- The default rule will automatically be placed as the last rule when you add it to a policy.
- The default rule can't use any basic or advanced filters.

- The default rule must apply to all object versions.
- The default rule should create replicated copies.



Don't use a rule that creates erasure-coded copies as the default rule for a policy. Erasurecoding rules should use an advanced filter to prevent smaller objects from being erasurecoded.

- In general, the default rule should retain objects forever.
- If you are using (or you plan to enable) the global S3 Object Lock setting, the default rule must be compliant.

Steps

- 1. Select ILM > Rules.
- 2. Select Create.

Step 1 (Enter details) of the Create ILM rule wizard appears.

- 3. Enter a unique name for the rule in the Rule name field.
- 4. Optionally, enter a short description for the rule in the **Description** field.
- 5. Leave the **Tenant accounts** field blank.

The default rule must apply to all tenant accounts.

6. Leave the Bucket name drop-down selection as applies to all buckets.

The default rule must apply to all S3 buckets and Swift containers.

- 7. Keep the default answer, **No**, for the question, "Apply this rule to older object versions only (in S3 buckets with versioning enabled)?"
- 8. Don't add advanced filters.

The default rule can't specify any filters.

9. Select Next.

Step 2 (Define placements) appears.

10. For Reference time, select any option.

If you kept the default answer, **No**, for the question, "Apply this rule to older object versions only?" Noncurrent time will not be included in the pull-down list. The default rule must apply all object versions.

- 11. Specify the placement instructions for the default rule.
 - The default rule should retain objects forever. A warning appears when you activate a new policy if the default rule does not retain objects forever. You must confirm this is the behavior you expect.
 - The default rule should create replicated copies.



Don't use a rule that creates erasure-coded copies as the default rule for a policy. Erasure-coding rules should include the **Object size (MB) greater than 200 KB** advanced filter to prevent smaller objects from being erasure-coded.

- If you are using (or you plan to enable) the global S3 Object Lock setting, the default rule must be compliant:
 - It must create at least two replicated object copies or one erasure-coded copy.
 - These copies must exist on Storage Nodes for the entire duration of each line in the placement instructions.
 - Object copies can't be saved in a Cloud Storage Pool.
 - Object copies can't 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."
- 12. Look at the Retention diagram to confirm your placement instructions.
- 13. Select Continue.

Step 3 (Select ingest behavior) appears.

14. Select the ingest option to use, and select Create.

Manage ILM policies

ILM policies: Overview

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



An ILM policy that has been incorrectly configured can result in unrecoverable data loss. Before activating an ILM policy, carefully review the ILM policy and its ILM rules, and then simulate the ILM policy. Always confirm that the ILM policy will work as intended.

Default ILM policy

When you install StorageGRID and add sites, a default ILM policy is automatically created, as follows:

- If your grid contains one site, the default policy contains a default rule that replicates two copies of each object at that site.
- If your grid contains more than one site, the default rule replicates one copy of each object at each site.

If the default policy does not meet your storage requirements, you can create your own rules and policy. See Create an ILM rule and Create an ILM policy.

One or many active ILM policies?

You can have one or more active ILM policies at a time.

One policy

If your grid will use a simple data protection scheme with few tenant-specific and bucket-specific rules, use a single active ILM policy. The ILM rules can contain filters to manage different buckets or tenants.



When you have only one policy and a tenant's requirements change, you must create a new ILM policy or clone the existing policy to apply changes, simulate, and then activate the new ILM policy. Changes to the ILM policy could result in object moves that could take many days and cause system latency.

Multiple policies

To provide different quality-of-service options to tenants, you can have more than one active policy at a time. Each policy can manage specific tenants, S3 buckets, and objects. When you apply or change one policy for a specific set of tenants or objects, the policies applied to other tenants and objects are not affected.

ILM policy tags

If you want to allow tenants to easily switch between multiple data protection policies on a per-bucket basis, use multiple ILM policies with *ILM policy tags*. You assign each ILM policy to a tag, then tenants tag a bucket to apply the policy to that bucket. You can set ILM policy tags on S3 buckets only.

For example, you might have three tags named Gold, Silver, and Bronze. You can assign an ILM policy to each tag, based on how long and where that policy stores objects. Tenants can choose which policy to use by tagging their buckets. A bucket tagged Gold is managed by the Gold policy and receives the Gold level of data protection and performance.

Default ILM policy tag

A default ILM policy tag is automatically created when you install StorageGRID. Every grid must have one active policy that is assigned to the Default tag. The default policy applies to all objects in Swift containers, and any untagged S3 buckets.



How does an ILM policy evaluate objects?

An active ILM policy controls the placement, duration, and data protection of objects.

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

- 1. If the filters for the first rule in the policy match an object, the object is ingested according to that rule's ingest behavior and stored according to that rule's placement instructions.
- 2. If the filters for the first rule don't match the object, the object is evaluated against each subsequent rule in the policy until a match is made.
- 3. If no rules match an object, the ingest behavior and placement instructions for the default rule in the policy are applied. The default rule is the last rule in a policy. The default rule must apply to all tenants, all S3 buckets or Swift containers, and all object versions and can't use any advanced filters.

Example ILM policy

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

Rule 1: Replicated copies for Tenant A

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

 Make two replicated copies of all objects not matched by Rule 1 or Rule 2 (objects not belonging to Tenant A that are 1 MB or smaller).



What are active and inactive policies?

Every StorageGRID system must have at least one active ILM policy. If you want to have more than one active ILM policy, you create ILM policy tags and assign a policy to each tag. Tenants then apply tags to S3 buckets. The default policy is applied to all objects in buckets that do not have a policy tag assigned.

When you first create an ILM policy, you select one or more ILM rules and arrange them in a specific order. After you have simulated the policy to confirm its behavior, you activate it.

When you activate one ILM policy, StorageGRID uses that policy to manage all objects, including existing objects and newly ingested objects. Existing objects might be moved to new locations when the ILM rules in the new policy are implemented.

If you activate more than one ILM policy at a time, and tenants apply policy tags to S3 buckets, the objects in each bucket are managed according to the policy assigned to the tag.

A StorageGRID system tracks the history of policies that have been activated or deactivated.

Considerations for creating an ILM policy

 Only use the system-provided policy, Baseline 2 copies policy, in test systems. For StorageGRID 11.6 and earlier, the Make 2 Copies rule in this policy uses the All Storage Nodes storage pool, which contains all sites. If your StorageGRID system has more than one site, two copies of an object might be placed on the same site.



The All Storage Nodes storage pool is automatically created during the installation of StorageGRID 11.6 and earlier. If you upgrade to a later version of StorageGRID, the All Storage Nodes pool will still exist. If you install StorageGRID 11.7 or later as a new installation, the All Storage Nodes pool is not created.

- When designing a new policy, consider all of the different types of objects that might be ingested into your grid. Make sure the policy includes rules to match and place these objects as required.
- Keep the ILM policy as simple as possible. This avoids potentially dangerous situations where object data is not protected as intended when changes are made to the StorageGRID system over time.
- Make sure that the rules in the policy are in the correct order. When the policy is activated, new and existing objects are evaluated by the rules in the order listed, starting at the top. For example, if the first rule in a policy matches an object, that object will not be evaluated by any other rule.
- The last rule in every ILM policy is the default ILM rule, which can't use any filters. If an object has not been matched by another rule, the default rule controls where that object is placed and for how long it is retained.
- Before activating a new policy, review any changes that the policy is making to the placement of existing objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

Create ILM policies

Create one or more ILM policies to meet your quality-of-service requirements.

Having one active ILM policy allows you to apply the same ILM rules to all tenants and buckets.

Having multiple active ILM policies allows you to apply the appropriate ILM rules to specific tenants and buckets to fulfill multiple quality-of-service requirements.

Create an ILM policy

About this task

Before creating your own policy, verify that the default ILM policy does not meet your storage requirements.



Only use the system-provided policies, 2 copies Policy (for one-site grids) or 1 Copy per Site (for multi-site grids), in test systems. For StorageGRID 11.6 and earlier, the default rule in this policy uses the All Storage Nodes storage pool, which contains all sites. If your StorageGRID system has more than one site, two copies of an object might be placed on the same site.



If the global S3 Object Lock setting has been enabled, you must ensure that the ILM policy is compliant with the requirements of buckets that have S3 Object Lock enabled. In this section, follow the instructions that mention having S3 Object Lock enabled.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the required access permissions.
- You have created ILM rules based on whether S3 Object Lock is enabled.


 Optionally, you have watched the video: Video: Information lifecycle management policies in StorageGRID 11.8

Information lifecycle management policies in StorageGRID 11.8 Genng Same	■ NetApp
And the second sec	

See also Create an ILM policy: Overview.

Steps

1. Select ILM > Policies.

If the global S3 Object Lock setting is enabled, the ILM policies page indicates which ILM rules are compliant.

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

Create new policy

a. Select Create policy.

Clone existing policy

a. Select the checkbox for the policy you want to start with, then select Clone.

Edit existing policy

- a. If a policy is inactive, you can edit it. Select the checkbox for the inactive policy you want to start with, then select **Edit**.
- 3. In the **Policy name** field, enter a unique name for the policy.

- 4. Optionally, in the **Reason for change** field, enter the reason you are creating a new policy.
- 5. To add rules to the policy, select **Select rules**. Select a rule name to view the settings for that rule.

If you are cloning a policy:

- The rules used by the policy you are cloning are selected.
- If the policy you are cloning used any rules with no filters that were not the default rule, you are prompted to remove all but one of those rules.
- If the default rule used a filter, you are prompted to select a new default rule.
- If the default rule was not the last rule, you can move the rule to the end of the new policy.

S3 Object Lock not enabled

a. Select one default rule for the policy. To create a new default rule, select ILM rules page.

The default rule applies to any objects that don't match another rule in the policy. The default rule can't use any filters and is always evaluated last.



Don't use the Make 2 Copies rule as the default rule for a policy. The Make 2 Copies rule uses a single storage pool, All Storage Nodes, which contains all sites. If your StorageGRID system has more than one site, two copies of an object might be placed on the same site.

S3 Object Lock enabled

a. Select one default rule for the policy. To create a new default rule, select **ILM rules page**.

The list of rules contains only the rules that are compliant and don't use any filters.



Don't use the Make 2 Copies rule as the default rule for a policy. The Make 2 Copies rule uses a single storage pool, All Storage Nodes, which contains all sites. If you use this rule, multiple copies of an object might be placed on the same site.

b. If you need a different "default" rule for objects in non-compliant S3 buckets, select **Include a rule without filters for non-compliant S3 buckets**, and select one non-compliant rule that does not use a filter.

For example, you might want to use a Cloud Storage Pool to store objects in buckets that don't have S3 Object Lock enabled.



You can only select one non-compliant rule that does not use a filter.

See also Example 7: Compliant ILM policy for S3 Object Lock.

- 6. When you are done selecting the default rule, select **Continue**.
- 7. For the Other rules step, select any other rules you want to add to the policy. These rules use at least one filter (tenant account, bucket name, advanced filter, or the Noncurrent reference time). Then select **Select**.

The Create a policy window now lists the rules you selected. The default rule is at the end, with the other

rules above it.

If S3 Object Lock is enabled and you also selected a non-compliant "default" rule, that rule is added as the second-to-last rule in the policy.



A warning appears if any rule does not retain objects forever. When you activate this policy, you must confirm that you want StorageGRID to delete objects when the placement instructions for the default rule elapse (unless a bucket lifecycle keeps the objects for a longer time period).

8. Drag the rows for the non-default rules to determine the order in which these rules will be evaluated.

You can't move the default rule. If S3 Object Lock is enabled, you also can't move the non-compliant "default" rule if one was selected.



You must confirm that the ILM rules are in the correct order. When the policy is activated, new and existing objects are evaluated by the rules in the order listed, starting at the top.

- 9. As required, select Select rules to add or remove rules.
- 10. When you are done, select **Save**.
- 11. Repeat these steps to create additional ILM policies.
- 12. Simulate an ILM policy. You should always simulate a policy before activating it to ensure it works as expected.

Simulate a policy

Simulate a policy on test objects before activating the policy and applying it to your production data.

Before you begin

• You know the S3 bucket/object-key or the Swift container/object-name for each object you want to test.

Steps

- 1. Using an S3 or Swift client or the S3 Console, ingest the objects required to test each rule.
- 2. On the ILM policies page, select the checkbox for the policy, then select Simulate.
- 3. In the Object field, enter the S3 bucket/object-key or the Swift container/object-name for a test object. For example, bucket-01/filename.png.
- 4. If S3 versioning is enabled, optionally enter a version ID for the object in the Version ID field.
- 5. Select Simulate.
- 6. In the Simulation results section, confirm that each object was matched by the correct rule.
- 7. To determine which storage pool or erasure-coding profile is in effect, select the name of the matched rule to go to the rule details page.



Review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

Results

Any edits to the policy's rules will be reflected in the Simulation results and show the new match and previous

match. The Simulate policy window retains the objects you tested until you select either **Clear all** or the remove icon \times for each object in the Simulation results list.

Related information

Example ILM policy simulations

Activate a policy

When you activate a single new ILM policy, existing objects and newly ingested objects are managed by that policy. When you activate multiple policies, ILM policy tags assigned to buckets determine the objects to be managed.

Before you activate a new policy:

- 1. Simulate the policy to confirm that it behaves as you expect.
- 2. Review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.



Errors in an ILM policy can cause unrecoverable data loss.

About this task

When you activate an ILM policy, the system distributes the new policy to all nodes. However, the new active policy might not actually take effect until all grid nodes are available to receive the new policy. In some cases, the system waits to implement a new active policy to ensure that grid objects aren't accidentally removed. Specifically:

- If you make policy changes that **increase data redundancy or durability**, those changes are implemented immediately. For example, if you activate a new policy that includes a three-copies rule instead of a two-copies rule, that policy will be implemented right away because it increases data redundancy.
- If you make policy changes that **could decrease data redundancy or durability**, those changes will not be implemented until all grid nodes are available. For example, if you activate a new policy that uses a two-copies rule instead of a three-copies rule, the new policy will appear in the Active policy tab but it will not take effect until all nodes are online and available.

Steps

Follow the steps for activating one policy or multiple policies:

Activate one policy

Follow these steps if you will have only one active policy. If you already have one or more active policies and you are activating additional policies, follow the steps for activating multiple policies.

1. When you are ready to activate a policy, select **ILM > Policies**.

Alternatively, you can activate a single policy from the **ILM > Policy tags** page.

- 2. On the Policies tab, select the checkbox for the policy you want to activate, then select Activate.
- 3. Follow the appropriate step:
 - If a warning message prompts you to confirm that you want to activate the policy, select OK.
 - If a warning message containing details about the policy appears:
 - a. Review the details to ensure the policy would manage data as expected.
 - b. If the default rule stores objects for a limited number of days, review the retention diagram and then type in that number of days into the text box.
 - c. If the default rule stores objects forever, but one or more other rules has limited retention, type **yes** in the text box.
 - d. Select Activate policy.

Activate multiple policies

To activate multiple policies, you must create tags and assign a policy to each tag.



When multiple tags are in use, if tenants frequently reassign policy tags to buckets, grid performance might be impacted. If you have untrusted tenants, consider using only the Default tag.

- 1. Select **ILM > Policy tags**.
- 2. Select Create.
- 3. In the Create policy tag dialog box, type a tag name and, optionally, a description for the tag.



Tag names and descriptions are visible to tenants. Choose values that will help tenants make an informed decision when selecting policy tags to assign to their buckets. For example, if the assigned policy will delete objects after a period of time, you could communicate that in the description. Do not include sensitive information in these fields.

4. Select Create tag.

- 5. In the ILM policy tags table, use the pull-down to select a policy to assign to the tag.
- 6. If warnings appear in the Policy limitations column, select View policy details to review the policy.
- 7. Ensure each policy would manage data as expected.
- 8. Select Activate assigned policies. Or, select Clear changes to remove the policy assignment.
- 9. In the Activate policies with new tags dialog box, review the descriptions of how each tag, policy, and rule will manage objects. Make changes as needed to ensure the policies will manage objects as expected.
- 10. When you are sure you want to activate the policies, type yes in the text box, then select Activate

Related information

Example 6: Changing an ILM policy

Example ILM policy simulations

The examples of ILM policy simulations provide guidelines for structuring and modifying simulations for your environment.

Example 1: Verify rules when simulating an ILM policy

This example describes how to verify rules when simulating a policy.

In this example, the **Example ILM policy** is being simulated against the ingested objects in two buckets. The policy includes three rules, as follows:

- The first rule, Two copies, two years for bucket-a, applies only to objects in bucket-a.
- The second rule, **EC objects > 1 MB**, applies to all buckets but filters on objects greater than 1 MB.
- The third rule, **Two copies, two data centers**, is the default rule. It does not include any filters and does not use the Noncurrent reference time.

After simulating the policy, confirm that each object was matched by the correct rule.

: [Simulation results Use this table to confirm the results of applying this policy to the selected objects. Clear all						
	Object 🗢	Version ID 🖨	Rule matched 🚷 🗢	Previous match 💡 ≑	Actions		
	bucket-a/bucket-a object.pdf	_	Two copies, two years for bucket-a	_	×		
	bucket-b/test object greater than 1 MB.pdf	_	EC objects > 1 MB	_	×		
	bucket-b/test object less than 1 MB.pdf	_	Two copies, two data centers	_	×		

In this example:

- bucket-a/bucket-a object.pdf correctly matched the first rule, which filters on objects in bucketa.
- bucket-b/test object greater than 1 MB.pdf is in bucket-b, so it did not match the first rule. Instead, it was correctly matched by the second rule, which filters on objects greater than 1 MB.
- bucket-b/test object less than 1 MB.pdf did not match the filters in the first two rules, so it will be placed by the default rule, which includes no filters.

Example 2: Reorder rules when simulating an ILM policy

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

In this example, the **Demo** policy is being simulated. This policy, which is intended to find objects that have series=x-men user metadata, includes three rules, as follows:

- The first rule, PNGs, filters for key names that end in .png.
- The second rule, **X-men**, applies only to objects for Tenant A and filters for series=x-men user metadata.
- The last rule, **Two copies two data centers**, is the default rule, which matches any objects that don't match the first two rules.

Steps

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

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

Simulation results Use this table to confirm the results of applying this policy to the selected objects.						
Clear all						
Object 🗢	Version ID ≑	Rule matched 🕜 ≑	Previous match 存 ≑	Actions		
photos/Havok.png	_	PNGs	_	×		

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

- 3. To resolve the issue, reorder the rules.
 - a. Select Finish to close the Simulate ILM Policy window.
 - b. Select Edit to edit the policy.
 - c. Drag the X-men rule to the top of the list.
 - d. Select Save.
- 4. Select Simulate.

The objects you previously tested are re-evaluated against the updated policy, and the new simulation results are shown. In the example, the Rule matched column shows that the Havok.png object now matches the X-men metadata rule, as expected. The Previous match column shows that the PNGs rule matched the object in the previous simulation.

Simulation results Use this table to confirm the results of applying this policy to the selected objects. Clear all					
Object 🗢	Version ID 🜲	Rule matched 👔 💠	Previous match 🧿 ≑	Actions	
photos/Havok.png	_	X-men	PNGs	×	

Example 3: Correct a rule when simulating an ILM policy

This example shows how to simulate a policy, correct a rule in the policy, and continue the simulation.

In this example, the **Demo** policy is being simulated. This policy is intended to find objects that have series=x-men user metadata. However, unexpected results occurred when simulating this policy against the Beast.jpg object. Instead of matching the X-men metadata rule, the object matched the default rule, Two copies two data centers.

S	Simulation results se this table to confirm the results of applying this policy to the Clear all	selected objects.			
	Object 🗢	Version ID 🜲	Rule matched 👔 💠	Previous match 👔 ≑	Actions
	photos/Beast.jpg		Two copies two data centers	_	×

When a test object is not matched by the expected rule in the policy, you must examine each rule in the policy and correct any errors.

Steps

- 1. Select **Finish** to close the Simulate policy dialog. On the details page for the policy, select **Retention diagram**. Then select **Expand all** or **View details** for each rule as needed.
- 2. Review the rule's tenant account, reference time, and filtering criteria.

As an example, suppose the metadata for the X-men rule was entered as "x-men01" instead of "x-men."

- 3. To resolve the error, correct the rule as follows:
 - If the rule is part of the policy, you can either clone the rule or remove the rule from the policy and then edit it.
 - If the rule is part of the active policy, you must clone the rule. You can't edit or remove a rule from the active policy.
- 4. Perform the simulation again.

In this example, the corrected X-men rule now matches the Beast.jpg object based on the series=x-men user metadata, as expected.

Simulation results Use this table to confirm the results of applying this policy to the selected objects.						
Clear all						
Object 🗢	Version ID 🜲	Rule matched 🥝 🗢	Previous match 🥝 🗢	Actions		
photos/Beast.jpg	_	X-men	-	×		

Manage ILM policy tags

You can view ILM policy tag details, edit a tag, or remove a tag.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the required access permissions.

View ILM policy tag details

To view the details for a tag:

- 1. Select **ILM > Policy tags**.
- 2. Select the name of the policy from the table. The details page for the tag appears.
- 3. On the details page, view the previous history of assigned policies.
- 4. View a policy by selecting it.

Edit ILM policy tag



Tag names and descriptions are visible to tenants. Choose values that will help tenants make an informed decision when selecting policy tags to assign to their buckets. For example, if the assigned policy will delete objects after a period of time, you could communicate that in the description. Do not include sensitive information in these fields.

To edit the description for an existing tag:

- 1. Select ILM > Policy tags.
- 2. Select the checkbox for the tag, then select Edit.

Alternatively, select the name of the tag. The details page for the tag appears, and you can select **Edit** on that page.

- 3. Change the tag description as needed
- 4. Select Save.

Remove ILM policy tag

When you remove a policy tag, any buckets that are assigned that tag will have the Default policy applied.

To remove a tag:

1. Select ILM > Policy tags.

2. Select the checkbox for the tag, then select **Remove**. A confirmation dialog box appears.

Alternatively, select the name of the tag. The details page for the tag appears, and you can select **Remove** on that page.

3. Select **Yes** to delete the tag.

Verify an ILM policy with object metadata lookup

After you have activated an ILM policy, 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 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. If the S3 bucket is versioned and you want to look up a specific version of an S3 object using the bucket and object key, you have the version ID.
 - **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.
- 4. Optionally, enter a version ID for the object (S3 only).
- 5. Select Look Up.

The object metadata lookup results appear. This page lists the following types of information:

- System metadata, including:
 - object ID (UUID)
 - object name
 - name of the container
 - result type (object, delete marker, S3 bucket or Swift container)
 - tenant account name or ID
 - logical size of the object
 - date and time the object was first created
 - date and time the object was last modified

- · Any custom user metadata key-value pairs associated with the object.
- For S3 objects, any object tag key-value pairs associated with the object.
- For replicated object copies, the current storage location of each copy.
- For erasure-coded object copies, the current storage location of each fragment.
- For object copies in a Cloud Storage Pool, the location of the object, including the name of the external bucket and the object's unique identifier.
- For segmented objects and multipart objects, a list of object segments including segment identifiers and data sizes. For objects with more than 100 segments, only the first 100 segments are shown.
- All object metadata in the unprocessed, internal storage format. This raw metadata includes internal system metadata that is not guaranteed to persist from release to release.

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



The following screenshot is an example. Your results will vary depending on your StorageGRID version.

System Metadata

Object ID	A12E96FF-B13F-4905-9E9E-45373F6E7DA8
Name	testobject
Container	source
Account	t-1582139188
Size	5.24 MB
Creation Time	2020-02-19 12:15:59 PST
Modified Time	2020-02-19 12:15:59 PST

Replicated Copies

Node	Disk Path
99-97	/var/local/rangedb/2/p/06/0B/00nM8H\$iTFbnQQ}ICV2E
99-99	/var/local/rangedb/1/p/12/0A/00nM8H\${TFboW28{CXG%

Raw Metadata

```
{
    "TYPE": "CTNT",
    "CHND": "A12E96FF-B13F-4905-9E9E-45373F6E7DA8",
    "NAME": "testobject",
    "CBID": "0x8823DE7EC7C10416",
    "PHND": "FEA0AE51-534A-11EA-9FCD-31FF00C36D56",
    "PPTH": "source",
    "META": {
        "BASE": {
            "PAWS": "2",
        }
        }
    }
}
```

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



If the Audit option is enabled, you can also monitor the audit log for the ORLM Object Rules Met message. The ORLM audit message can provide you with more information about the status of the ILM evaluation process, but it can't give you information about the correctness of the object data's placement or the completeness of the ILM policy. You must evaluate this yourself. For details, see Review audit logs.

Related information

- Use S3 REST API
- Use Swift REST API

Work with ILM policies and ILM rules

As your storage requirements change, you might need to put additional policies in place or modify the ILM rules associated with a policy. You can view ILM metrics to determine system performance.

Before you begin

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

View ILM policies

To view active and inactive ILM policies and policy activation history:

- 1. Select **ILM > Policies**.
- 2. Select **Policies** to view a list of active and inactive policies. The table lists the name of each policy, the tags the policy is assigned to, and whether the policy is active or inactive.
- 3. Select Activation history to view a list of activation start and end dates for policies.
- 4. Select a policy name to view the details for the policy.



If you view the details for a policy whose status is Edited or Deleted, a message appears explaining that you are viewing the version of the policy that was active for the specified time span and has since been edited or deleted.

Edit an ILM policy

You can only edit an inactive policy. If you want to edit an active policy, deactivate it or create a clone and edit the clone.

To edit a policy:

- 1. Select **ILM > Policies**.
- 2. Select the checkbox for the policy you want to edit, then select Edit.
- 3. Edit the policy by following the instructions in Create ILM policies.
- 4. Simulate the policy before you re-activate it.



An ILM policy that has been incorrectly configured can result in unrecoverable data loss. Before activating an ILM policy, carefully review the ILM policy and its ILM rules, and then simulate the ILM policy. Always confirm that the ILM policy will work as intended.

Clone an ILM policy

To clone an ILM policy:

- 1. Select ILM > Policies.
- 2. Select the checkbox for the policy you want to clone, then select Clone.
- 3. Create a new policy starting with the policy you've cloned by following the instructions in Create ILM policies.



An ILM policy that has been incorrectly configured can result in unrecoverable data loss. Before activating an ILM policy, carefully review the ILM policy and its ILM rules, and then simulate the ILM policy. Always confirm that the ILM policy will work as intended.

Remove an ILM policy

You can only remove an ILM policy if it is inactive. To remove a policy:

- 1. Select **ILM > Policies**.
- 2. Select the checkbox for the inactive policy you want to remove.
- 3. Select Remove.

View ILM rule details

To view the details for an ILM rule, including the retention diagram and placement instructions for the rule:

- 1. Select ILM > Rules.
- 2. Select the name of the rule whose details you want to view. Example:

2 copies 2 c	lata centers	
Compliant:	No	
Ingest behavior:	Strict	
Reference time:	Noncurrent time	
Clone Edit	Remove	
Rule detail	Used in policies	
Retention diagram	Time period Storage pool	Replicated copy Series Erasure-coded (EC) copy
Rule analysis:	Objects processed by this rule will not be deleted by ILM.	
Reference time: Nonc	urrent time Ingest behavior: Strict Day 0	
Day 0 - forever	2 replicated copies - Data Center 1	
	EC 2+1 - Data Center 1	
Duration		Forever

Additionally, you can use the details page to clone, edit, or remove a rule. You can't edit or remove a rule if it's used in any policy.

Clone an ILM rule

You can clone an existing rule if you want to create a new rule that uses some of the settings of the existing rule. If you need to edit a rule that's used in any policy, you clone the rule instead and make changes to the clone. After you make changes to the clone, you can remove the original rule from the policy and replace it with the modified version as required.



You can't clone an ILM rule if it was created using StorageGRID version 10.2 or earlier.

Steps

- 1. Select ILM > Rules.
- 2. Select the checkbox for the rule you want to clone, then select **Clone**. Alternatively, select the rule name, then select **Clone** from the rule details page.
- 3. Update the cloned rule by following the steps for editing an ILM rule and using advanced filters in ILM rules.

When cloning an ILM rule, you must enter a new name.

Edit an ILM rule

You might need to edit an ILM rule to change a filter or placement instruction.

You can't edit a rule if it is used in any ILM policy. Instead, you can clone the rule and make any required

changes to the cloned copy.



An ILM policy that has been incorrectly configured can result in unrecoverable data loss. Before activating an ILM policy, carefully review the ILM policy and its ILM rules, and then simulate the ILM policy. Always confirm that the ILM policy will work as intended.

Steps

- 1. Select ILM > Rules.
- 2. Confirm that the rule you want to edit is not used in any ILM policy.
- 3. If the rule you want to edit is not in use, select the checkbox for the rule and select **Actions** > **Edit**. Alternatively, select the name of the rule, then select **Edit** on the rule details page.
- 4. Complete the steps of the Edit ILM rule wizard. As necessary, follow the steps for creating an ILM rule and using advanced filters in ILM rules.

When editing an ILM rule, you can't change its name.

Remove an ILM rule

To keep the list of current ILM rules manageable, remove any ILM rules that you aren't likely to use.

Steps

To remove an ILM rule that is currently used in an active policy:

- 1. Clone the policy.
- 2. Remove the ILM rule from the policy clone.
- 3. Save, simulate, and activate the new policy to make sure objects are protected as expected.
- 4. Go to the steps for removing an ILM rule that is currently used in an inactive policy.

To remove an ILM rule that is currently used in an inactive policy:

- 1. Select the inactive policy.
- 2. Remove the ILM rule from the policy or remove the policy.
- 3. Go to the steps for removing an ILM rule that is not currently used.

To remove an ILM rule that is not currently used:

- 1. Select ILM > Rules.
- 2. Confirm that the rule you want to remove is not used in any policy.
- 3. If the rule you want to remove is not in use, select the rule and select **Actions** > **Remove**. You can select multiple rules and remove all of them at the same time.
- 4. Select **Yes** to confirm that you want to remove the ILM rule.

View ILM metrics

You can view metrics for ILM, such as the number of objects in the queue and the evaluation rate. You can monitor these metrics to determine system performance. A large queue or evaluation rate might indicate that the system is not able to keep up with the ingest rate, the load from the client applications is excessive, or that some abnormal condition exists.

Steps

1. Select **Dashboard > ILM**.



Because the dashboard can be customized, the ILM tab might not be available.

2. Monitor the metrics on the ILM tab.

You can select the question mark ? to see a description of the items on the ILM tab.

	Month 💙	- - ·	ILM evaluation rate (objects/second)	🗂 Month 💙	L
2			1.2		
9			0.9		
6 No dat	а		0.6 No data		
3			0:3		
0			0		
Nov27		Nov 28	Nov27	r	lov 28
To be scanned 🌘 To be deleted			Scans		
.M information 🥥					

Use S3 Object Lock

Manage objects with S3 Object Lock

As a grid administrator, you can enable S3 Object Lock for your StorageGRID system and implement a compliant ILM policy to help ensure that objects in specific S3 buckets aren't deleted or overwritten for a specified amount of time.

What is S3 Object Lock?

The StorageGRID S3 Object Lock feature is an object-protection solution that is equivalent to S3 Object Lock in Amazon Simple Storage Service (Amazon S3).

As shown in the figure, when the global S3 Object Lock setting is enabled for a StorageGRID system, an S3 tenant account can create buckets with or without S3 Object Lock enabled. If a bucket has S3 Object Lock enabled, bucket versioning is required and is enabled automatically.

If a bucket has S3 Object Lock enabled, S3 client applications can optionally specify retention settings for any object version saved to that bucket.

In addition, a bucket that has S3 Object Lock enabled can optionally have a default retention mode and retention period. The default settings apply only to objects that are added to the bucket without their own



StorageGRID with S3 Object Lock setting enabled

Retention modes

The StorageGRID S3 Object Lock feature supports two retention modes to apply different levels of protection to objects. These modes are equivalent to the Amazon S3 retention modes.

- · In compliance mode:
 - The object can't be deleted until its retain-until-date is reached.
 - The object's retain-until-date can be increased, but it can't be decreased.
 - The object's retain-until-date can't be removed until that date is reached.
- In governance mode:
 - Users with special permission can use a bypass header in requests to modify certain retention settings.
 - These users can delete an object version before its retain-until-date is reached.
 - These users can increase, decrease, or remove an object's retain-until-date.

Retention settings for object versions

If a bucket is created with S3 Object Lock enabled, users can use the S3 client application to optionally specify the following retention settings for each object that is added to the bucket:

- Retention mode: Either compliance or governance.
- **Retain-until-date**: If an object version's retain-until-date is in the future, the object can be retrieved, but it can't be deleted.
- Legal hold: Applying a legal hold to an object version immediately locks that object. For example, you might need to put a legal hold on an object that is related to an investigation or legal dispute. A legal hold has no expiration date, but remains in place until it is explicitly removed. Legal holds are independent of the retain-until-date.



If an object is under a legal hold, no one can delete the object, regardless of its retention mode.

For details on the object settings, see Use S3 REST API to configure S3 Object Lock.

Default retention setting for buckets

If a bucket is created with S3 Object Lock enabled, users can optionally specify the following default settings for the bucket:

- Default retention mode: Either compliance or governance.
- **Default retention period**: How long new object versions added to this bucket should be retained, starting from the day they are added.

The default bucket settings apply only to new objects that don't have their own retention settings. Existing bucket objects aren't affected when you add or change these default settings.

See Create an S3 bucket and Update S3 Object Lock default retention.

Comparing S3 Object Lock to legacy Compliance

The S3 Object Lock replaces the Compliance feature that was available in earlier StorageGRID versions. Because the S3 Object Lock feature conforms to Amazon S3 requirements, it deprecates the proprietary StorageGRID Compliance feature, which is now referred to as "legacy Compliance."



The global Compliance setting is deprecated. If you enabled this setting using a previous version of StorageGRID, the S3 Object Lock setting is enabled automatically. You can continue to use StorageGRID to manage the settings of existing compliant buckets; however, you can't create new compliant buckets. For details, see NetApp Knowledge Base: How to manage legacy Compliant buckets in StorageGRID 11.5.

If you used the legacy Compliance feature in a previous version of StorageGRID, refer to the following table to learn how it compares to the S3 Object Lock feature in StorageGRID.

	S3 Object Lock	Compliance (legacy)
How is the feature enabled globally?	From the Grid Manager, select CONFIGURATION > System > S3 Object Lock.	No longer supported.
How is the feature enabled for a bucket?	Users must enable S3 Object Lock when creating a new bucket using the Tenant Manager, the Tenant Management API, or the S3 REST API.	No longer supported.
Is bucket versioning supported?	Yes. Bucket versioning is required and is enabled automatically when S3 Object Lock is enabled for the bucket.	No.

	S3 Object Lock	Compliance (legacy)
How is object retention set?	Users can set a retain-until-date for each object version, or they can set a default retention period for each bucket.	Users must set a retention period for the entire bucket. The retention period applies to all objects in the bucket.
Can the retention period be changed?	 In compliance mode, the retain- until-date for an object version can be increased but never decreased. In governance mode, users with special permissions can decrease or even remove an object's retention settings. 	A bucket's retention period can be increased but never decreased.
Where is legal hold controlled?	Users can place a legal hold or lift a legal hold for any object version in the bucket.	A legal hold is placed on the bucket and affects all objects in the bucket.
When can objects be deleted?	 In compliance mode, an object version can be deleted after the retain-until-date is reached, assuming the object is not under legal hold. In governance mode, users with special permissions can delete an object before its retain-until-date is reached, assuming the object is not under legal hold. 	An object can be deleted after the retention period expires, assuming the bucket is not under legal hold. Objects can be deleted automatically or manually.
Is bucket lifecycle configuration supported?	Yes	No

Workflow for S3 Object Lock

As a grid administrator, you must coordinate closely with tenant users to ensure that the objects are protected in a manner that satisfies their retention requirements.

The workflow diagram shows the high-level steps for using S3 Object Lock. These steps are performed by the grid administrator and by tenant users.



Grid administrator tasks

As the workflow diagram shows, a grid administrator must perform two high-level tasks before S3 tenant users can use S3 Object Lock:

- 1. Create at least one compliant ILM rule and make that rule the default rule in an active ILM policy.
- 2. Enable the global S3 Object Lock setting for the entire StorageGRID system.

Tenant user tasks

After the global S3 Object Lock setting has been enabled, tenants can perform these tasks:

- 1. Create buckets that have S3 Object Lock enabled.
- 2. Optionally, specify default retention settings for the bucket. Any default bucket settings are applied only to new objects that don't have their own retention settings.
- 3. Add objects to those buckets and optionally specify object-level retention periods and legal hold settings.
- 4. As required, update default retention for the bucket or update the retention period or the legal hold setting for an individual object.

Requirements for S3 Object Lock

You must review the requirements for enabling the global S3 Object Lock setting, the requirements for creating compliant ILM rules and ILM policies, and the restrictions StorageGRID places on buckets and objects that use S3 Object Lock.

Requirements for using the global S3 Object Lock setting

- You must enable the global S3 Object Lock setting using the Grid Manager or the Grid Management API before any S3 tenant can create a bucket with S3 Object Lock enabled.
- Enabling the global S3 Object Lock setting allows all S3 tenant accounts to create buckets with S3 Object Lock enabled.
- After you enable the global S3 Object Lock setting, you can't disable the setting.
- You can't enable the global S3 Object Lock unless the default rule in all active ILM policies is *compliant* (that is, the default rule must comply with the requirements of buckets with S3 Object Lock enabled).
- When the global S3 Object Lock setting is enabled, you can't create a new ILM policy or activate an existing ILM policy unless the default rule in the policy is compliant. After the global S3 Object Lock setting has been enabled, the ILM rules and ILM policies pages indicate which ILM rules are compliant.

Requirements for compliant ILM rules

If you want to enable the global S3 Object Lock setting, you must ensure that the default rule in all active ILM policies is compliant. A compliant rule satisfies the requirements of both buckets with S3 Object Lock enabled and any existing buckets that have legacy Compliance enabled:

- It must create at least two replicated object copies or one erasure-coded copy.
- These copies must exist on Storage Nodes for the entire duration of each line in the placement instructions.
- Object copies can't be saved in a Cloud Storage Pool.
- Object copies can't 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."

Requirements for ILM policies

When the global S3 Object Lock setting is enabled, active and inactive ILM policies can include both compliant and non-compliant rules.

- The default rule in an active or inactive ILM policy must be compliant.
- Non-compliant rules only apply to objects in buckets that don't have S3 Object Lock enabled or that don't have the legacy Compliance feature enabled.
- Compliant rules can apply to objects in any bucket; S3 Object Lock or legacy Compliance does not need to be enabled for the bucket.

A compliant ILM policy might include these three rules:

- 1. A compliant rule that creates erasure-coded copies of the objects in a specific bucket with S3 Object Lock enabled. The EC copies are stored on Storage Nodes from day 0 to forever.
- 2. A non-compliant rule that creates two replicated object copies on Storage Nodes for a year and then

moves one object copy to Archive Nodes and stores that copy forever. This rule only applies to buckets that don't have S3 Object Lock or legacy Compliance enabled because it stores only one object copy forever and it uses Archive Nodes.

3. A default, compliant rule that creates two replicated object copies on Storage Nodes from day 0 to forever. This rule applies to any object in any bucket that was not filtered out by the first two rules.

Requirements for buckets with S3 Object Lock enabled

- If the global S3 Object Lock setting is enabled for the StorageGRID system, you can use the Tenant Manager, the Tenant Management API, or the S3 REST API to create buckets with S3 Object Lock enabled.
- If you plan to use S3 Object Lock, you must enable S3 Object Lock when you create the bucket. You can't enable S3 Object Lock for an existing bucket.
- When S3 Object Lock is enabled for a bucket, StorageGRID automatically enables versioning for that bucket. You can't disable S3 Object Lock or suspend versioning for the bucket.
- Optionally, you can specify a default retention mode and retention period for each bucket using the Tenant Manager, the Tenant Management API, or the S3 REST API. The bucket's default retention settings apply only to new objects added to the bucket that don't have their own retention settings. You can override these default settings by specifying a retention mode and retain-until-date for each object version when it is uploaded.
- Bucket lifecycle configuration is supported for buckets with S3 Object Lock enabled.
- CloudMirror replication is not supported for buckets with S3 Object Lock enabled.

Requirements for objects in buckets with S3 Object Lock enabled

- To protect an object version, you can specify default retention settings for the bucket, or you can specify retention settings for each object version. Object-level retention settings can be specified using the S3 client application or the S3 REST API.
- Retention settings apply to individual object versions. An object version can have both a retain-until-date and a legal hold setting, one but not the other, or neither. Specifying a retain-until-date or a legal hold setting for an object protects only the version specified in the request. You can create new versions of the object, while the previous version of the object remains locked.

Lifecycle of objects in buckets with S3 Object Lock enabled

Each object that is saved in a bucket with S3 Object Lock enabled goes through these stages:

1. Object ingest

When an object version is added to bucket that has S3 Object Lock enabled, retention settings are applied as follows:

- If retention settings are specified for the object, the object-level settings are applied. Any default bucket settings are ignored.
- If no retention settings are specified for the object, the default bucket settings are applied, if they exist.
- If no retention settings are specified for the object or the bucket, the object is not protected by S3 Object Lock.

If retention settings are applied, both the object and any S3 user-defined metadata are protected.

2. Object retention and deletion

Multiple copies of each protected object are stored by StorageGRID for the specified retention period. The exact number and type of object copies and the storage locations are determined by the compliant rules in the active ILM policies. Whether a protected object can be deleted before its retain-until-date is reached depends on its retention mode.

• If an object is under a legal hold, no one can delete the object, regardless of its retention mode.

Related information

- Create an S3 bucket
- Update S3 Object Lock default retention
- Use S3 REST API to configure S3 Object Lock
- Example 7: Compliant ILM policy for S3 Object Lock

Enable S3 Object Lock globally

If an S3 tenant account needs to comply with regulatory requirements when saving object data, you must enable S3 Object Lock for your entire StorageGRID system. Enabling the global S3 Object Lock setting allows any S3 tenant user to create and manage buckets and objects with S3 Object Lock.

Before you begin

- You have the Root access permission.
- You are signed in to the Grid Manager using a supported web browser.
- You have reviewed the S3 Object Lock workflow, and you understand the considerations.
- You have confirmed that the default rule in the active ILM policy is compliant. See Create a default ILM rule for details.

About this task

A grid administrator must enable the global S3 Object Lock setting to allow tenant users to create new buckets that have S3 Object Lock enabled. After this setting is enabled, it can't be disabled.



The global Compliance setting is deprecated. If you enabled this setting using a previous version of StorageGRID, the S3 Object Lock setting is enabled automatically. You can continue to use StorageGRID to manage the settings of existing compliant buckets; however, you can't create new compliant buckets. For details, see NetApp Knowledge Base: How to manage legacy Compliant buckets in StorageGRID 11.5.

Steps

1. Select CONFIGURATION > System > S3 Object Lock.

The S3 Object Lock Settings page appears.

2. Select Enable S3 Object Lock.

3. Select Apply.

A confirmation dialog box appears and reminds you that you can't disable S3 Object Lock after it is enabled.

4. If you are sure you want to permanently enable S3 Object Lock for your entire system, select **OK**.

When you select OK:

- If the default rule in the active ILM policy is compliant, S3 Object Lock is now enabled for the entire grid and can't be disabled.
- If the default rule is not compliant, an error appears. You must create and activate a new ILM policy that includes a compliant rule as its default rule. Select **OK**. Then, create a new policy, simulate it, and activate it. See Create ILM policy for instructions.

Resolve consistency errors when updating the S3 Object Lock or legacy Compliance configuration

If a data center site or multiple Storage Nodes at a site become unavailable, you might need to help S3 tenant users apply changes to the S3 Object Lock or legacy Compliance configuration.

Tenant users who have buckets with S3 Object Lock (or legacy Compliance) enabled can change certain settings. For example, a tenant user using S3 Object Lock might need to put an object version under legal hold.

When a tenant user updates the settings for an S3 bucket or an object version, StorageGRID attempts to immediately update the bucket or object metadata across the grid. If the system is unable to update the metadata because a data center site or multiple Storage Nodes are unavailable, it returns an error:

```
503: Service Unavailable
Unable to update compliance settings because the settings can't be
consistently applied on enough storage services. Contact your grid
administrator for assistance.
```

To resolve this error, follow these steps:

- 1. Attempt to make all Storage Nodes or sites available again as soon as possible.
- 2. If you are unable to make enough of the Storage Nodes at each site available, contact technical support, who can help you recover nodes and ensure that changes are consistently applied across the grid.
- 3. Once the underlying issue has been resolved, remind the tenant user to retry their configuration changes.

Related information

- Use a tenant account
- Use S3 REST API
- Recover and maintain

Example ILM rules and policies

Example 1: ILM rules and policy for object storage

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



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss.

ILM rule 1 for example 1: Copy object data to two sites

This example ILM rule copies object data to storage pools in two sites.

Rule definition	Example value
One-site storage pools	Two storage pools, each containing different sites, named Site 1 and Site 2.
Rule name	Two Copies Two Sites
Reference time	Ingest time
Placements	On Day 0 to forever, keep one replicated copy at Site 1 and one replicated copy at Site 2.

The Rule analysis section of the Retention diagram states:

- StorageGRID site-loss protection will apply for the duration of this rule.
- · Objects processed by this rule will not be deleted by ILM.

eference time 🥝		
Ingest time	*	
ime period and p	lacements	Sort by start date
you want a rule to a ne criteria in the filte	apply only to specific objects, select Previous and add advanced filters. When objects are evaluated, the er.	e rule is applied if the object's metadata matches
Time period 1	From Day 0 Store forever	×
Store objects by	replicating V 1 C copies at Site 1 X X	
and store objects	by replicating V 1 C copies at Site 2 X X	
Add other type or	rlocation	
ld anoth <mark>e</mark> r time per	riod	
etention diagram	n	Replicated of the second se
e analysis: • Sto • Ob	orageGRID site-loss protection will apply for the duration of this rule. Ojects processed by this rule will not be deleted by ILM.	
eference time: Ingest	t time Day 0	
Day 0 - forever	1 replicated copy - Site 1	
	1 replicated copy - Site 2	
Juration	Enrever	

ILM rule 2 for example 1: Erasure-coding profile with bucket matching

This example ILM rule uses an erasure-coding profile and an S3 bucket to determine where and how long the object is stored.

Rule definition	Example value
Storage pool with multiple sites	 One storage pool across three sites (Sites 1, 2, 3) Use 6+3 erasure-coding scheme
Rule name	S3 Bucket finance-records
Reference time	Ingest time
Placements	For objects in the S3 bucket named finance-records, create one erasure-coded copy in the pool specified by the erasure-coding profile. Keep this copy forever.

Time period and placements

Sort by start date

If you want a rule to apply only to specific objects, select **Previous** and add advanced filters. When objects are evaluated, the rule is applied if the object's metadata matches the criteria in the filter.

Time period 1	From Day 0	© store	forever 🗸		×
Store objects by	y erasure codin	ng 💙 using 6	+3 EC scheme at Sites 1,	2,3 🧪 🗙	
Add other type	or location				
Add another time p	eriod				
Retention diagra	ım				Erasure-coded (EC) copy
Rule analysis: •	StorageGRID site-loss pr Objects processed by th	rotection will apply for the is rule will not be deleted t	duration of this rule. by ILM.		
Reference time: Inge	est time Day 0				
Day 0 - forever	EC 6+3 - Sites	1, 2, 3			×
Duration				Forever	

ILM policy for example 1

In practice, most ILM policies are simple, even though the StorageGRID system allows you to design sophisticated and complex ILM policies.

A typical ILM policy for a multi-site grid might include ILM rules such as the following:

- At ingest, store all objects belonging to the S3 bucket named finance-records in a storage pool that contains three sites. Use 6+3 erasure coding.
- If an object does not match the first ILM rule, use the policy's default ILM rule, Two Copies Two Data Centers, to store one copy of that object in Site 1, and one copy in Site 2.

Proposed policy na	me		
Object Storage Po	licy		
Reason for change			
example 1			
Manage rules 1. Select the rules you w 2. Determine the order i Select rules	rant to add to the policy. n which the rules will be evaluated by dragging and dropping the rows. The default r	ule will be automatically placed at the end of the policy and cannot be moved.	
Rule order	Rule name	Filters	
1	 S3 Bucket finance-records Ø 	Tenant is Finance Bucket name is finance-records	
Default	Two Copies Two Data Centers		

Related information

- ILM policies: Overview
- Create ILM policies

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

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



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss.

ILM rule 1 for example 2: Use EC for objects greater than 1 MB

This example ILM rule erasure codes objects that are greater than 1 MB.



Erasure coding is best suited for objects greater than 1 MB. Don't use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.

Rule definition	Example value
Rule name	EC Only Objects > 1 MB
Reference time	Ingest time
Advanced filter for Object size	Object size greater than 1 MB
Placements	Create a 2+1 erasure-coded copy using three sites

ilter group 1	Objects with all of follow	ing metadata will be eva	aluated by thi	s rule:					×
Object size	~	greater than	~	1	0	MB	~	×	

ILM rule 2 for example 2: Two replicated copies

This example ILM rule creates two replicated copies and does not filter by object size. This rule is the default rule for the policy. Because the first rule filters out all objects greater than 1 MB, this rule only applies to objects that are 1 MB or smaller.

Rule definition	Example value
Rule name	Two Replicated Copies
Reference time	Ingest time
Advanced filter for Object size	None
Placements	On Day 0 to forever, keep one replicated copy at Site 1 and one replicated copy at Site 2.

ILM policy for example 2: Use EC for objects greater than 1 MB

This example ILM policy includes two ILM rules:

- The first rule erasure codes all objects that are greater than 1 MB.
- The second (default) ILM rule creates two replicated copies. Because objects greater than 1 MB have been filtered out by rule 1, rule 2 only applies to objects that are 1 MB or smaller.

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

You can use the following example rules and policy to ensure that images greater than 1 MB are erasure-coded and that two copies are made of smaller images.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss.

ILM rule 1 for example 3: Use EC for image files greater than 1 MB

This example ILM rule uses advanced filtering to erasure code all image files greater than 1 MB.



Erasure coding is best suited for objects greater than 1 MB. Don't use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.

Rule definition	Example value
Rule name	EC Image Files > 1 MB

Rule definition	Example value
Reference time	Ingest time
Advanced filter for Object size	Object size greater than 1 MB
Advanced filters for Key	Ends with .jpgEnds with .png
Placements	Create a 2+1 erasure-coded copy using three sites

Object size	~	greater than	~	1	0	MB	~	×	
		0							
and Key		 ends with 		.jpg			×		
Filter group	2 Objects with all of f	ollowing metadata will	be evaluated b	ov this rule:					
Filter group	2 Objects with all of f	ollowing metadata will	be evaluated b	by this rule:					
Filter group	2 Objects with all of fo	ollowing metadata will greater than	be evaluated b	by this rule:	0	MB	~	×	

Because this rule is configured as the first rule in the policy, the erasure-coding placement instruction only applies to .jpg and .png files that are greater than 1 MB.

ILM rule 2 for example 3: Create 2 replicated copies for all remaining image files

This example ILM rule uses advanced filtering to specify that smaller image files be replicated. Because the first rule in the policy has already matched image files greater than 1 MB, this rule applies to image files that are 1 MB or smaller.

Rule definition	Example value
Rule name	2 Copies for Image Files
Reference time	Ingest time
Advanced filters for Key	Ends with .jpgEnds with .png
Placements	Create 2 replicated copies in two storage pools

ILM policy for example 3: Better protection for image files

This example ILM policy includes three rules:

- The first rule erasure codes all image files greater than 1 MB.
- The second rule creates two copies of any remaining image files (that is, images that are 1 MB or smaller).
- The default rule applies to all remaining objects (that is, any non-image files).

Rule order	Rule name	Filters
1	EC image files > 1 MB	Object size is greater than 1 MB
2	2 copies for small images	Object size is less than or equal to 200 KB
Default	Default rule	-

Example 4: ILM rules and policy for S3 versioned objects

If you have an S3 bucket with versioning enabled, you can manage the noncurrent object versions by including rules in your ILM policy that use "Noncurrent time" as the reference time.



If you specify a limited retention time for objects, those objects will be deleted permanently after the time period is reached. Make sure you understand how long the objects will be retained.

As this example shows, you can control the amount of storage used by versioned objects by using different placement instructions for noncurrent object versions.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss.



To perform ILM policy simulation on a noncurrent version of an object, you must know the object version's UUID or CBID. To find the UUID and CBID, use object metadata lookup while the object is still current.

Related information

• How objects are deleted

ILM rule 1 for example 4: Save three copies for 10 years

This example ILM rule stores a copy of each object at three sites for 10 years.

This rule applies to all objects, whether or not they are versioned.

Rule definition	Example value
Storage pools	Three storage pools, each consisting of different data centers, named Site 1, Site 2, and Site 3.
Rule name	Three Copies Ten Years
Reference time	Ingest time
Placements	On Day 0, keep three replicated copies for 10 years (3,652 days), one in Site 1, one in Site 2, and one in Site 3. At the end of 10 years, delete all copies of the object.

ILM rule 2 for example 4: Save two copies of noncurrent versions for 2 years

This example ILM rule stores two copies of the noncurrent versions of an S3 versioned object for 2 years.

Because ILM rule 1 applies to all versions of the object, you must create another rule to filter out any noncurrent versions.

To create a rule that uses "Noncurrent time" as the reference time, select **Yes** for the question, "Apply this rule to older object versions only (in S3 buckets with versioning enabled)?" in Step 1 (Enter details) of the Create an ILM rule wizard. When you select **Yes**, *Noncurrent time* is automatically selected for the reference time, and you can't select a different reference time.

Rule name	
Older Object Version	as: Two Copies Two Years
Description (optional)	
Older versions only	
Basic filters (optional) Specify which tenant account	s and buckets this rule applies to.
Tenant accounts 👩	Select tenant accounts
Bucket name 💡	matches all 🗸 🗸

In this example, only two copies of the noncurrent versions are stored, and those copies will be stored for two years.

Rule definition	Example value
Storage Pools	Two storage pools, each at different data centers, Site 1 and Site 2.
Rule name	Noncurrent Versions: Two Copies Two Years
Reference time	Noncurrent time Automatically selected when you select Yes for the question, "Apply this rule to older object versions only (in S3 buckets with versioning enabled)?" in the Create an ILM rule wizard.
Placements	On Day 0 relative to noncurrent time (that is, starting from the day the object version becomes the noncurrent version), keep two replicated copies of the noncurrent object versions for 2 years (730 days), one in Site 1 and one in Site 2. At the end of 2 years, delete the noncurrent versions.

ILM policy for example 4: S3 versioned objects

If you want to manage older versions of an object differently than the current version, rules that use "Noncurrent time" as the reference time must appear in the ILM policy before rules that apply to the current object version.

An ILM policy for S3 versioned objects might include ILM rules such as the following:

• Keep any older (noncurrent) versions of each object for 2 years, starting from the day the version became noncurrent.



The "Noncurrent time" rules must appear in the policy before the rules that apply to the current object version. Otherwise, the noncurrent object versions will never be matched by the "Noncurrent time" rule.

• At ingest, create three replicated copies and store one copy at each of three sites. Keep copies of the current object version for 10 years.

When you simulate the example policy, you would expect test objects to be evaluated as follows:

- Any noncurrent object versions would be matched by the first rule. If a noncurrent object version is older than 2 years, it is permanently deleted by ILM (all copies of the noncurrent version removed from the grid).
- The current object version would be matched by the second rule. When the current object version has been stored for 10 years, the ILM process adds a delete marker as the current version of the object, and it makes the previous object version "noncurrent". The next time ILM evaluation occurs, this noncurrent version is matched by the first rule. As a result, the copy at Site 3 is purged and the two copies at Site 1 and Site 2 are stored for 2 more years.

Example 5: ILM rules and policy for Strict ingest behavior

You can use a location filter and the Strict ingest behavior in a rule to prevent objects from being saved at a particular data center location.

In this example, a Paris-based tenant does not want to store some objects outside of the EU because of regulatory concerns. Other objects, including all objects from other tenant accounts, can be stored at either the Paris data center or the US data center.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss.

Related information

- Ingest options
- Create ILM rule: Select ingest behavior

ILM rule 1 for example 5: Strict ingest to guarantee Paris data center

This example ILM rule uses the Strict ingest behavior to guarantee that objects saved by a Paris-based tenant to S3 buckets with the region set to eu-west-3 region (Paris) are never stored at the US data center.

This rule applies to objects that belong to the Paris tenant and that have the S3 bucket region set to eu-west-3 (Paris).

Rule definition	Example value
Tenant account	Paris tenant
Advanced filter	Location constraint equals eu-west-3
Storage pools	Site 1 (Paris)
Rule name	Strict ingest to guarantee Paris data center
Reference time	Ingest time
Placements	On Day 0, keep two replicated copies forever in Site 1 (Paris)
Ingest behavior	Strict. Always use this rule's placements on ingest. Ingest fails if it is not possible to store two copies of the object at the Paris data center.

Strict inges	st to guarantee Pa	ris data cen	ter		
Compliant:	Yes		Ingest behavior:	Strict	
Used in active policy:	No		Reference time:	Ingest time	
Used in proposed pol	icy: No				
Clone Edit	Remove				
Filters					
This rule applies if: • Tenant is Paris tenan And it only applies if obje • Location constraint	nt cts have this metadata: is eu-west-3				*
Time period and place Retention diago	ram Placement instructio	ons			
Sort placements by	me period Storage pool				Replicated copy
Rule analysis: • Sto • Ob	prageGRID site-loss protection will not apply fro jects processed by this rule will not be deleted i	om Day 0 - Forever. by ILM.			
Reference time: Ingest	time Ingest behavior: Strict				
Day 0 - forever	2 replicated copies - Site 1				
Duration			Forever		

ILM rule 2 for example 5: Balanced ingest for other objects

This example ILM rule uses the Balanced ingest behavior to provide optimum ILM efficiency for any objects not matched by the first rule. Two copies of all objects matched by this rule will be stored—one at the US data center and one at the Paris data center. If the rule can't be satisfied immediately, interim copies are stored at any available location.

This rule applies to objects that belong to any tenant and any region.

Rule definition	Example value
Tenant account	Ignore
Advanced filter	Not specified
Storage pools	Site 1 (Paris) and Site 2 (US)
Rule name	2 Copies 2 Data Centers
Reference time	Ingest time
Placements	On Day 0, keep two replicated copies forever at two data centers

Rule definition	Example value
Ingest behavior	Balanced. Objects that match this rule are placed according to the rule's placement instructions if possible. Otherwise, interim copies are made at any available location.

ILM policy for example 5: Combining ingest behaviors

The example ILM policy includes two rules that have different ingest behaviors.

An ILM policy that uses two different ingest behaviors might include ILM rules such as the following:

- Store objects that belong to the Paris tenant and that have the S3 bucket region set to eu-west-3 (Paris) only in the Paris data center. Fail ingest if the Paris data center is not available.
- Store all other objects (including those that belong to the Paris tenant but that have a different bucket region) in both the US data center and the Paris data center. Make interim copies in any available location if the placement instruction can't be satisfied.

When you simulate the example policy, you expect test objects to be evaluated as follows:

- Any objects that belong to the Paris tenant and that have the S3 bucket region set to eu-west-3 are matched by the first rule and are stored at the Paris data center. Because the first rule uses Strict ingest, these objects are never stored at the US data center. If the Storage Nodes at the Paris data center aren't available, ingest fails.
- All other objects are matched by the second rule, including objects that belong to the Paris tenant and that don't have the S3 bucket region set to eu-west-3. One copy of each object is saved at each data center. However, because the second rule uses Balanced ingest, if one data center is unavailable, two interim copies are saved at any available location.

Example 6: Change an ILM policy

If your data protection needs to be changed or you add new sites, you can create and activate a new ILM policy.

Before changing a policy, you must understand how changes in ILM placements can temporarily affect the overall performance of a StorageGRID system.

In this example, a new StorageGRID site has been added in an expansion, and a new active ILM policy needs to be implemented to store data at the new site. To implement a new active policy, first create a policy. Afterward, you must simulate and then activate the new policy.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss.

How changing an ILM policy affects performance

When you activate a new ILM policy, the performance of your StorageGRID system might be temporarily affected, especially if the placement instructions in the new policy require many existing objects to be moved to new locations.

When you activate a new ILM policy, StorageGRID uses it to manage all objects, including existing objects and

newly ingested objects. Before activating a new ILM policy, review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

To ensure a new ILM policy does not affect the placement of existing replicated and erasure-coded objects, you can create an ILM rule with an ingest time filter. For example, **Ingest time** *is on or after <date and time>*, so that the new rule applies only to objects ingested on or after the date and time specified.

The types of ILM policy changes that can temporarily affect StorageGRID performance include the following:

• Applying a different erasure-coding profile to existing erasure-coded objects.



StorageGRID considers each erasure-coding profile to be unique and does not reuse erasure-coding fragments when a new profile is used.

- Changing the type of copies required for existing objects; for example, converting a large percentage of replicated objects to erasure-coded objects.
- Moving copies of existing objects to a completely different location; for example, moving a large number of objects to or from a Cloud Storage Pool or to or from a remote site.

Active ILM policy for example 6: Data protection at two sites

In this example, the active ILM policy was initially designed for a two-site StorageGRID system and uses two ILM rules.

Active policy	Policy history		
olicy name:	Data Protection for Two Sites (2 rules)		
ason for change :	Data protection for two sites (using 2 rules)		
art date:	2022-10-11 10:37:11 MDT		
Simulate			
Policy rules	Retention diagram		
Policy rules	Retention diagram Rule name	Filters 🚱	
Policy rules Rule order @	Retention diagram Rule name One-Site Erasure Coding for Tenant A	Filters 🎯 Tenant is Tenant A	

In this ILM policy, objects belonging to Tenant A are protected by 2+1 erasure coding at a single site, while objects belonging to all other tenants are protected across two sites using 2-copy replication.

Rule 1: One-site erasure coding for Tenant A

Rule definition	Example value
Rule name	One-Site Erasure Coding for Tenant A
Rule definition	Example value
-----------------	--
Tenant Account	Tenant A
Storage Pool	Site 1
Placements	2+1 erasure coding in Site 1 from day 0 to forever

Rule 2: Two-site replication for other tenants

Rule definition	Example value
Rule name	Two-Site Replication for Other Tenants
Tenant Account	Ignore
Storage Pools	Site 1 and Site 2
Placements	Two replicated copies from day 0 to forever: one copy at Site 1 and one copy at Site 2.

ILM policy for example 6: Data protection at three sites

In this example, the ILM policy is being replaced with a new policy for a three-site StorageGRID system.

After performing an expansion to add the new site, the grid administrator created two new storage pools: a storage pool for Site 3 and a storage pool containing all three sites (not the same as the All Storage Nodes default storage pool). Then, the administrator created two new ILM rules and a new ILM policy, which is designed to protect data at all three sites.

When this new ILM policy is activated, objects belonging to Tenant A will be protected by 2+1 erasure coding at three sites, while objects belonging to other tenants (and smaller objects belonging to Tenant A) will be protected across three sites using 3-copy replication.

Rule 1: Three-site erasure coding for Tenant A

Rule definition	Example value
Rule name	Three-Site Erasure Coding for Tenant A
Tenant Account	Tenant A
Storage Pool	All 3 Sites (includes Site 1, Site 2, and Site 3)
Placements	2+1 erasure coding in All 3 Sites from day 0 to forever

Rule 2: Three-site replication for other tenants

Rule definition	Example value
Rule name	Three-Site Replication for Other Tenants
Tenant Account	Ignore
Storage Pools	Site 1, Site 2, and Site 3
Placements	Three replicated copies from day 0 to forever: one copy at Site 1, one copy at Site 2, and one copy at Site 3.

Activating the ILM policy for example 6

When you activate a new ILM policy, existing objects might be moved to new locations or new object copies might be created for existing objects, based on the placement instructions in any new or updated rules.



Errors in an ILM policy can cause unrecoverable data loss. Carefully review and simulate the policy before activating it to confirm that it will work as intended.



When you activate a new ILM policy, StorageGRID uses it to manage all objects, including existing objects and newly ingested objects. Before activating a new ILM policy, review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

What happens when erasure-coding instructions change

In the currently active ILM policy for this example, objects belonging to Tenant A are protected using 2+1 erasure coding at Site 1. In the new ILM policy, objects belonging to Tenant A will be protected using 2+1 erasure coding at Sites 1, 2, and 3.

When the new ILM policy is activated, the following ILM operations occur:

- New objects ingested by Tenant A are split into two data fragments and one parity fragment is added. Then, each of the three fragments is stored at a different site.
- The existing objects belonging to Tenant A are re-evaluated during the ongoing ILM scanning process. Because the ILM placement instructions use a new erasure-coding profile, entirely new erasure-coded fragments are created and distributed to the three sites.



The existing 2+1 fragments at Site 1 aren't reused. StorageGRID considers each erasurecoding profile to be unique and does not reuse erasure-coding fragments when a new profile is used.

What happens when replication instructions change

In the currently active ILM policy for this example, objects belonging to other tenants are protected using two replicated copies in storage pools at Sites 1 and 2. In the new ILM policy, objects belonging to other tenants will be protected using three replicated copies in storage pools at Sites 1, 2, and 3.

When the new ILM policy is activated, the following ILM operations occur:

- When any tenant other than Tenant A ingests a new object, StorageGRID creates three copies and saves one copy at each site.
- Existing objects belonging to these other tenants are re-evaluated during the ongoing ILM scanning process. Because the existing object copies at Site 1 and Site 2 continue to satisfy the replication requirements of the new ILM rule, StorageGRID only needs to create one new copy of the object for Site 3.

Performance impact of activating this policy

When the ILM policy in this example is activated, the overall performance of this StorageGRID system will be temporarily affected. Higher than normal levels of grid resources will be required to create new erasure-coded fragments for Tenant A's existing objects and new replicated copies at Site 3 for other tenants' existing objects.

As a result of the ILM policy change, client read and write requests might temporarily experience higher than normal latencies. Latencies will return to normal levels after the placement instructions are fully implemented across the grid.

To avoid resource issues when activating a new ILM policy, you can use the Ingest time advanced filter in any rule that might change the location of large numbers of existing objects. Set Ingest time to be greater than or equal to the approximate time when the new policy will go into effect to ensure that existing objects aren't moved unnecessarily.



Contact technical support if you need to slow or increase the rate at which objects are processed after an ILM policy change.

Example 7: Compliant ILM policy for S3 Object Lock

You can use the S3 bucket, ILM rules, and ILM policy in this example as a starting point when defining an ILM policy to meet the object protection and retention requirements for objects in buckets with S3 Object Lock enabled.



If you used the legacy Compliance feature in previous StorageGRID releases, you can also use this example to help manage any existing buckets that have the legacy Compliance feature enabled.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss.

Related information

- Manage objects with S3 Object Lock
- Create an ILM policy

Bucket and objects for S3 Object Lock example

In this example, an S3 tenant account named Bank of ABC has used the Tenant Manager to create a bucket with S3 Object Lock enabled to store critical bank records.

Bucket definition	Example value
Tenant Account Name	Bank of ABC
Bucket Name	bank-records
Bucket Region	us-east-1 (default)

Each object and object version that is added to the bank-records bucket will use the following values for retain-until-date and legal hold settings.

Setting for each object	Example value
retain-until-date	"2030-12-30T23:59:59Z" (December 30, 2030) Each object version has its own retain-until-date setting. This setting can be increased, but not decreased.
legal hold	"OFF" (Not in effect) A legal hold can be placed or lifted on any object version at any time during the retention period. If an object is under a legal hold, the object can't be deleted even if the retain-until-date has been reached.

ILM rule 1 for S3 Object Lock example: Erasure-coding profile with bucket matching

This example ILM rule applies only to the S3 tenant account named Bank of ABC. It matches any object in the bank-records bucket and then uses erasure coding to store the object on Storage Nodes at three data center sites using a 6+3 erasure-coding profile. This rule satisfies the requirements of buckets with S3 Object Lock enabled: a copy is kept on Storage Nodes from day 0 to forever, using Ingest time as the reference time.

Rule definition	Example value
Rule name	Compliant Rule: EC Objects in bank-records Bucket - Bank of ABC
Tenant Account	Bank of ABC
Bucket Name	bank-records
Advanced filter	Object Size (MB) greater than 1 Note: This filter ensures that erasure coding is not used for objects 1 MB or smaller.
Rule definition	Example value

Rule definition	Example value
Reference time	Ingest time

Rule definition	Example value
Placements	From day 0 store forever
Erasure-coding profile	 Create an erasure-coded copy on Storage Nodes at three data center sites Uses 6+3 erasure-coding scheme

ILM rule 2 for S3 Object Lock example: Non-compliant rule

This example ILM rule initially stores two replicated object copies on Storage Nodes. After one year, it stores one copy on a Cloud Storage Pool forever. Because this rule uses a Cloud Storage Pool, it is not compliant and will not apply to the objects in buckets with S3 Object Lock enabled.

Rule definition	Example value
Rule name	Non-compliant rule: Use Cloud Storage Pool
Tenant accounts	Not specified
Bucket name	Not specified, but will only apply to buckets that don't have S3 Object Lock (or the legacy Compliance feature) enabled.
Advanced filter	Not specified

Rule definition	Example value
Reference time	Ingest time
Placements	 On Day 0, keep two replicated copies on Storage Nodes in Data Center 1 and Data Center 2 for 365 days After 1 year, keep one replicated copy in a Cloud Storage Pool forever

ILM rule 3 for S3 Object Lock example: Default rule

This example ILM rule copies object data to storage pools in two data centers. This compliant rule is designed to be the default rule in the ILM policy. It does not include any filters, does not use the Noncurrent reference time, and satisfies the requirements of buckets with S3 Object Lock enabled: two object copies are kept on Storage Nodes from day 0 to forever, using Ingest as the reference time.

Rule definition	Example value
Rule name	Default compliant rule: Two Copies Two Data Centers
Tenant account	Not specified

Rule definition	Example value
Bucket name	Not specified
Advanced filter	Not specified
Rule definition	Example value
Reference time	Ingest time
Placements	From Day 0 to forever, keep two replicated copies—one on Storage

Compliant ILM policy for S3 Object Lock example

To create an ILM policy that will effectively protect all objects in your system, including those in buckets with S3 Object Lock enabled, you must select ILM rules that satisfy the storage requirements for all objects. Then, you must simulate and activate the policy.

Nodes in Data Center 1 and one on Storage Nodes in Data Center 2.

Add rules to the policy

In this example, the ILM policy includes three ILM rules, in the following order:

- 1. A compliant rule that uses erasure coding to protect objects greater than 1 MB in a specific bucket with S3 Object Lock enabled. The objects are stored on Storage Nodes from day 0 to forever.
- A non-compliant rule that creates two replicated object copies on Storage Nodes for a year and then moves one object copy to a Cloud Storage Pool forever. This rule does not apply to buckets with S3 Object Lock enabled because it uses a Cloud Storage Pool.
- 3. The default compliant rule that creates two replicated object copies on Storage Nodes from day 0 to forever.

Simulate the policy

After you have added rules to your policy, chosen a default compliant rule, and arranged the other rules, you should simulate the policy by testing objects from the bucket with S3 Object Lock enabled and from other buckets. For example, when you simulate the example policy, you would expect test objects to be evaluated as follows:

- The first rule will only match test objects that are greater than 1 MB in the bucket bank-records for the Bank of ABC tenant.
- The second rule will match all objects in all non-compliant buckets for all other tenant accounts.
- The default rule will match these objects:
 - Objects 1 MB or smaller in the bucket bank-records for the Bank of ABC tenant.
 - Objects in any other bucket that has S3 Object Lock enabled for all other tenant accounts.

Activate the policy

When you are completely satisfied that the new policy protects object data as expected, you can activate it.

Example 8: Priorities for S3 bucket lifecycle and ILM policy

Depending on your lifecycle configuration, objects follow the retention settings of either the S3 bucket lifecycle or an ILM policy.

Example of bucket lifecycle taking priority over ILM policy

ILM policy

- Rule based on noncurrent-time reference: On Day 0, keep X copies for 20 days
- Rule based on ingest-time reference (default): On Day 0, keep X copies for 50 days

Bucket Lifecycle

```
• Filter: {Prefix: "docs/"}, Expiration: Days: 100,
NoncurrentVersionExpiration: Days: 5
```

Result

- An object named "docs/text" is ingested. It matches the bucket lifecycle filter of "docs/" prefix.
 - After 100 days a delete-marker is created and "docs/text" becomes noncurrent.
 - After 5 days, a total of 105 days since ingest, "docs/text" is deleted.
- An object named "video/movie" is ingested. It does not match the filter and uses the ILM retention policy.
 - After 50 days a delete-marker is created and "video/movie" becomes noncurrent.
 - After 20 days, a total of 70 days since the ingest, "video/movie" is deleted.

Example of bucket lifecycle implicitly keeping-forever

ILM policy

- Rule based on noncurrent-time reference: On Day 0, keep X copies for 20 days
- Rule based on ingest-time reference (default): On Day 0, keep X copies for 50 days

Bucket Lifecycle

• Filter: {Prefix: "docs/"}, Expiration: ExpiredObjectDeleteMarker: true

Result

• An object named "docs/text" is ingested. It matches the bucket lifecycle filter of "docs/" prefix.

The Expiration action applies only to expired delete markers, which implies keeping everything else forever (starting with "docs/").

Delete markers that start with "docs/" are removed when they become expired.

- An object named "video/movie" is ingested. It does not match the filter and uses the ILM retention policy.
 - After 50 days a delete-marker is created and "video/movie" becomes noncurrent.
 - After 20 days, a total of 70 days since the ingest, "video/movie" is deleted.

Example of using bucket lifecycle to duplicate ILM and clean up expired delete markers

ILM policy

- Rule based on noncurrent-time reference: On Day 0, keep X copies for 20 days
- Rule based on ingest-time reference (default): On Day 0, keep X copies for 50 days

Bucket Lifecycle

```
• Filter: {}, Expiration: Days: 50, NoncurrentVersionExpiration: Days: 20
```

Result

- The ILM policy is duplicated in the bucket lifecycle.
- An object is ingested. No filter means that the bucket lifecycle applies to all objects and overrides the ILM retention settings.
 - $\circ\,$ After 50 days a delete-marker is created and the object becomes noncurrent.
 - After 20 days, a total of 70 days since the ingest, the noncurrent object is deleted and the deletemarker becomes expired.
 - After 30 days, a total of 100 days since the ingest, the expired delete-marker is deleted.

System hardening

System hardening: Overview

System hardening is the process of eliminating as many security risks as possible from a StorageGRID system.

This document provides an overview of the hardening guidelines that are specific to StorageGRID. These guidelines are a supplement to industry-standard best practices for system hardening. For example, these guidelines assume that you use strong passwords for StorageGRID, use HTTPS instead of HTTP, and enable certificate-based authentication where available.

As you install and configure StorageGRID, you can use these guidelines to help you meet any prescribed security objectives for information system confidentiality, integrity, and availability.

StorageGRID follows the NetApp Vulnerability Handling Policy. Reported vulnerabilities are verified and addressed according to the product security incident response process.

General considerations for hardening StorageGRID systems

When hardening a StorageGRID system, you must consider the following:

- Which of the three StorageGRID networks you have implemented. All StorageGRID systems must use the Grid Network, but you might also be using the Admin Network, the Client Network, or both. Each network has different security considerations.
- The type of platforms you use for the individual nodes in your StorageGRID system. StorageGRID nodes can be deployed on VMware virtual machines, within a container engine on Linux hosts, or as dedicated hardware appliances. Each type of platform has its own set of hardening best practices.
- How trusted the tenant accounts are. If you are a service provider with untrusted tenant accounts, you will have different security concerns than if you only use trusted, in-house tenants.
- Which security requirements and conventions are followed by your organization. You might need to comply with specific regulatory or corporate requirements.

Hardening guidelines for software upgrades

You must keep your StorageGRID system and related services up to date to defend against attacks.

Upgrades to StorageGRID software

Whenever possible, you should upgrade StorageGRID software to the most recent major release or to the previous major release. Keeping StorageGRID up to date helps reduce the amount of time that known vulnerabilities are active and reduces the overall attack surface area. In addition, the most recent releases of StorageGRID often contain security hardening features that aren't included in earlier releases.

Consult the NetApp Interoperability Matrix Tool (IMT) to determine which version of StorageGRID software you should be using. When a hotfix is required, NetApp prioritizes creating updates for the most recent releases. Some patches might not be compatible with earlier releases.

- To download the most recent StorageGRID releases and hotfixes, go to NetApp Downloads: StorageGRID.
- To upgrade StorageGRID software, see the upgrade instructions.
- To apply a hotfix, see the StorageGRID hotfix procedure.

Upgrades to external services

External services can have vulnerabilities that affect StorageGRID indirectly. You should ensure that the services that StorageGRID depends on are kept up to date. These services include LDAP, KMS (or KMIP server), DNS, and NTP.

For a list of supported versions, see the NetApp Interoperability Matrix Tool.

Upgrades to hypervisors

If your StorageGRID nodes are running on VMware or another hypervisor, you must ensure that the hypervisor software and firmware are up to date.

For a list of supported versions, see the NetApp Interoperability Matrix Tool.

Upgrades to Linux nodes

If your StorageGRID nodes are using Linux host platforms, you must ensure that security updates and kernel updates are applied to the host OS. Additionally, you must apply firmware updates to vulnerable hardware when these updates become available.

For a list of supported versions, see the NetApp Interoperability Matrix Tool.

Hardening guidelines for StorageGRID networks

The StorageGRID system supports up to three network interfaces per grid node, allowing you to configure the networking for each individual grid node to match your security and access requirements.

For detailed information about StorageGRID networks, see the StorageGRID network types.

Guidelines for Grid Network

You must configure a Grid Network for all internal StorageGRID traffic. All grid nodes are on the Grid Network, and they must be able to talk to all other nodes.

When configuring the Grid Network, follow these guidelines:

- Ensure that the network is secured from untrusted clients, such as those on the open internet.
- When possible, use the Grid Network exclusively for internal traffic. Both the Admin Network and the Client Network have additional firewall restrictions that block external traffic to internal services. Using the Grid Network for external client traffic is supported, but this use offers fewer layers of protection.
- If the StorageGRID deployment spans multiple data centers, use a virtual private network (VPN) or equivalent on the Grid Network to provide additional protection for internal traffic.
- Some maintenance procedures require secure shell (SSH) access on port 22 between the primary Admin Node and all other grid nodes. Use an external firewall to restrict SSH access to trusted clients.

Guidelines for Admin Network

The Admin Network is typically used for administrative tasks (trusted employees using the Grid Manager or SSH) and for communicating with other trusted services such as LDAP, DNS, NTP, or KMS (or KMIP server). However, StorageGRID does not enforce this usage internally.

If you are using the Admin Network, follow these guidelines:

- Block all internal traffic ports on the Admin Network. See the list of internal ports.
- If untrusted clients can access the Admin Network, block access to StorageGRID on the Admin Network with an external firewall.

Guidelines for Client Network

The Client Network is typically used for tenants and for communicating with external services, such as the CloudMirror replication service or another platform service. However, StorageGRID does not enforce this usage internally.

If you are using the Client Network, follow these guidelines:

- Block all internal traffic ports on the Client Network. See the list of internal ports.
- Accept inbound client traffic only on explicitly configured endpoints. See the information about managing firewall controls.

Hardening guidelines for StorageGRID nodes

StorageGRID nodes can be deployed on VMware virtual machines, within a container engine on Linux hosts, or as dedicated hardware appliances. Each type of platform and each type of node has its own set of hardening best practices.

Control remote IPMI access to BMC

You can enable or disable remote IPMI access for all appliances containing a BMC. The remote IPMI interface allows low-level hardware access to your StorageGRID appliances by anyone with a BMC account and password. If you do not need remote IPMI access to the BMC, disable this option.

- To control remote IPMI access to the BMC in Grid Manager, go to CONFIGURATION > Security > Security settings > Appliances:
 - Clear the Enable remote IPMI access checkbox to disable IPMI access to the BMC.
 - Select the Enable remote IPMI access checkbox to enable IPMI access to the BMC.

Firewall configuration

As part of the system hardening process, you must review external firewall configurations and modify them so that traffic is accepted only from the IP addresses and on the ports from which it is strictly needed.

StorageGRID includes an internal firewall on each node that enhances the security of your grid by enabling you to control network access to the node. You should manage internal firewall controls to prevent network access on all ports except those necessary for your specific grid deployment. The configuration changes you make on the Firewall control page are deployed to each node.

Specifically, you can manage these areas:

- **Privileged addresses**: You can allow selected IP addresses or subnets to access ports that are closed by settings on the Manage external access tab.
- Manage external access: You can close ports that are open by default, or reopen ports previously closed.
- **Untrusted Client Network**: You can specify whether a node trusts inbound traffic from the Client Network as well as the additional ports you want open when untrusted Client Network is configured.

While this internal firewall provides an additional layer of protection against some common threats, it does not remove the need for an external firewall.

For a list of all internal and external ports used by StorageGRID, see Network port reference.

Disable unused services

For all StorageGRID nodes, you should disable or block access to unused services. For example, if you aren't planning to configure client access to the audit shares for NFS, block or disable access to these services.

Virtualization, containers, and shared hardware

For all StorageGRID nodes, avoid running StorageGRID on the same physical hardware as untrusted software. Don't assume that hypervisor protections will prevent malware from accessing StorageGRID-protected data if both StorageGRID and the malware exist on the same the physical hardware. For example, the Meltdown and Spectre attacks exploit critical vulnerabilities in modern processors and allow programs to steal data in memory on the same computer.

Protect nodes during installation

Don't allow untrusted users to access StorageGRID nodes over the network when the nodes are being installed. Nodes aren't fully secure until they have joined the grid.

Guidelines for Admin Nodes

Admin Nodes provide management services such as system configuration, monitoring, and logging. When you sign in to the Grid Manager or the Tenant Manager, you are connecting to an Admin Node.

Follow these guidelines to secure the Admin Nodes in your StorageGRID system:

- Secure all Admin Nodes from untrusted clients, such as those on the open internet. Ensure that no untrusted client can access any Admin Node on the Grid Network, the Admin Network, or the Client Network.
- StorageGRID Groups control access to Grid Manager and Tenant Manager features. Grant each Group of users the minimum required permissions for their role, and use the read-only access mode to prevent users from changing configuration.
- When using StorageGRID load balancer endpoints, use Gateway Nodes instead of Admin Nodes for untrusted client traffic.
- If you have untrusted tenants, don't allow them to have direct access to the Tenant Manager or the Tenant Management API. Instead, have any untrusted tenants use a tenant portal or an external tenant management system, which interacts with the Tenant Management API.
- Optionally, use an admin proxy for more control over AutoSupport communication from Admin Nodes to NetApp Support. See the steps for creating an admin proxy.
- Optionally, use the restricted 8443 and 9443 ports to separate Grid Manager and Tenant Manager communications. Block the shared port 443 and limit tenant requests to port 9443 for additional protection.
- Optionally, use separate Admin Nodes for grid administrators and tenant users.

For more information, see the instructions for administering StorageGRID.

Guidelines for Storage Nodes

Storage Nodes manage and store object data and metadata. Follow these guidelines to secure the Storage Nodes in your StorageGRID system.

- Don't allow untrusted clients to connect directly to Storage Nodes. Use a load balancer endpoint served by a Gateway Node or a third party load balancer.
- Don't enable outbound services for untrusted tenants. For example, when creating the account for an untrusted tenant, don't allow the tenant to use its own identity source and don't allow the use of platform services. See the steps for creating a tenant account.
- Use a third-party load balancer for untrusted client traffic. Third-party load balancing offers more control and additional layers of protection against attack.
- Optionally, use a storage proxy for more control over Cloud Storage Pools and platform services communication from Storage Nodes to external services. See the steps for creating a storage proxy.
- Optionally, connect to external services using the Client Network. Then, select CONFIGURATION > Security > Firewall control > Untrusted Client Networks and indicate that the Client Network on the Storage Node is untrusted. The Storage Node no longer accepts any incoming traffic on the Client Network, but it continues to allow outbound requests for Platform Services.

Guidelines for Gateway Nodes

Gateway Nodes provide an optional load-balancing interface that client applications can use to connect to StorageGRID. Follow these guidelines to secure any Gateway Nodes in your StorageGRID system:

- Configure and use load balancer endpoints. See Considerations for load balancing.
- Use a third-party load balancer between the client and the Gateway Node or Storage Nodes for untrusted client traffic. Third-party load balancing offers more control and additional layers of protection against attack. If you do use a third-party load balancer, network traffic can still optionally be configured to go through an internal load balancer endpoint or be sent directly to Storage Nodes.
- If you are using load balancer endpoints, optionally have clients connect over the Client Network. Then,

select **CONFIGURATION** > **Security** > **Firewall control** > **Untrusted Client Networks** and indicate that the Client Network on the Gateway Node is untrusted. The Gateway Node only accepts inbound traffic on the ports explicitly configured as load balancer endpoints.

Guidelines for hardware appliance nodes

StorageGRID hardware appliances are specially designed for use in a StorageGRID system. Some appliances can be used as Storage Nodes. Other appliances can be used as Admin Nodes or Gateway Nodes. You can combine appliance nodes with software-based nodes or deploy fully engineered, all-appliance grids.

Follow these guidelines to secure any hardware appliance nodes in your StorageGRID system:

- If the appliance uses SANtricity System Manager for storage controller management, prevent untrusted clients from accessing SANtricity System Manager over the network.
- If the appliance has a baseboard management controller (BMC), be aware that the BMC management port allows low-level hardware access. Connect the BMC management port only to a secure, trusted, internal management network. If no such network is available, leave the BMC management port unconnected or blocked, unless a BMC connection is requested by technical support.
- If the appliance supports remote management of the controller hardware over Ethernet using the Intelligent Platform Management Interface (IPMI) standard, block untrusted traffic on port 623.

You can enable or disable remote IPMI access for all appliances containing a BMC. The remote IPMI interface allows low-level hardware access to your StorageGRID appliances by anyone with a BMC account and password. If you do not need remote IPMI access to the BMC, disable this option using one of the following methods:

In Grid Manager, go to **CONFIGURATION > Security > Security settings > Appliances** and clear the **Enable remote IPMI access** checkbox.

In the Grid management API, use the private endpoint: PUT /private/bmc.

- For appliance models containing SED, FDE, or FIPS NL-SAS drives that you manage with SANtricity System Manager, enable and configure SANtricity Drive Security.
- For appliance models containing SED or FIPS NVMe SSDs that you manage using the StorageGRID Appliance Installer and Grid Manager, enable and configure StorageGRID drive encryption.
- For appliances without SED, FDE, or FIPS drives, enable and configure StorageGRID software node encryption using a Key Management Server (KMS).

Hardening guidelines for TLS and SSH

You should replace the default certificates created during installation and select the appropriate security policy for TLS and SSH connections.

Hardening guidelines for certificates

i

You should replace the default certificates created during installation with your own custom certificates.

For many organizations, the self-signed digital certificate for StorageGRID web access is not compliant with their information security policies. On production systems, you should install a CA-signed digital certificate for use in authenticating StorageGRID.

Specifically, you should use custom server certificates instead of these default certificates:

- **Management interface certificate**: Used to secure access to the Grid Manager, the Tenant Manager, the Grid Management API, and the Tenant Management API.
- S3 and Swift API certificate: Used to secure access to Storage Nodes and Gateway Nodes, which S3 and Swift client applications use to upload and download object data.

See Manage security certificates for details and instructions.



StorageGRID manages the certificates used for load balancer endpoints separately. To configure load balancer certificates, see Configure load balancer endpoints.

When using custom server certificates, follow these guidelines:

- Certificates should have a *subjectAltName* that matches DNS entries for StorageGRID. For details, see section 4.2.1.6, "Subject Alternative Name," in RFC 5280: PKIX Certificate and CRL Profile.
- When possible, avoid the use of wildcard certificates. An exception to this guideline is the certificate for an S3 virtual hosted style endpoint, which requires the use of a wildcard if bucket names aren't known in advance.
- When you must use wildcards in certificates, you should take additional steps to reduce the risks. Use a wildcard pattern such as *.s3.example.com, and don't use the s3.example.com suffix for other applications. This pattern also works with path-style S3 access, such as dc1-s1.s3.example.com/mybucket.
- Set the certificate expiration times to be short (for example, 2 months), and use the Grid Management API to automate certificate rotation. This especially important for wildcard certificates.

In addition, clients should use strict hostname checking when communicating with StorageGRID.

Hardening guidelines for TLS and SSH policy

You can select a security policy to determine which protocols and ciphers are used to establish secure TLS connections with client applications and secure SSH connections to internal StorageGRID services.

The security policy controls how TLS and SSH encrypt data in motion. As a best practice, you should disable encryption options that aren't required for application compatibility. Use the default Modern policy, unless your system needs to be Common Criteria-compliant or you need to use other ciphers.

See Manage the TLS and SSH policy for details and instructions.

Other hardening guidelines

In addition to following the hardening guidelines for StorageGRID networks and nodes, you should follow the hardening guidelines for other areas of the StorageGRID system.

Logs and audit messages

Always protect StorageGRID logs and audit message output in a secure manner. StorageGRID logs and audit messages provide invaluable information from a support and system availability standpoint. In addition, the information and details contained in StorageGRID logs and audit message output are generally of a sensitive nature.

Configure StorageGRID to send security events to an external syslog server. If using syslog export, select TLS and RELP/TLS for the transport protocols.

See the Log files reference for more information about StorageGRID logs. See Audit messages for more information about StorageGRID audit messages.

NetApp AutoSupport

The AutoSupport feature of StorageGRID allows you to proactively monitor the health of your system and automatically send packages to the NetApp Support Site, your organization's internal support team, or a support partner. By default, sending AutoSupport packages to NetApp is enabled when StorageGRID is configured for the first time.

The AutoSupport feature can be disabled. However, NetApp recommends enabling it because AutoSupport helps speed problem identification and resolution should an issue arise on your StorageGRID system.

AutoSupport supports HTTPS, HTTP, and SMTP for transport protocols. Because of the sensitive nature of AutoSupport packages, NetApp strongly recommends using HTTPS as the default transport protocol for sending AutoSupport packages to NetApp.

Cross-origin resource sharing (CORS)

You can configure cross-origin resource sharing (CORS) for an S3 bucket if you want that bucket and objects in that bucket to be accessible to web applications in other domains. In general, don't enable CORS unless it is required. If CORS is required, restrict it to trusted origins.

See the steps for configuring cross-origin resource sharing (CORS).

External security devices

A complete hardening solution must address security mechanisms outside of StorageGRID. Using additional infrastructure devices for filtering and limiting access to StorageGRID is an effective way to establish and maintain a stringent security posture. These external security devices include firewalls, intrusion prevention systems (IPSs), and other security devices.

A third-party load balancer is recommended for untrusted client traffic. Third-party load balancing offers more control and additional layers of protection against attack.

Ransomware mitigation

Help protect your object data from ransomware attacks by following the recommendations in Ransomware defense with StorageGRID.

Configure StorageGRID for FabricPool

Configure StorageGRID for FabricPool: Overview

If you use NetApp ONTAP software, you can use NetApp FabricPool to tier inactive data to a NetApp StorageGRID object storage system.

Use these instructions to:

- Learn the considerations and best practices for configuring StorageGRID for a FabricPool workload.
- Learn how to configure a StorageGRID object storage system for use with FabricPool.
- Learn how to provide the required values to ONTAP when attaching StorageGRID as a FabricPool cloud

tier.

Quick start for configuring StorageGRID for FabricPool



Plan your configuration

- Decide which FabricPool volume tiering policy you will use to tier inactive ONTAP data to StorageGRID.
- Plan and install a StorageGRID system to meet your storage capacity and performance needs.
- Become familiar with StorageGRID system software, including the Grid Manager and the Tenant Manager.
- Review the FabricPool best practices for HA groups, load balancing, ILM, and more.
- Review these additional resources, which provide details about using and configuring ONTAP and FabricPool:

TR-4598: FabricPool Best Practices in ONTAP

ONTAP 9: FabricPool tier management overview with System Manager



Perform prerequisite tasks

Obtain the information needed to attach StorageGRID as a cloud tier, including:

- · IP addresses
- Domain names
- SSL certificate

Optionally, configure identity federation and single sign-on.



Configure StorageGRID settings

Use StorageGRID to obtain the values ONTAP needs to connect to the grid.

Using the FabricPool setup wizard is the recommended and the fastest way to configure all items, but you can also configure each entity manually, if required.

1.
4
\sim

Configure ONTAP and DNS

Use ONTAP to add a cloud tier that uses the StorageGRID values. Then, configure DNS entries to associate IP addresses to any domain names you plan to use.



Monitor and manage

When your system is up and running, perform ongoing tasks in ONTAP and StorageGRID to manage and monitor FabricPool data tiering over time.

What is FabricPool?

FabricPool is an ONTAP hybrid storage solution that uses a high-performance flash aggregate as the

performance tier and an object store as the cloud tier. Using FabricPool-enabled aggregates helps you reduce storage cost without compromising performance, efficiency, or protection.

FabricPool associates a cloud tier (an external object store, such as StorageGRID) with a local tier (an ONTAP storage aggregate) to create a composite collection of discs. Volumes inside the FabricPool can then take advantage of the tiering by keeping active (hot) data on high-performance storage (the local tier) and tiering inactivate (cold) data to the external object store (the cloud tier).

No architectural changes are required, and you can continue managing your data and application environment from the central ONTAP storage system.

What is StorageGRID?

NetApp StorageGRID is a storage architecture that manages data as objects, as opposed to other storage architectures such as file or block storage. Objects are kept inside a single container (such as a bucket) and aren't nested as files inside a directory inside other directories. Although object storage generally provides lower performance than file or block storage, it is significantly more scalable. StorageGRID buckets can hold petabytes of data and billions of objects.

Why use StorageGRID as a FabricPool cloud tier?

FabricPool can tier ONTAP data to a number of object storage providers, including StorageGRID. Unlike public clouds that might set a maximum number of supported input/output operations per second (IOPS) at the bucket or container level, StorageGRID performance scales with the number of nodes in a system. Using StorageGRID as a FabricPool cloud tier allows you to keep your cold data in your own private cloud for highest performance and complete control over your data.

In addition, a FabricPool license is not required when you use StorageGRID as the cloud tier.

Information needed to attach StorageGRID as a cloud tier

Before you can attach StorageGRID as a cloud tier for FabricPool, you must perform configuration steps in StorageGRID and obtain certain values for use in ONTAP.

What values do I need?

The following table shows the values you must configure in StorageGRID and how those values are used by ONTAP and the DNS server.

Value	Where value is configured	Where value is used
Virtual IP (VIP) addresses	StorageGRID > HA group	DNS entry
Port	StorageGRID > Load balancer endpoint	ONTAP System Manager > Add Cloud Tier
SSL certificate	StorageGRID > Load balancer endpoint	ONTAP System Manager > Add Cloud Tier
Server name (FQDN)	StorageGRID > Load balancer endpoint	DNS entry

Value	Where value is configured	Where value is used
Access key ID and secret access key	StorageGRID > Tenant and bucket	ONTAP System Manager > Add Cloud Tier
Bucket/Container name	StorageGRID > Tenant and bucket	ONTAP System Manager > Add Cloud Tier

How do I get these values?

Depending on your requirements, you can do either of the following to obtain the information you need:

- Use the FabricPool setup wizard. The FabricPool setup wizard helps you to quickly configure the required values in StorageGRID and outputs a file that you can use to configure ONTAP System Manager. The wizard guides you through the required steps and helps to make sure your settings conform to StorageGRID and FabricPool best practices.
- Configure each item manually. Then, enter the values into ONTAP System Manager or the ONTAP CLI. Follow these steps:
 - 1. Configure a high availability (HA) group for FabricPool.
 - 2. Create a load balancer endpoint for FabricPool.
 - 3. Create a tenant account for FabricPool.
 - 4. Sign in to the tenant account, and create the bucket and access keys for the root user.
 - 5. Create an ILM rule for FabricPool data and add it to your active ILM policies. See Configure ILM for FabricPool data.
 - 6. Optionally, create a traffic classification policy for FabricPool.

Use FabricPool setup wizard

Use FabricPool setup wizard: Considerations and requirements

You can use the FabricPool setup wizard to configure StorageGRID as the object storage system for a FabricPool cloud tier. After you complete the setup wizard, you can enter the required details into ONTAP System Manager.

When to use the FabricPool setup wizard

The FabricPool setup wizard guides you through each step of configuring StorageGRID for use with FabricPool and automatically configures certain entities for you, such as the ILM and traffic classification policies. As part of completing the wizard, you download a file that you can use to enter values into ONTAP System Manager. Use the wizard to configure your system more quickly and to make sure your settings conform to StorageGRID and FabricPool best practices.

Assuming you have Root access permission, you can complete the FabricPool setup wizard when you start using the StorageGRID Grid Manager, or you can access and complete the wizard at any later time. Depending on your requirements, you can also configure some or all of the required items manually and then use the wizard to assemble the values that ONTAP needs into a single file.



Use the FabricPool setup wizard unless you know you have special requirements or your implementation will require significant customization.

Before using the wizard

Confirm you have completed these prerequisite steps.

Review best practices

- You have a general understanding of the information needed to attach StorageGRID as a cloud tier.
- You have reviewed the FabricPool best practices for:
 - High availability (HA) groups
 - Load balancing
 - ILM rules and policy

Obtain IP addresses and set up VLAN interfaces

If you will configure an HA group, you know which nodes ONTAP will connect to and which StorageGRID network will be used. You also know which values to enter for the subnet CIDR, gateway IP address, and virtual IP (VIP) addresses.

If you plan to use a virtual LAN to segregate FabricPool traffic, you have already configured the VLAN interface. See Configure VLAN interfaces.

Configure identity federation and SSO

If you plan to use identity federation or single sign-on (SSO) for your StorageGRID system, you have enabled these features. You also know which federated group should have root access for the tenant account that ONTAP will use. See Use identity federation and Configure single sign-on.

Obtain and configure domain names

- You know which fully qualified domain name (FQDN) to use for StorageGRID. Domain name server (DNS) entries will map this FQDN to the virtual IP (VIP) addresses of the HA group that you create using the wizard. See Configure DNS server.
- If you plan to use S3 virtual hosted-style requests, you have configured S3 endpoint domain names. ONTAP uses path-style URLs by default, but using virtual hosted-style requests is recommended.

Review load balancer and security certificate requirements

If you plan to use the StorageGRID load balancer, you have reviewed the general considerations for load balancing. You have the certificates you will upload or the values you need to generate a certificate.

If you plan to use an external (third-party) load balancer endpoint, you have the fully qualified domain name (FQDN), port, and certificate for that load balancer.

Confirm ILM storage pool configuration

if you initially installed StorageGRID 11.6 or earlier, you have configured the storage pool you will use. In general, you should create a storage pool for each StorageGRID site you will use to store ONTAP data.



This prerequisite does not apply if you initially installed StorageGRID 11.7 or 11.8. When you initially install either of these versions, storage pools are automatically created for each site.

Relationship between ONTAP and the StorageGRID cloud tier

The FabricPool wizard guides you through the process of creating a single StorageGRID cloud tier that includes one StorageGRID tenant, one set of access keys, and one StorageGRID bucket. You can attach this StorageGRID cloud tier to one or more ONTAP local tiers.

Attaching a single cloud tier to multiple local tiers in a cluster is the general best practice. However, depending on your requirements, you might want to use more than one bucket or even more than one StorageGRID tenant for the local tiers in a single cluster. Using different buckets and tenants allows you to isolate data and data access between ONTAP local tiers, but is somewhat more complex to configure and manage.

NetApp does not recommend attaching a single cloud tier to local tiers in multiple clusters.



For the best practices for using StorageGRID with NetApp MetroCluster™ and FabricPool Mirror, see TR-4598: FabricPool Best Practices in ONTAP.

Optional: Use a different bucket for each local tier

To use more than one bucket for the local tiers in an ONTAP cluster, add more than one StorageGRID cloud tier in ONTAP. Each cloud tier shares the same HA group, load balancer endpoint, tenant, and access keys, but uses a different container (StorageGRID bucket). Follow these general steps:

- 1. From StorageGRID Grid Manager, complete the FabricPool setup wizard for the first cloud tier.
- 2. From ONTAP System Manager, add a cloud tier and use the file you downloaded from StorageGRID to provide the required values.
- 3. From StorageGRID Tenant Manager, sign in to the tenant that was created by the wizard, and create a second bucket.
- 4. Complete the FabricPool wizard again. Select the existing HA group, load balancer endpoint, and tenant. Then, select the new bucket you created manually. Create a new ILM rule for the new bucket and activate an ILM policy to include that rule.
- 5. From ONTAP, add a second cloud tier but provide the new bucket name.

Optional: Use a different tenant and bucket for each local tier

To use more than one tenant and different sets of access keys for the local tiers in an ONTAP cluster, add more than one StorageGRID cloud tier in ONTAP. Each cloud tier shares the same HA group, load balancer endpoint, but uses a different tenant, access keys, and container (StorageGRID bucket). Follow these general steps:

- 1. From StorageGRID Grid Manager, complete the FabricPool setup wizard for the first cloud tier.
- 2. From ONTAP System Manager, add a cloud tier and use the file you downloaded from StorageGRID to provide the required values.
- Complete the FabricPool wizard again. Select the existing HA group and load balancer endpoint. Create a new tenant and bucket. Create a new ILM rule for the new bucket and activate an ILM policy to include that rule.
- 4. From ONTAP, add a second cloud tier but provide the new access key, secret key, and bucket name.

Access and complete the FabricPool setup wizard

You can use the FabricPool setup wizard to configure StorageGRID as the object storage system for a FabricPool cloud tier.

Before you begin

• You have reviewed the considerations and requirements for using the FabricPool setup wizard.



If you want to configure StorageGRID for use with any other S3 client application, go to Use S3 setup wizard.

• You have the Root access permission.

Access the wizard

You can complete the FabricPool setup wizard when you start using the StorageGRID Grid Manager, or you can access and complete the wizard at any later time.

Steps

- 1. Sign in to the Grid Manager using a supported web browser.
- 2. If the **FabricPool and S3 setup wizard** banner appears on the dashboard, select the link in the banner. If the banner no longer appears, select the help icon from the header bar in the Grid Manager and select **FabricPool and S3 setup wizard**.



3. In the FabricPool section of the FabricPool and S3 setup wizard page, select **Configure now**.

Step 1 of 9: Configure HA group appears.

Step 1 of 9: Configure HA group

A high availability (HA) group is a collection of nodes that each contain the StorageGRID Load Balancer service. An HA group can contain Gateway Nodes, Admin Nodes, or both.

You can use an HA group to help keep FabricPool data connections available. An HA group uses virtual IP addresses (VIPs) to provide highly available access to the Load Balancer service. If the active interface in the

HA group fails, a backup interface can manage the workload with little impact to FabricPool operations

For details about this task, see Manage high availability groups and Best practices for high availability groups.

- 1. If you plan to use an external load balancer, you don't need to create an HA group. Select **Skip this step** and go to Step 2 of 9: Configure load balancer endpoint.
- 2. To use the StorageGRID load balancer, create a new HA group or use an existing HA group.

Create HA group

- a. To create a new HA group, select **Create HA group**.
- b. For the Enter details step, complete the following fields.

Field	Description
HA group name	A unique display name for this HA group.
Description (optional)	The description of this HA group.

c. For the Add interfaces step, select the node interfaces you want to use in this HA group.

Use the column headers to sort the rows, or enter a search term to locate interfaces more quickly.

You can select one or more nodes, but you can select only one interface for each node.

d. For the **Prioritize interfaces** step, determine the Primary interface and any backup interfaces for this HA group.

Drag rows to change the values in the **Priority order** column.

The first interface in the list is the Primary interface. The Primary interface is the active interface unless a failure occurs.

If the HA group includes more than one interface and the active interface fails, the virtual IP (VIP) addresses move to the first backup interface in the priority order. If that interface fails, the VIP addresses move to the next backup interface, and so on. When failures are resolved, the VIP addresses move back to highest priority interface available.

e. For the Enter IP addresses step, complete the following fields.

Field	Description
Subnet CIDR	The address of the VIP subnet in CIDR notation—an IPv4 address followed by a slash and the subnet length (0-32). The network address must not have any host bits set. For example, 192.16.0.0/22.
Gateway IP address (optional)	Optional. If the ONTAP IP addresses used to access StorageGRID aren't on the same subnet as the StorageGRID VIP addresses, enter the StorageGRID VIP local gateway IP address. The local gateway IP address must be within the VIP subnet.
Virtual IP address	Enter at least one and no more than ten VIP addresses for the active interface in the HA group. All VIP addresses must be within the VIP subnet and all will be active at the same time on the active interface. At least one address must be IPv4. Optionally, you can specify additional IPv4 and IPv6 addresses.

- f. Select Create HA group and then select Finish to return to the FabricPool setup wizard.
- g. Select **Continue** to go to the load balancer step.

Use existing HA group

- a. To use an existing HA group, select the HA group name from the **Select an HA group** drop-down list.
- b. Select **Continue** to go to the load balancer step.

Step 2 of 9: Configure load balancer endpoint

StorageGRID uses a load balancer to manage the workload from client applications, such as FabricPool. Load balancing maximizes speed and connection capacity across multiple Storage Nodes.

You can use the StorageGRID Load Balancer service, which exists on all Gateway and Admin Nodes, or you can connect to an external (third-party) load balancer. Using the StorageGRID load balancer is recommended.

For details about this task, see the general considerations for load balancing and the best practices for load balancing for FabricPool.

Steps

1. Select or create a StorageGRID load balancer endpoint or use an external load balancer.

Create endpoint

- a. Select Create endpoint.
- b. For the **Enter endpoint details** step, complete the following fields.

Field	Description
Name	A descriptive name for the endpoint.
Port	The StorageGRID port you want to use for load balancing. This field defaults to 10433 for the first endpoint you create, but you can enter any unused external port. If you enter 80 or 443, the endpoint is configured only on Gateway Nodes, because these ports are reserved on Admin Nodes. Note: Ports used by other grid services aren't permitted. See the Network port reference.
Client type	Must be S3 .
Network protocol	Select HTTPS . Note : Communicating with StorageGRID without TLS encryption is supported but not recommended.

c. For the **Select binding mode** step, specify the binding mode. The binding mode controls how the endpoint is accessed using any IP address or using specific IP addresses and network interfaces.

Mode	Description
Global (default)	Clients can access the endpoint using the IP address of any Gateway Node or Admin Node, the virtual IP (VIP) address of any HA group on any network, or a corresponding FQDN. Use the Global setting (default) unless you need to restrict the accessibility of this endpoint.
Virtual IPs of HA groups	Clients must use a virtual IP address (or corresponding FQDN) of an HA group to access this endpoint. Endpoints with this binding mode can all use the same port number, as long as the HA groups you select for the endpoints don't overlap.
Node interfaces	Clients must use the IP addresses (or corresponding FQDNs) of selected node interfaces to access this endpoint.
Node type	Based on the type of node you select, clients must use either the IP address (or corresponding FQDN) of any Admin Node or the IP address (or corresponding FQDN) of any Gateway Node to access this endpoint.

d. For the Tenant access step, select one of the following:

Field	Description
Allow all tenants (default)	All tenant accounts can use this endpoint to access their buckets. Allow all tenants is almost always the appropriate option for the load balancer endpoint used for FabricPool. You must select this option if you are using the FabricPool setup wizard for a new StorageGRID system and you have not yet created any tenant accounts.
Allow selected tenants	Only the selected tenant accounts can use this endpoint to access their buckets.
Block selected tenants	The selected tenant accounts can't use this endpoint to access their buckets. All other tenants can use this endpoint.

e. For the Attach certificate step, select one of the following:

Field	Description
Upload certificate (recommended)	Use this option to upload a CA-signed server certificate, certificate private key, and optional CA bundle.
Generate certificate	Use this option to generate a self-signed certificate. See Configure load balancer endpoints for details of what to enter.
Use StorageGRID S3 and Swift certificate	This option is available only if you have already uploaded or generated a custom version of the StorageGRID global certificate. See Configure S3 and Swift API certificates for details.

- f. Select **Finish** to return to the FabricPool setup wizard.
- g. Select **Continue** to go to the tenant and bucket step.



Changes to an endpoint certificate can take up to 15 minutes to be applied to all nodes.

Use existing load balancer endpoint

- a. Select the name of an existing endpoint from the **Select a load balancer endpoint** drop-down list.
- b. Select **Continue** to go to the tenant and bucket step.

Use external load balancer

a. Complete the following fields for the external load balancer.

Field	Description
FQDN	The fully qualified domain name (FQDN) of the external load balancer.
Port	The port number that FabricPool will use to connect to the external load balancer.
Certificate	Copy the server certificate for the external load balancer and paste it into this field.

b. Select **Continue** to go to the tenant and bucket step.

Step 3 of 9: Tenant and bucket

A tenant is an entity that can use S3 applications to store and retrieve objects in StorageGRID. Each tenant has its own users, access keys, buckets, objects, and a specific set of capabilities. You must create a StorageGRID tenant before you can create the bucket that FabricPool will use.

A bucket is a container used to store a tenant's objects and object metadata. Although some tenants might have many buckets, the wizard lets you create or select only one tenant and one bucket at a time. You can use the Tenant Manager later to add any additional buckets you need.

You can create a new tenant and bucket for FabricPool use, or you can select an existing tenant and bucket. If you create a new tenant, the system automatically creates the access key ID and secret access key for the tenant's root user.

For details about this task, see Create a tenant account for FabricPool and Create an S3 bucket and obtain an access key.

Steps

Create a new tenant and bucket or select an existing tenant.

New tenant and bucket

- 1. To create a new tenant and bucket, enter a Tenant name. For example, FabricPool tenant.
- 2. Define root access for the tenant account, based on whether your StorageGRID system uses identity federation, single sign-on (SSO), or both.

Option	Do this
If identity federation is not enabled	Specify the password to use when signing into the tenant as the local root user.
If identity federation is enabled	 Select an existing federated group to have Root access permission for the tenant. Optionally, specify the password to use when signing in to the tenant as the local root user.
If both identity federation and single sign-on (SSO) are enabled	Select an existing federated group to have Root access permission for the tenant. No local users can sign in.

3. For **Bucket name**, enter the name of the bucket FabricPool will use to store ONTAP data. For example, fabricpool-bucket.



You can't change the bucket name after creating the bucket.

4. Select the **Region** for this bucket.

Use the default region (us-east-1) unless you expect to use ILM in the future to filter objects based on the bucket's region.

5. Select Create and Continue to create the tenant and bucket and to go to the download data step

Select tenant and bucket

The existing tenant account must have at least one bucket that does not have versioning enabled. You can't select an existing tenant account if no bucket exists for that tenant.

- 1. Select the existing tenant from the Tenant name drop-down list.
- 2. Select the existing bucket from the Bucket name drop-down list.

FabricPool does not support object versioning, so buckets that have versioning enabled aren't shown.



Don't select a bucket that has S3 Object Lock enabled for use with FabricPool.

3. Select **Continue** to go to the download data step.

Step 4 of 9: Download ONTAP settings

During this step, you download a file that you can use to enter values into ONTAP System Manager.

Steps

1. Optionally, select the copy icon (()) to copy both the access key ID and secret access key to the clipboard.

These values are included in the download file, but you might want to save them separately.

2. Select Download ONTAP settings to download a text file that contains the values you've entered so far.

The ONTAP_FabricPool_settings_bucketname.txt file includes the information you need to configure StorageGRID as the object storage system for a FabricPool cloud tier, including:

- Load balancer connection details, including the server name (FQDN), port, and certificate
- Bucket name
- · Access key ID and secret access key for the root user of the tenant account
- 3. Save the copied keys and downloaded file to a secure location.



Don't close this page until you have copied both access keys, downloaded the ONTAP settings, or both. The keys will not be available after you close this page. Make sure to save this information in a secure location because it can be used to obtain data from your StorageGRID system.

- 4. Select the checkbox to confirm you have downloaded or copied the access key ID and secret access key.
- 5. Select Continue to go to the ILM storage pool step.

Step 5 of 9: Select a storage pool

A storage pool is a group of Storage Nodes. When you select a storage pool, you determine which nodes StorageGRID will use to store the data tiered from ONTAP.

For details about this step, see Create a storage pool.

Steps

- 1. From the Site drop-down list, select the StorageGRID site you want to use for the data tiered from ONTAP.
- 2. From the **Storage pool** drop-down list, select the storage pool for that site.

The storage pool for a site includes all Storage Nodes at that site.

3. Select Continue to go to the ILM rule step.

Step 6 of 9: Review ILM rule for FabricPool

Information lifecycle management (ILM) rules control the placement, duration, and ingest behavior for all objects in your StorageGRID system.

The FabricPool setup wizard automatically creates the recommended ILM rule for FabricPool use. This rule applies only to the bucket you specified. It uses 2+1 erasure coding at a single site to store the data that is tiered from ONTAP.

For details about this step, see Create ILM rule and Best practices for using ILM with FabricPool data.

1. Review the rule details.

Field	Description
Rule name	Automatically generated and can't be changed
Description	Automatically generated and can't be changed
Filter	The bucket name This rule only applies to objects that are saved in the bucket you specified.
Reference time	Ingest time The placement instruction starts when objects are initially saved to the bucket.
Placement instruction	Use 2+1 erasure coding

- 2. Sort the retention diagram by **Time period** and **Storage pool** to confirm the placement instruction.
 - The **Time period** for the rule is **Day 0 forever**. **Day 0** means that the rule is applied when data is tiered from ONTAP. Forever means that StorageGRID ILM will not delete data that has been tiered from ONTAP.
 - The **Storage pool** for the rule is the storage pool you selected. **EC 2+1** means the data will stored using 2+1 erasure coding. Each object will be saved as two data fragments and one parity fragment. The three fragments for each object will be saved to different Storage Nodes at a single site.
- 3. Select Create and Continue to create this rule and to go to the ILM policy step.

Step 7 of 9: Review and activate ILM policy

After the FabricPool setup wizard creates the ILM rule for FabricPool use, it creates an ILM policy. You must carefully simulate and review this policy before activating it.

For details about this step, see Create ILM policy and Best practices for using ILM with FabricPool data.



When you activate a new ILM policy, StorageGRID uses that policy to manage the placement, duration, and data protection of all objects in the grid, including existing objects and newly ingested objects. In some cases, activating a new policy can cause existing objects to be moved to new locations.



To avoid data loss, do not use an ILM rule that will expire or delete FabricPool cloud tier data. Set the retention period to **forever** to ensure that FabricPool objects aren't deleted by StorageGRID ILM.

- 1. Optionally, update the system-generated **Policy name**. By default, the system appends "+ FabricPool" to the name of your active or inactive policy, but you can provide your own name.
- 2. Review the list of rules in the inactive policy.
 - If your grid doesn't have an inactive ILM policy, the wizard creates an inactive policy by cloning your active policy and adding the new rule to the top.

- If your grid already has an inactive ILM policy and that policy uses the same rules and same order as the active ILM policy, the wizard adds the new rule to the top of the inactive policy.
- If your inactive policy contains different rules or a different order than the active policy, the wizard creates a new inactive policy by cloning your active policy and adding the new rule to the top.
- 3. Review the order of the rules in the new inactive policy.

Because the FabricPool rule is the first rule, any objects in the FabricPool bucket are placed before the other rules in the policy are evaluated. Objects in any other buckets are placed by subsequent rules in the policy.

- 4. Review the retention diagram to learn how different objects will be retained.
 - a. Select Expand all to see a retention diagram for each rule in the inactive policy.
 - b. Select **Time period** and **Storage pool** to review the retention diagram. Confirm that any rules that apply to the FabricPool bucket or tenant retain objects **forever**.
- 5. When you have reviewed the inactive policy, select **Activate and continue** to activate the policy and go to the traffic classification step.



Errors in an ILM policy can cause irreparable data loss. Review the policy carefully before activating.

Step 8 of 9: Create traffic classification policy

As an option, the FabricPool setup wizard can create a traffic classification policy that you can use to monitor the FabricPool workload. The system-created policy uses a matching rule to identify all network traffic related to the bucket you created. This policy monitors traffic only; it does not limit traffic for FabricPool or any other clients.

For details about this step, see Create a traffic classification policy for FabricPool.

Steps

- 1. Review the policy.
- 2. If you want to create this traffic classification policy, select Create and continue.

As soon as FabricPool begins tiering data to StorageGRID, you can go to the Traffic Classification Policies page to view network traffic metrics for this policy. Later, you can also add rules to limit other workloads and ensure that the FabricPool workload has most of the bandwidth.

3. Otherwise, select Skip this step.

Step 9 of 9: Review summary

The summary provides details about the items you configured, including the name of the load balancer, tenant, and bucket, the traffic classification policy, and the active ILM policy,

- 1. Review the summary.
- 2. Select Finish.

Next steps

After completing the FabricPool wizard, perform these additional steps.

Steps

- 1. Go to Configure ONTAP System Manager to enter the saved values and to complete the ONTAP side of the connection. You must add StorageGRID as a cloud tier, attach the cloud tier to a local tier to create a FabricPool, and set volume tiering policies.
- 2. Go to Configure the DNS server and ensure that the DNS includes a record to associate the StorageGRID server name (fully qualified domain name) to each StorageGRID IP address you will use.
- 3. Go to Other best practices for StorageGRID and FabricPool to learn the best practices for StorageGRID audit logs and other global configuration options.

Configure StorageGRID manually

Create a high availability (HA) group for FabricPool

When configuring StorageGRID for use with FabricPool, you can optionally create one or more high availability (HA) groups.

An HA group is a collection of nodes that each contain the StorageGRID Load Balancer service. An HA group can contain Gateway Nodes, Admin Nodes, or both.

You can use an HA group to help keep FabricPool data connections available. An HA group uses virtual IP addresses (VIPs) to provide highly available access to the Load Balancer service. If the active interface in the HA group fails, a backup interface can manage the workload with little impact to FabricPool operations.

For details about this task, see Manage high availability groups. To use the FabricPool setup wizard to complete this task, go to Access and complete the FabricPool setup wizard.

Before you begin

- You have reviewed the best practices for high availability groups.
- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.
- If you plan to use a VLAN, you have created the VLAN interface. See Configure VLAN interfaces.

Steps

- 1. Select CONFIGURATION > Network > High availability groups.
- 2. Select Create.
- 3. For the Enter details step, complete the following fields.

Field	Description
HA group name	A unique display name for this HA group.
Description (optional)	The description of this HA group.

4. For the Add interfaces step, select the node interfaces you want to use in this HA group.

Use the column headers to sort the rows, or enter a search term to locate interfaces more quickly.

You can select one or more nodes, but you can select only one interface for each node.

5. For the **Prioritize interfaces** step, determine the Primary interface and any backup interfaces for this HA group.

Drag rows to change the values in the **Priority order** column.

The first interface in the list is the Primary interface. The Primary interface is the active interface unless a failure occurs.

If the HA group includes more than one interface and the active interface fails, the virtual IP (VIP) addresses move to the first backup interface in the priority order. If that interface fails, the VIP addresses move to the next backup interface, and so on. When failures are resolved, the VIP addresses move back to highest priority interface available.

6. For the Enter IP addresses step, complete the following fields.

Field	Description
Subnet CIDR	The address of the VIP subnet in CIDR notation—an IPv4 address followed by a slash and the subnet length (0-32).
	The network address must not have any host bits set. For example, 192.16.0.0/22.
Gateway IP address (optional)	Optional. If the ONTAP IP addresses used to access StorageGRID aren't on the same subnet as the StorageGRID VIP addresses, enter the StorageGRID VIP local gateway IP address. The local gateway IP address must be within the VIP subnet.
Virtual IP address	Enter at least one and no more than ten VIP addresses for the active interface in the HA group. All VIP addresses must be within the VIP subnet.
	At least one address must be IPv4. Optionally, you can specify additional IPv4 and IPv6 addresses.

7. Select Create HA group and then select Finish.

Create a load balancer endpoint for FabricPool

StorageGRID uses a load balancer to manage the workload from client applications, such as FabricPool. Load balancing maximizes speed and connection capacity across multiple Storage Nodes.

When configuring StorageGRID for use with FabricPool, you must configure a load balancer endpoint and upload or generate a load balancer endpoint certificate, which is used to secure the connection between ONTAP and StorageGRID.

To use the FabricPool setup wizard to complete this task, go to Access and complete the FabricPool setup wizard.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.
- You have reviewed the general considerations for load balancing as well as the best practices for load balancing for FabricPool.

Steps

- 1. Select CONFIGURATION > Network > Load balancer endpoints.
- 2. Select Create.
- 3. For the Enter endpoint details step, complete the following fields.

Field	Description
Name	A descriptive name for the endpoint.
Port	The StorageGRID port you want to use for load balancing. This field defaults to 10433 for the first endpoint you create, but you can enter any unused external port. If you enter 80 or 443, the endpoint is configured only on Gateway Nodes. These ports are reserved on Admin Nodes. Note: Ports used by other grid services aren't permitted. See the Network port reference. You will provide this number to ONTAP when you attach StorageGRID as a FabricPool cloud tier.
Client type	Select S3 .
Network protocol	Select HTTPS . Note : Communicating with StorageGRID without TLS encryption is supported but not recommended.

4. For the **Select binding mode** step, specify the binding mode. The binding mode controls how the endpoint is accessed using any IP address or using specific IP addresses and network interfaces.

Mode	Description
Global (default)	Clients can access the endpoint using the IP address of any Gateway Node or Admin Node, the virtual IP (VIP) address of any HA group on any network, or a corresponding FQDN.
	Use the Global setting (default) unless you need to restrict the accessibility of this endpoint.

Mode	Description
Virtual IPs of HA groups	Clients must use a virtual IP address (or corresponding FQDN) of an HA group to access this endpoint.
	Endpoints with this binding mode can all use the same port number, as long as the HA groups you select for the endpoints don't overlap.
Node interfaces	Clients must use the IP addresses (or corresponding FQDNs) of selected node interfaces to access this endpoint.
Node type	Based on the type of node you select, clients must use either the IP address (or corresponding FQDN) of any Admin Node or the IP address (or corresponding FQDN) of any Gateway Node to access this endpoint.

5. For the **Tenant access** step, select one of the following:

Field	Description
Allow all tenants (default)	All tenant accounts can use this endpoint to access their buckets. Allow all tenants is almost always the appropriate option for the load balancer endpoint used for FabricPool. You must select this option if you have not yet created any tenant accounts.
Allow selected tenants	Only the selected tenant accounts can use this endpoint to access their buckets.
Block selected tenants	The selected tenant accounts can't use this endpoint to access their buckets. All other tenants can use this endpoint.

6. For the **Attach certificate** step, select one of the following:

Field	Description
Upload certificate (recommended)	Use this option to upload a CA-signed server certificate, certificate private key, and optional CA bundle.
Generate certificate	Use this option to generate a self-signed certificate. See Configure load balancer endpoints for details of what to enter.
Use StorageGRID S3 and Swift certificate	This option is available only if you have already uploaded or generated a custom version of the StorageGRID global certificate. See Configure S3 and Swift API certificates for details.

7. Select Create.



Changes to an endpoint certificate can take up to 15 minutes to be applied to all nodes.

Create a tenant account for FabricPool

You must create a tenant account in the Grid Manager for FabricPool use.

Tenant accounts allow client applications to store and retrieve objects on StorageGRID. Each tenant account has its own account ID, authorized groups and users, buckets, and objects.

For details about this task, see Create tenant account. To use the FabricPool setup wizard to complete this task, go to Access and complete the FabricPool setup wizard.

Before you begin

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

Steps

- 1. Select TENANTS.
- 2. Select Create.
- 3. For the Enter details steps, enter the following information.

Field	Description
Name	A name for the tenant account. Tenant names don't need to be unique. When the tenant account is created, it receives a unique, numeric account ID.
Description (optional)	A description to help identify the tenant.
Client type	Must be S3 for FabricPool.
Storage quota (optional)	Leave this field blank for FabricPool.

4. For the Select permissions step:

a. Don't select Allow platform services.

FabricPool tenants don't typically need to use platform services, such as CloudMirror replication.

- b. Optionally, select Use own identity source.
- c. Don't select Allow S3 Select.

FabricPool tenants don't typically need to use S3 Select.

- d. Optionally, select Use grid federation connection to allow the tenant to use a grid federation connection for account clone and cross-grid replication. Then, select the grid federation connection to use.
- For the Define root access step, specify which user will have the initial Root access permission for the tenant account, based on whether your StorageGRID system uses identity federation, single sign-on (SSO), or both.
| Option | Do this |
|--|---|
| If identity federation is not enabled | Specify the password to use when signing into the tenant as the local root user. |
| If identity federation is enabled | Select an existing federated group to have Root access
permission for the tenant. |
| | Optionally, specify the password to use when signing in to the
tenant as the local root user. |
| If both identity federation and single sign-on (SSO) are enabled | Select an existing federated group to have Root access permission
for the tenant. No local users can sign in. |

6. Select Create tenant.

Create an S3 bucket and obtain access keys

Before using StorageGRID with a FabricPool workload, you must create an S3 bucket for your FabricPool data. You also need to obtain an access key and secret access key for the tenant account you will use for FabricPool.

For details about this task, see Create S3 bucket and Create your own S3 access keys. To use the FabricPool setup wizard to complete this task, go to Access and complete the FabricPool setup wizard.

Before you begin

- You have created a tenant account for FabricPool use.
- You have Root access to the tenant account.

Steps

1. Sign in to the Tenant Manager.

You can do either of the following:

- From the Tenant Accounts page in the Grid Manager, select the **Sign in** link for the tenant, and enter your credentials.
- Enter the URL for the tenant account in a web browser, and enter your credentials.
- 2. Create an S3 bucket for FabricPool data.

You must create a unique bucket for each ONTAP cluster you plan to use.

- a. Select View buckets from the dashboard, or select STORAGE (S3) > Buckets.
- b. Select Create bucket.
- c. Enter the name of the StorageGRID bucket you want to use with FabricPool. For example, fabricpool-bucket.



You can't change the bucket name after creating the bucket.

d. Select the region for this bucket.

By default, all buckets are created in the us-east-1 region.

- e. Select Continue.
- f. Select Create bucket.



Don't select **Enable object versioning** for the FabricPool bucket. Similarly, don't edit a FabricPool bucket to use **Available** or a non-default consistency. The recommended bucket consistency for FabricPool buckets is **Read-after-new-write**, which is the default consistency for a new bucket.

- 3. Create an access key and a secret access key.
 - a. Select STORAGE (S3) > My access keys.
 - b. Select Create key.
 - c. Select Create access key.
 - d. Copy the access key ID and the secret access key to a safe location, or select **Download .csv** to save a spreadsheet file containing the access key ID and secret access key.

You will enter these values in ONTAP when you configure StorageGRID as a FabricPool cloud tier.



If you generate a new access key and secret access key in StorageGRID in the future, enter the new keys into ONTAP before deleting the old values from StorageGRID. Otherwise, ONTAP might temporarily lose its access to StorageGRID.

Configure ILM for FabricPool data

You can use this simple example policy as a starting point for your own ILM rules and policy.

This example assumes you are designing the ILM rules and an ILM policy for a StorageGRID system that has four Storage Nodes at a single data center in Denver, Colorado. The FabricPool data in this example uses a bucket named fabricpool-bucket.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss. To learn more, see Manage objects with ILM.



To avoid data loss, do not use an ILM rule that will expire or delete FabricPool cloud tier data. Set the retention period to **forever** to ensure that FabricPool objects aren't deleted by StorageGRID ILM.

Before you begin

- You have reviewed the best practices for using ILM with FabricPool data.
- You are signed in to the Grid Manager using a supported web browser.
- You have the ILM or Root access permission.
- If you upgraded to StorageGRID 11.8 from a previous StorageGRID version, you have configured the storage pool you will use. In general, you should create a storage pool for each StorageGRID site you will use to store data.



This prerequisite does not apply if you initially installed StorageGRID 11.7 or 11.8. When you initially install either of these versions, storage pools are automatically created for each site.

Steps

1. Create an ILM rule that applies only to the data in fabricpool-bucket. This example rule creates erasure-coded copies.

Rule definition	Example value
Rule name	2 + 1 erasure coding for FabricPool data
Bucket name	fabricpool-bucket You could also filter on the FabricPool tenant account.
Advanced filters	Object size greater than 0.2 MB. Note: FabricPool only writes 4 MB objects, but you must add an Object size filter because this rule uses erasure coding.
Reference time	Ingest time
Time period and placements	From Day 0 store forever Store objects by erasure coding using 2+1 EC scheme at Denver and retain those objects in StorageGRID forever. Image: Comparison of the expire o
Ingest behavior	Balanced

2. Create a default ILM rule that will create two replicated copies of any objects not matched by the first rule. Don't select a basic filter (tenant account or bucket name) or any advanced filters.

Rule definition	Example value
Rule name	Two replicated copies
Bucket name	none
Advanced filters	none
Reference time	Ingest time

Rule definition	Example value
Time period and placements	From Day 0 store forever Store objects by replicating 2 copies at Denver.
Ingest behavior	Balanced

- 3. Create an ILM policy and select the two rules. Because the replication rule does not use any filters, it can be the default (last) rule for the policy.
- 4. Ingest test objects into the grid.
- 5. Simulate the policy with the test objects to verify the behavior.
- 6. Activate the policy.

When this policy is activated, StorageGRID places object data as follows:

- The data tiered from FabricPool in fabricpool-bucket will be erasure-coded using the 2+1 erasurecoding scheme. Two data fragments and one parity fragment will be placed on three different Storage Nodes.
- All objects in all other buckets will be replicated. Two copies will be created and placed on two different Storage Nodes.
- The copies will be maintained in StorageGRID forever. StorageGRID ILM won't delete these objects.

Create a traffic classification policy for FabricPool

You can optionally design a StorageGRID traffic classification policy to optimize quality of service for the FabricPool workload.

For details about this task, see Manage traffic classification policies. To use the FabricPool setup wizard to complete this task, go to Access and complete the FabricPool setup wizard.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.

About this task

The best practices for creating a traffic classification policy for FabricPool depend on the workload, as follows:

• If you plan to tier FabricPool primary workload data to StorageGRID, you should ensure that the FabricPool workload has most of the bandwidth. You can create a traffic classification policy to limit all other workloads.



In general, FabricPool read operations are more important to prioritize than write operations.

For example, if other S3 clients use this StorageGRID system, you should create a traffic classification policy. You can limit network traffic for the other buckets, tenants, IP subnets, or load balancer endpoints.

*Generally, you should not impose quality of service limits on any FabricPool workload; you should only limit the other workloads.

• The limits placed on other workloads should account for the behavior of those workloads. The limits imposed will also vary based on the sizing and capabilities of your grid and what the expected amount of utilization is.

Steps

- 1. Select CONFIGURATION > Network > Traffic classification.
- 2. Select Create.
- 3. Enter a name and a description (optional) for the policy and select **Continue**.
- 4. For the Add matching rules step, add at least one rule.
 - a. Select Add rule
 - b. For Type, select **Load balancer endpoint**, and select the load balancer endpoint you created for FabricPool.

You can also select the FabricPool tenant account or bucket.

- c. If you want this traffic policy to limit traffic for the other endpoints, select Inverse match.
- 5. Optionally, add one or more limits to control the network traffic matched by the rule.



StorageGRID collects metrics even if you don't add any limits, so you can understand traffic trends.

- a. Select Add a limit.
- b. Select the type of traffic you want to limit and the limit to apply.
- 6. Select Continue.
- 7. Read and review the Traffic classification policy. Use the **Previous** button to go back and make changes as required. When you are satisfied with the policy, select **Save and continue**.

After your finish

View network traffic metrics to verify that the polices are enforcing the traffic limits you expect.

Configure ONTAP System Manager

After you have obtained the required StorageGRID information, you can go to ONTAP to add StorageGRID as a cloud tier.

Before you begin

- If you completed the FabricPool setup wizard, you have the ONTAP_FabricPool_settings_bucketname.txt file you downloaded.
- If you configured StorageGRID manually, you have the fully qualified domain name (FQDN) you are using for StorageGRID or the virtual IP (VIP) address for the StorageGRID HA group, the port number for the load balancer endpoint, the load balancer certificate, the access key ID and secret key for the root user of the tenant account, and the name of the bucket ONTAP will use in that tenant.

Access ONTAP System Manager

These instructions describe how to use ONTAP System Manager to add StorageGRID as a cloud tier. You can complete the same configuration using the ONTAP CLI. For instructions, go to ONTAP 9: FabricPool tier management with the CLI.

Steps

- 1. Access System Manager for the ONTAP cluster you want to tier to StorageGRID.
- 2. Sign in as an administrator for the cluster.
- 3. Navigate to **STORAGE** > **Tiers** > **Add Cloud Tier**.
- 4. Select StorageGRID from the list of object store providers.

Enter StorageGRID values

See ONTAP 9: FabricPool tier management overview with System Manager for more information.

Steps

1. Complete the Add Cloud Tier form, using the ONTAP_FabricPool_settings_bucketname.txt file or the values you obtained manually.

Field	Description
Name	Enter a unique name for this cloud tier. You can accept the default value.
URL style	If you configured S3 endpoint domain names, select Virtual Hosted-Style URL. Path-Style URL is the default for ONTAP, but using virtual hosted-style requests is recommended for StorageGRID. You must use Path-Style URL if you provide an IP address instead of a domain name for the Server name (FQDN) field.
Server name (FQDN)	 Enter the fully qualified domain name (FQDN) you are using for StorageGRID or the virtual IP (VIP) address for the StorageGRID HA group. For example, s3.storagegrid.company.com. Note the following: The IP address or domain name that you specify here must match the certificate you uploaded or generated for the StorageGRID load balancer endpoint. If you provide a domain name, the DNS record must map to each IP address you will use to connect to StorageGRID. See Configure the DNS server.
SSL	Enabled (default).
Object store certificate	Paste the certificate PEM you are using for the StorageGRID load balancer endpoint, including: BEGIN CERTIFICATE andEND CERTIFICATE Note: If an intermediate CA issued the StorageGRID certificate, you must provide the intermediate CA certificate. If the StorageGRID certificate was issued directly by the Root CA, you must provide the Root CA certificate.

Field	Description
Port	Enter the port used by the StorageGRID load balancer endpoint. ONTAP will use this port when it connects to StorageGRID. For example, 10433.
Access key and secret key	Enter the access key ID and secret access key for the root user of the StorageGRID tenant account. Tip : If you generate a new access key and secret access key in StorageGRID in the future, enter the new keys into ONTAP before deleting the old values from StorageGRID. Otherwise, ONTAP might temporarily lose its access to StorageGRID.
Container name	Enter the name of the StorageGRID bucket you created for use with this ONTAP tier.

- 2. Complete the final FabricPool configuration in ONTAP.
 - a. Attach one or more aggregates to the cloud tier.
 - b. Optionally, create a volume tiering policy.

Configure the DNS server

After configuring high availability groups, load balancer endpoints, and S3 endpoint domain names, you must ensure that the DNS includes the necessary entries for StorageGRID. You must include a DNS entry for each name in the security certificate and for each IP address you might use.

See Considerations for load balancing.

DNS entries for StorageGRID server name

Add DNS entries to associate the StorageGRID server name (fully qualified domain name) to each StorageGRID IP address you will use.

The IP addresses you enter in the DNS depend on whether you are using an HA group of load-balancing nodes:

- If you have configured an HA group, ONTAP will connect to the virtual IP addresses of that HA group.
- If you aren't using an HA group, ONTAP can connect to the StorageGRID Load Balancer service using the IP address of any Gateway Node or Admin Node.
- If the server name resolves to more than one IP address, ONTAP establishes client connections with all IP addresses (up to a maximum of 16 IP addresses). The IP addresses are picked up in a round-robin method when connections are established.

DNS entries for virtual hosted-style requests

If you have defined S3 endpoint domain names and you will use virtual hosted-style requests, add DNS entries for all required S3 endpoint domain names, including any wildcard names.

StorageGRID best practices for FabricPool

Best practices for high availability (HA) groups

Before attaching StorageGRID as a FabricPool cloud tier, learn about StorageGRID high availability (HA) groups and review the best practices for using HA groups with FabricPool.

What is an HA group?

A high availability (HA) group is a collection of interfaces from multiple StorageGRID Gateway Nodes, Admin Nodes, or both. An HA group helps to keep client data connections available. If the active interface in the HA group fails, a backup interface can manage the workload with little impact on FabricPool operations.

Each HA group provides highly available access to the shared services on the associated nodes. For example, an HA group that consists of interfaces only on Gateway Nodes or on both Admin Nodes and Gateway Nodes provides highly available access to the shared Load Balancer service.

To learn more about high availability groups, see Manage high availability (HA) groups.

Using HA groups

The best practices for creating a StorageGRID HA group for FabricPool depend on the workload.

- If you plan to use FabricPool with primary workload data, you must create an HA group that includes at least two load-balancing nodes to prevent data retrieval interruption.
- If you plan to use the FabricPool snapshot-only volume tiering policy or non-primary local performance tiers (for example, disaster recovery locations or NetApp SnapMirror® destinations), you can configure an HA group with only one node.

These instructions describe setting up an HA group for Active-Backup HA (one node is active and one node is backup). However, you might prefer to use DNS Round Robin or Active-Active HA. To learn the benefits of these other HA configurations, see Configuration options for HA groups.

Best practices for load balancing for FabricPool

Before attaching StorageGRID as a FabricPool cloud tier, review the best practices for using load balancers with FabricPool.

To learn general information about the StorageGRID load balancer and the load balancer certificate, see Considerations for load balancing.

Best practices for tenant access to the load balancer endpoint used for FabricPool

You can control which tenants can use a specific load balancer endpoint to access their buckets. You can allow all tenants, allow some tenants, or block some tenants. When creating a load balance endpoint for FabricPool use, select **Allow all tenants**. ONTAP encrypts the data that is placed in StorageGRID buckets, so little additional security would be provided by this extra security layer.

Best practices for the security certificate

When you create a StorageGRID load balancer endpoint for FabricPool use, you provide the security certificate that will allow ONTAP to authenticate with StorageGRID.

In most cases, the connection between ONTAP and StorageGRID should use Transport Layer Security (TLS) encryption. Using FabricPool without TLS encryption is supported but not recommended. When you select the network protocol for the StorageGRID load balancer endpoint, select **HTTPS**. Then provide the security certificate that will allow ONTAP to authenticate with StorageGRID.

To learn more about the server certificate for a load balancing endpoint:

- Manage security certificates
- Considerations for load balancing
- Hardening guidelines for server certificates

Add certificate to ONTAP

When you add StorageGRID as a FabricPool cloud tier, you must install the same certificate on the ONTAP cluster, including the root and any subordinate certificate authority (CA) certificates.

Manage certificate expiration



If the certificate used to secure the connection between ONTAP and StorageGRID expires, FabricPool will temporarily stop working and ONTAP will temporarily lose access to data tiered to StorageGRID.

To avoid certificate expiration issues, follow these best practices:

- Carefully monitor any alerts that warn of approaching certificate expiration dates, such as the **Expiration of load balancer endpoint certificate** and **Expiration of global server certificate for S3 and Swift API** alerts.
- Always keep the StorageGRID and ONTAP versions of the certificate in sync. If you replace or renew the certificate used for a load balancer endpoint, you must replace or renew the equivalent certificate used by ONTAP for the cloud tier.
- Use a publicly signed CA certificate. If you use a certificate signed by a CA, you can use the Grid Management API to automate certificate rotation. This allows you to replace soon-to-expire certificates nondisruptively.
- If you have generated a self-signed StorageGRID certificate and that certificate is about to expire, you must
 manually replace the certificate in both StorageGRID and in ONTAP before the existing certificate expires.
 If a self-signed certificate has already expired, turn off certificate validation in ONTAP to prevent access
 loss.

See NetApp Knowledge Base: How to configure a new StorageGRID self-signed server certificate on an existing ONTAP FabricPool deployment for instructions.

Best practices for using ILM with FabricPool data

If you are using FabricPool to tier data to StorageGRID, you must understand the requirements for using StorageGRID information lifecycle management (ILM) with FabricPool data.



FabricPool has no knowledge of StorageGRID ILM rules or policies. Data loss can occur if the StorageGRID ILM policy is misconfigured. For detailed information, see Create an ILM rule: Overview and Create an ILM policy: Overview.

Guidelines for using ILM with FabricPool

When you use the FabricPool setup wizard, the wizard automatically creates a new ILM rule for each S3 bucket you create and adds that rule to an inactive policy. You are prompted to activate the policy. The automatically created rule follows the recommended best practices: it uses 2+1 erasure coding at a single site.

If you are configuring StorageGRID manually instead of using the FabricPool setup wizard, review these guidelines to ensure that your ILM rules and ILM policy are suitable for FabricPool data and your business requirements. You might need to create new rules and update your active ILM policies to meet these guidelines.

• You can use any combination of replication and erasure-coding rules to protect cloud tier data.

The recommended best practice is to use 2+1 erasure coding within a site for cost-efficient data protection. Erasure coding uses more CPU, but offers significantly less storage capacity, than replication. The 4+1 and 6+1 schemes use less capacity than the 2+1 scheme. However, the 4+1 and 6+1 schemes are less flexible if you need to add Storage Nodes during grid expansion. For details, see Add storage capacity for erasure-coded objects.

• Each rule applied to FabricPool data must either use erasure coding or it must create at least two replicated copies.



An ILM rule that creates only one replicated copy for any time period puts data at risk of permanent loss. If only one replicated copy of an object exists, that object is lost if a Storage Node fails or has a significant error. You also temporarily lose access to the object during maintenance procedures such as upgrades.

• If you need to remove FabricPool data from StorageGRID, use ONTAP to retrieve all data for the FabricPool volume and promote it to the performance tier.



To avoid data loss, do not use an ILM rule that will expire or delete FabricPool cloud tier data. Set the retention period in each ILM rule to **forever** to ensure that FabricPool objects aren't deleted by StorageGRID ILM.

• Don't create rules that will move FabricPool cloud tier data out of the bucket to another location. You can't use a Cloud Storage Pool to move FabricPool data to another object store. Similarly, you can't archive FabricPool data to tape using an Archive Node.



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

• Starting with ONTAP 9.8, you can optionally create object tags to help classify and sort tiered data for easier management. For example, you can set tags only on FabricPool volumes attached to StorageGRID. Then, when you create ILM rules in StorageGRID, you can use the Object Tag advanced filter to select and place this data.

Other best practices for StorageGRID and FabricPool

When configuring a StorageGRID system for use with FabricPool, you might need to change other StorageGRID options. Before changing a global setting, consider how the change will affect other S3 applications.

Audit message and log destinations

FabricPool workloads often have a high rate of read operations, which can generate a high volume of audit messages.

- If you don't require a record of client read operations for FabricPool or any other S3 application, optionally
 go to CONFIGURATION > Monitoring > Audit and syslog server. Change the Client Reads setting to
 Error to decrease the number of audit messages recorded in the audit log. See Configure audit messages
 and log destinations for details.
- If you have a large grid, use multiple types of S3 applications, or want to retain all audit data, configure an external syslog server and save audit information remotely. Using an external server minimizes the performance impact of audit message logging without reducing the completeness of of audit data. See Considerations for external syslog server for details.

Object encryption

When configuring StorageGRID, you can optionally enable the global option for stored object encryption if data encryption is required for other StorageGRID clients. The data that is tiered from FabricPool to StorageGRID is already encrypted, so enabling the StorageGRID setting is not required. Client-side encryption keys are owned by ONTAP.

Object compression

When configuring StorageGRID, don't enable the global option to compress stored objects. The data that is tiered from FabricPool to StorageGRID is already compressed. Using the StorageGRID option will not further reduce an object's size.

Bucket consistency

For FabricPool buckets, the recommended bucket consistency is **Read-after-new-write**, which is the default consistency for a new bucket. Don't edit FabricPool buckets to use **Available** or **Strong-site**.

FabricPool tiering

If a StorageGRID node uses storage assigned from a NetApp ONTAP system, confirm that the volume does not have a FabricPool tiering policy enabled. For example, if a StorageGRID node is running on a VMware host, ensure the volume backing the datastore for the StorageGRID node does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.



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

Remove FabricPool data from StorageGRID

If you need to remove the FabricPool data that is currently stored in StorageGRID, you must use ONTAP to retrieve all data for the FabricPool volume and promote it to the performance tier.

Before you begin

• You have reviewed the instructions and considerations in Promote data to the performance tier.

- You are using ONTAP 9.8 or later.
- You are using a supported web browser.
- You belong to a StorageGRID user group for the FabricPool tenant account that has the Manage all buckets or Root access permission.

About this task

These instructions explain how to move data from StorageGRID back to FabricPool. You perform this procedure using ONTAP and StorageGRID Tenant Manager.

Steps

1. From ONTAP, issue the volume modify command.

Set tiering-policy to none to stop new tiering and set cloud-retrieval-policy to promote to return all data that was previously tiered to StorageGRID.

See Promote all data from a FabricPool volume to the performance tier.

2. Wait for the operation to complete.

You can use the volume object-store command with the tiering option to check the status of the performance tier promotion.

- 3. When the promote operation is complete, sign in to StorageGRID Tenant Manager for the FabricPool tenant account.
- 4. Select View buckets from the dashboard, or select STORAGE (S3) > Buckets.
- 5. Confirm that the FabricPool bucket is now empty.
- 6. If the bucket is empty, delete the bucket.

After you finish

When you delete the bucket, tiering from FabricPool to StorageGRID can no longer continue. However, because the local tier is still attached to the StorageGRID cloud tier, ONTAP System Manager will return error messages indicating that the bucket is inaccessible.

To prevent these error messages, do either of the following:

- Use FabricPool Mirror to attach a different cloud tier to the aggregate.
- Move the data from the FabricPool aggregate to a non-FabricPool aggregate and then delete the unused aggregate.

See the ONTAP documentation for FabricPool for instructions.

Copyright information

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

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

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

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

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

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

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

Trademark information

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