



Perform system administration

StorageGRID

NetApp
September 04, 2024

Table of Contents

- Perform system administration 1
 - Administer StorageGRID 1
 - Manage objects with ILM 291
 - System hardening 450
 - Configure FabricPool 458

Perform system administration

Administer StorageGRID

Administer StorageGRID: Overview

Use these instructions to configure and administer a StorageGRID system.

About these instructions

These instructions describe how to use the Grid Manager to set up groups and users, create tenant accounts to allow S3 and Swift client applications to store and retrieve objects, configure and manage StorageGRID networks, configure AutoSupport, manage node settings, and more.

These instructions are for technical personnel who will be configuring, administering, and supporting a StorageGRID system after it has been installed.

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 StorageGRID

Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	96
Microsoft Edge	96
Mozilla Firefox	94

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.

What you'll need

- You have your login credentials.
- You have the URL for the Grid Manager.
- You are using a [supported web browser](#).
- Cookies are enabled in your web browser.
- You have specific access permissions.

About this task

Each StorageGRID system includes one primary Admin Node and any number of non-primary Admin Nodes. You can sign in to the Grid Manager on any Admin Node to manage the StorageGRID system. However, the Admin Nodes are not exactly the same:

- Alarm acknowledgments (legacy system) made on one Admin Node are not copied to other Admin Nodes. For this reason, the information displayed for alarms might not look the same on each Admin Node.
- Some maintenance procedures can only be performed from the primary Admin Node.

If Admin Nodes are included in a high availability (HA) group, you connect using the virtual IP address of the HA group or a fully qualified domain name that maps to the virtual IP address. The primary Admin Node should be selected as the group's 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.

Steps

1. Launch a supported web browser.
2. In the browser's address bar, enter the URL for the Grid Manager:

```
https://FQDN_or_Admin_Node_IP/
```

where *FQDN_or_Admin_Node_IP* is a fully qualified domain name or the IP address of an Admin Node or the virtual IP address of an HA group of Admin Nodes.

If you must access the Grid Manager on a port other than the standard port for HTTPS (443), enter the following, where *FQDN_or_Admin_Node_IP* is a fully qualified domain name or IP address, and *port* is the port number:


```
https://FQDN_or_Admin_Node_IP:port/
```

3. If you are prompted with a security alert, install the certificate using the browser's installation wizard (see [About security certificates](#)).
4. Sign in to the Grid Manager:
 - If single sign-on (SSO) is not being used for your StorageGRID system:
 - i. Enter your username and password for the Grid Manager.
 - ii. Select **Sign In**.




The image shows the StorageGRID Grid Manager login page. On the left is the NetApp logo. On the right, the title "StorageGRID® Grid Manager" is displayed. Below the title are two input fields: "Username" and "Password". A "Sign in" button is located at the bottom right of the form area.

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



The image shows the StorageGRID Sign in page. On the left is the NetApp logo. On the right, the title "StorageGRID® Sign in" is displayed. Below the title is an "Account ID" input field containing a long string of zeros. Below this field is the text "For Grid Manager, leave this field blank." A "Sign in" button is located at the bottom right of the form area.

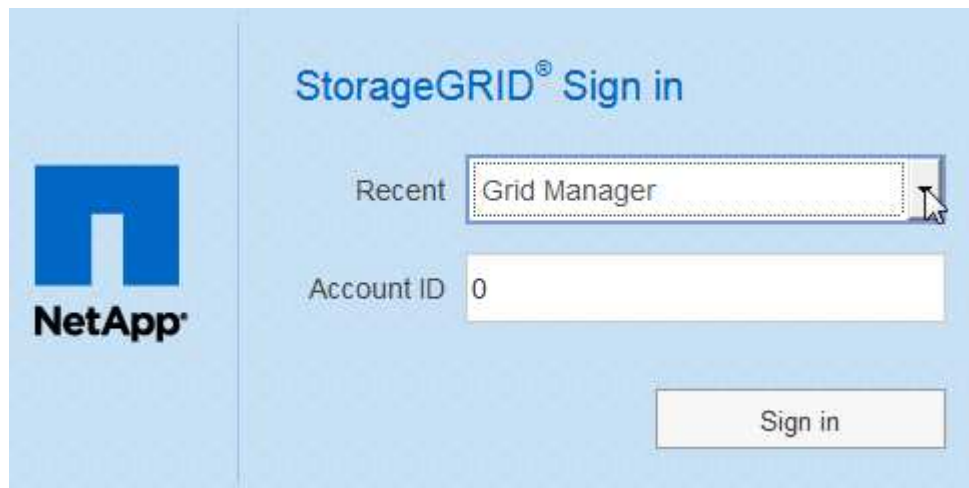
- ii. Enter your standard SSO credentials on your organization's SSO sign-in page. For example:



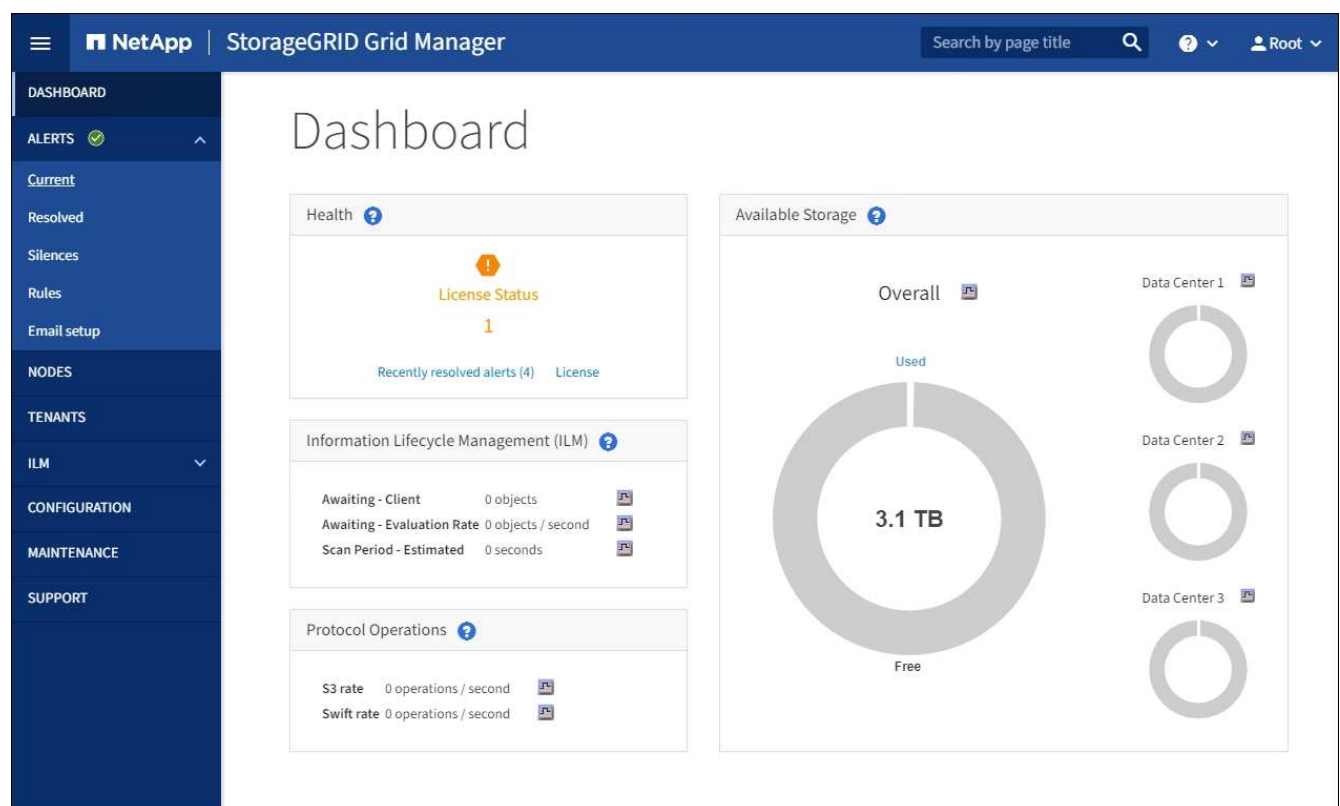
The image shows a sign-in form for an organizational account. The title is "Sign in with your organizational account". Below the title are two input fields: one for an email address (containing "someone@example.com") and one for a password (containing "Password"). A blue "Sign in" button is located at the bottom left of the form area.

- If SSO is enabled for your StorageGRID system and you have previously accessed the Grid Manager or a tenant account:
 - i. Do either of the following:

- Enter **0** (the account ID for the Grid Manager), and select **Sign in**.
- Select **Grid Manager** if it appears in the list of recent accounts, and select **Sign in**.



- ii. Sign in with your standard SSO credentials on your organization's SSO sign-in page. When you are signed in, the home page of the Grid Manager appears, which includes the Dashboard. To learn what information is provided, see [View the Dashboard](#).



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

Option	Steps
SSO not enabled	<ol style="list-style-type: none"> 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. Enter your username and password for the Grid Manager. Select Sign In.
SSO enabled	<p>In the browser's address bar, enter the fully qualified domain name or IP address of the other Admin Node.</p> <p>If you have signed in to one Admin Node, you can access other Admin Nodes without having to sign in again. However, if your SSO session expires, you are prompted for your credentials again.</p> <p>Note: SSO is not available on the restricted Grid Manager port. You must use the default HTTPS port (443) if you want users to authenticate with single sign-on.</p>

Related information

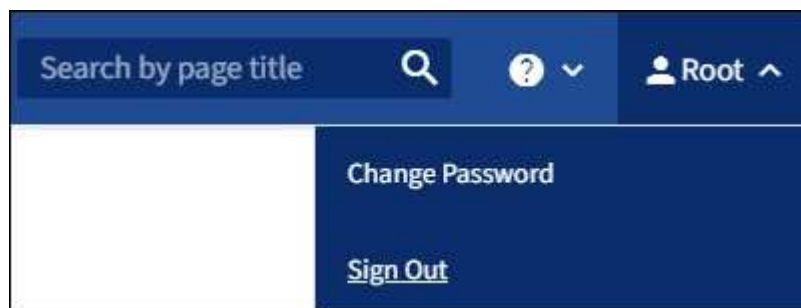
- [Control access through firewalls](#)
- [Configure single sign-on](#)
- [Manage admin groups](#)
- [Manage high availability groups](#)
- [Use a tenant account](#)
- [Monitor and troubleshoot](#)

Sign out of the Grid Manager

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

Steps

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



2. Select **Sign Out**.

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

Related information

- [Configure single sign-on](#)
- [Use a tenant account](#)

Change your password

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

What you'll need

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 cannot change your password in Grid Manager. Instead, you must change your password in the external identity source, for example, Active Directory or OpenLDAP.

Steps

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

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

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

Change the browser session timeout

You can control whether Grid Manager and Tenant Manager users are signed out if they are inactive for more than a certain amount of time.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

About this task

The GUI Inactivity Timeout defaults to 900 seconds (15 minutes). If a user's browser session is not active for this amount of time, the session times out.

As required, you can increase or decrease the timeout period by setting the GUI Inactivity Timeout display option.

If single sign-on (SSO) is enabled and a user's browser session times out, the system behaves as if the user selected **Sign Out** manually. The user must reenter their SSO credentials to access StorageGRID again. See [Configure single sign-on](#).

User session timeout can also be controlled by the following:



- A separate, non-configurable StorageGRID timer, which is included for system security. By default, each user's authentication token expires 16 hours after the user signs in. When a user's authentication expires, that user is automatically signed out, even if the value for the GUI Inactivity Timeout has not been reached. To renew the token, the user must sign back in.
- Timeout settings for the identity provider, assuming SSO is enabled for StorageGRID.

Steps

1. Select **CONFIGURATION > System > Display options**.
2. For **GUI Inactivity Timeout**, enter a timeout period of 60 seconds or more.

Set this field to 0 if you do not want to use this functionality. Users are signed out 16 hours after they sign in, when their authentication tokens expire.



Display Options

Updated: 2017-03-09 20:36:53 MST

Current Sender	ADMIN-DC1-ADM1
Preferred Sender	ADMIN-DC1-ADM1
GUI Inactivity Timeout	900
Notification Suppress All	<input type="checkbox"/>

Apply Changes

3. Select **Apply Changes**.

The new setting does not affect currently signed in users. Users must sign in again or refresh their browsers for the new timeout setting to take effect.

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.

What you'll need

- 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 panel on the Dashboard includes a License Status icon and a **License** link. The number indicates how many license-related issues there are.



Step

To view the license, do one of the following:

- From the Health panel on the Dashboard, select the License status icon or the **License** link. This link appears only if there is an issue with the license.
- Select **MAINTENANCE > System > License**.

The License Page appears and provides the following, read-only information about the current license:

- StorageGRID system ID, which is the unique identification number for this StorageGRID installation
- License serial number
- Licensed storage capacity of the grid
- Software license end date
- Support service contract end date
- Contents of the license text file



For licenses issued before StorageGRID 10.3, the licensed storage capacity is not included in the license file, and a "See License Agreement" message is displayed instead of a value.

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.

What you'll need

- 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. Enter the provisioning passphrase for your StorageGRID system in the **Provisioning Passphrase** text box.
3. Select **Browse**.
4. In the Open dialog box, locate and select the new license file (.txt), and select **Open**.

The new license file is validated and displayed.

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.

What you'll need

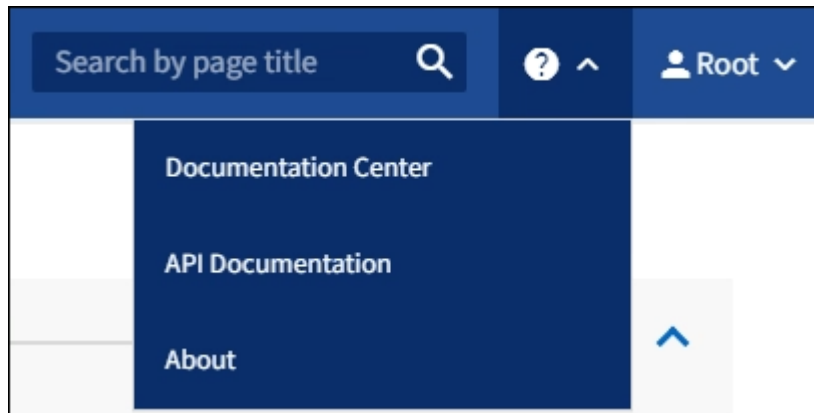
- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.



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

🔒

Parameters

Try it out

Name	Description
type string <small>(query)</small>	filter by group type Available values : local, federated <div>— ▼</div>
limit integer <small>(query)</small>	maximum number of results Default value : 25 <div>25</div>
marker string <small>(query)</small>	marker-style pagination offset (value is Group's URN) <div>marker - marker-style pagination offset (value</div>
includeMarker boolean <small>(query)</small>	if set, the marker element is also returned <div>— ▼</div>
order string <small>(query)</small>	pagination order (desc requires marker) Available values : asc, desc <div>— ▼</div>

Responses

Response content type application/json ▼

Code	Description
200	successfully retrieved Example Value Model <pre>{ "responseTime": "2021-03-29T14:22:19.673Z", "status": "success", "apiVersion": "3.3", "deprecated": false, "data": [{ "displayName": "Developers", </pre>

- 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.
- Determine if you need to modify the example request body. If so, you can select **Model** to learn the requirements for each field.
- Select **Try it out**.
- Provide any required parameters, or modify the request body as required.
- Select **Execute**.
- Review the response code to determine if the request was successful.

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*").



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

See “Protecting against Cross-Site Request Forgery” for information on improving authentication security.

- **client-certificates** — Operations to configure client certificates so that StorageGRID can be accessed securely using external monitoring tools.
- **config** — Operations related to the product release and versions of the Grid Management API. You can list the product release version and the major versions of the Grid Management API supported by that release, and you can disable deprecated versions of the API.
- **deactivated-features** — Operations to view features that might have been deactivated.
- **dns-servers** — Operations to list and change configured external DNS servers.
- **endpoint-domain-names** — Operations to list and change endpoint domain names.
- **erasure-coding** — Operations on Erasure Coding profiles.
- **expansion** — Operations on expansion (procedure-level).
- **expansion-nodes** — Operations on expansion (node-level).
- **expansion-sites** — Operations on expansion (site-level).
- **grid-networks** — Operations to list and change the Grid Network List.
- **grid-passwords** — Operations for grid password management.
- **groups** — Operations to manage local Grid Administrator Groups and to retrieve federated Grid Administrator Groups from an external LDAP server.

- **identity-source** — Operations to configure an external identity source and to manually synchronize federated group and user information.
- **ilm** — Operations on information lifecycle management (ILM).
- **license** — Operations to retrieve and update the StorageGRID license.
- **logs** — Operations for collecting and downloading log files.
- **metrics** — Operations on StorageGRID metrics including instant metric queries at a single point in time and range metric queries over a range of time. The Grid Management API uses the Prometheus systems monitoring tool as the backend data source. For information about constructing Prometheus queries, see the Prometheus web site.



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

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

Grid Management API versioning

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

For example, this Request URL specifies version 3 of the API.

`https://hostname_or_ip_address/api/v3/authorize`

The major version of the Tenant Management API is bumped when changes are made that are **not compatible** with older versions. The minor version of the Tenant Management API is bumped when changes are made that **are compatible** with older versions. Compatible changes include the addition of new endpoints or new properties. The following example illustrates how the API version is bumped based on the type of changes made.

Type of change to API	Old version	New version
Compatible with older versions	2.1	2.2

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

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



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

Outdated requests are marked as deprecated in the following ways:

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

```
Received call to deprecated v1 API at POST "/api/v1/authorize"
```

Determine which API versions are supported in the current release

Use the following API request to return a list of the supported API major versions:

```
GET https://{IP-Address}/api/versions
{
  "responseTime": "2019-01-10T20:41:00.845Z",
  "status": "success",
  "apiVersion": "3.0",
  "data": [
    2,
    3
  ]
}
```

Specify an API version for a request

You can specify the API version using a path parameter (/api/v3) or a header (Api-Version: 3). If you provide both values, the header value overrides the path value.

```
curl https://[IP-Address]/api/v3/grid/accounts

curl -H "Api-Version: 3" https://[IP-Address]/api/grid/accounts
```

Protect against Cross-Site Request Forgery (CSRF)

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

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

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

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

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

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

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

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

See the online API documentation for additional examples and details.



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

Use the API if single sign-on is enabled

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

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

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

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

What you'll need

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

About this task

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

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

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



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

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

Steps

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

When prompted, enter values for the following arguments:

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

```
python3 storagegrid-ssoauth.py
sso_method: adfs
saml_user: my-sso-username
saml_domain: my-domain
sg_address: storagegrid.example.com
tenant_account_id: 12345
Enter the user's SAML password:
*****

*****

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

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

3. If you want to use curl requests, use the following procedure.

a. Declare the variables needed to sign in.

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



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

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

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

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

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

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

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

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

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

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

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

The response includes the client request ID:

```
<form method="post" id="loginForm" autocomplete="off"
novalidate="novalidate" onKeyPress="if (event && event.keyCode == 13)
Login.submitLoginRequest();" action="/adfs/ls/?
SAMLRequest=fZHRTToMwFIZfhh...UJikvo77sXPw%3D%3D&RelayState=12345&clie
nt-request-id=00000000-0000-0000-ee02-0080000000de" >
```

- e. Save the client request ID from the response.

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

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

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


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



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

```
HTTP/1.1 302 Found
Content-Length: 0
Content-Type: text/html; charset=utf-8
Location:
https://adfs.example.com/adfs/ls/?SAMLRequest=fZHRT0MwFIZfhb...UJikvo
77sXPw%3D%3D&RelayState=12345&client-request-id=00000000-0000-0000-
ee02-0080000000de
Set-Cookie: MSISAuth=AAEAADAvsHpXk6ApV...pmP0aEiNtJvWY=; path=/adfs;
HttpOnly; Secure
Date: Tue, 06 Nov 2018 16:55:05 GMT
```

g. Save the MSISAuth cookie from the response.

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

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

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

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

```

HTTP/1.1 200 OK
Cache-Control: no-cache,no-store
Pragma: no-cache
Content-Length: 5665
Content-Type: text/html; charset=utf-8
Expires: -1
Server: Microsoft-HTTPAPI/2.0
P3P: ADFS doesn't have P3P policy, please contact your site's admin
for more details
Set-Cookie:
SamlSession=a3dpbnRlcnMtUHJpbWFyeS1BZG1pbi0xNzgmRmFsc2Umcng4NnJDZmFKV
XFxVWx3bk1lMnFuUSUzZCUzZCYmJiYmXzE3MjAyZTA5LTNmMDgtNDRkZC04Yzg5LTQ3ND
UxYzA3ZjkzYw==; path=/adfs; HttpOnly; Secure
Set-Cookie: MSISAuthenticated=MTEvNy8yMDE4IDQ6MzI6NTkgUE0=;
path=/adfs; HttpOnly; Secure
Set-Cookie: MSISLoopDetectionCookie=MjAxOC0xMS0wNzoxNjozMj01OVpcMQ==;
path=/adfs; HttpOnly; Secure
Date: Wed, 07 Nov 2018 16:32:59 GMT

<form method="POST" name="hiddenform"
action="https://storagegrid.example.com:443/api/saml-response">
  <input type="hidden" name="SAMLResponse"
value="PHNhbWxwOlJlc3Bvb3N...1scDpSZXNwb25zZT4=" /><input
type="hidden" name="RelayState" value="12345" />

```

- i. Save the SAMLResponse from the hidden field:

```
export SAMLResponse='PHNhbWxwOlJlc3Bvb3N...1scDpSZXNwb25zZT4='
```

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

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

```

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

```

The response includes the authentication token.

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

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

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

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

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

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

About this task

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

Steps

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

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

A logout URL is returned:

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

2. Save the logout URL.

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

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

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

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

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

4. Delete the StorageGRID bearer token.

Deleting the StorageGRID bearer token works the same way as without SSO. If 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

What you'll need

- 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 or CentOS, `./debs` for Ubuntu or Debian, and `./vsphere` for VMware).

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

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

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

Steps

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

```
npm install -g <module>
```

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

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

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

requested.

```
c:\Users\user\Documents\azure_sso>py storagegrid-azure-ssoauth.py --sso-email-address user@my-domain.com
--sg-address storagegrid.examp.e.com --tenant-account-id 0
Enter the user's SSO password:
*****

Watch for and approve a 2FA authorization request
*****

StorageGRID Auth Token: {'responseTime': '2021-10-04T21:30:48.807Z', 'status': 'success', 'apiVersion':
'3.4', 'data': '4807d93e-a3df-48f2-9680-906cd255979e'}
```



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

What you'll need

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

About this task

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

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

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



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

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

Steps

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

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

When prompted, enter values for the following arguments:

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

```
python3 storagegrid-ssoauth.py
sso_method: pingfederate
saml_user: my-sso-username
saml_domain:
sg_address: storagegrid.example.com
tenant_account_id: 12345
Enter the user's SAML password:
*****
*****
StorageGRID Auth Token: 56eb07bf-21f6-40b7-afob-5c6cacfb25e7
```

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

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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



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

g. Export the 'action' value:

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

h. Send cookies along with credentials:

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

i. Save the SAMLResponse from the hidden field:

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

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

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

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

The response includes the authentication token.

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

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

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

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

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

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

About this task

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

Steps

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

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

A logout URL is returned:

```
{
  "apiVersion": "3.0",
  "data": "https://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
```

Control access to StorageGRID

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.

What you'll need


- 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. Select **Make a change** under **Change provisioning passphrase**.



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.

Configuration > Grid passwords > Change provisioning passphrase

Success
Provisioning passphrase changed successfully

Change provisioning passphrase

The provisioning passphrase is required for any installation, expansion, or maintenance procedure that makes changes to the grid topology. This passphrase is also required to [download backups of the grid topology information and encryption keys for the StorageGRID system](#). After changing the provisioning passphrase, you must download a new [Recovery Package](#).

Current provisioning passphrase

New provisioning passphrase

Confirm new provisioning passphrase

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.

What you'll need

- 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 are not 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.



Do not 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 check box 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.

Control access through firewalls

When you want to control access through firewalls, you open or close specific ports at the external firewall.

Control access at the external firewall

You can control access to the user interfaces and APIs on StorageGRID Admin Nodes by opening or closing specific ports at the external firewall. For example, you might want to prevent tenants from being able to connect to the Grid Manager at the firewall, in addition to using other methods to control system access.

Port	Description	If port is open...
443	Default HTTPS port for Admin Nodes	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	<ul style="list-style-type: none">• Web browsers and management API clients can access the Grid Manager and the Grid Management API using HTTPS.• Web browsers and management API clients cannot access the Tenant Manager or the Tenant Management API.• Requests for internal content will be rejected.

Port	Description	If port is open...
9443	Restricted Tenant Manager port on Admin Nodes	<ul style="list-style-type: none"> • Web browsers and management API clients can access the Tenant Manager and the Tenant Management API using HTTPS. • Web browsers and management API clients cannot 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](#)

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.

What you'll need

- 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 for using 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 do not use their own identity source. Users in tenant groups can sign in to the Tenant Manager and perform tasks, based on the permissions assigned to the group in the Tenant Manager. See [Create tenant account](#) and [Use a tenant account](#) for details.

Enter the configuration

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

LDAP service type

Select the type of LDAP service you want to configure.

Active Directory	Azure	OpenLDAP	Other
------------------	-------	----------	-------

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

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



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

- **Username:** The full path of the distinguished name (DN) for the user that will connect to the LDAP server.

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

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

- `sAMAccountName` or `uid`
- `objectGUID`, `entryUUID`, or `nsuniqueid`
- `cn`
- `memberOf` or `isMemberOf`
- **Active Directory:** `objectSid`, `primaryGroupID`, `userAccountControl`, and `userPrincipalName`
- **Azure:** `accountEnabled` and `userPrincipalName`
- **Password:** The password associated with the username.
- **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 cannot 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.

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. Do not include any special characters in the username, such as @ or /.

A screenshot of a 'Test Connection' dialog box. The title bar says 'Test Connection' with a close button (X) on the right. The main text reads: 'To test the connection and the bind username format, enter the username and password of a federated user. For example, enter your own federated username and password. The test values are not saved.' Below this, there are two input fields. The first is labeled 'Test username' and contains the text 'myusername'. Below it is a smaller text label: 'The username of a federated user.' The second input field is labeled 'Test password' and contains a series of dots for masking. To the right of the password field is an eye icon to toggle visibility. At the bottom left is a 'Cancel' button, and at the bottom right is a blue 'Test Connection' button.

- 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** check box is disabled if single sign-on (SSO) is set to **Enabled** or **Sandbox Mode**. The SSO Status on the Single Sign-on page must be **Disabled** before you can disable identity federation. See [Disable single sign-on](#).

Steps

1. Go to the Identity federation page.
2. Uncheck the **Enable identity federation** check box.

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 are not 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 and 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.

What you'll need

- 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

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

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 cannot update later.
4. Select **Continue**.

Federated group

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

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 cannot make any changes or perform any operations in the Grid Manager or Grid Management API. Local read-only users can change their own passwords.



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

2. Select one or more [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)

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 check box 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 check box 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 check box for the group. b. Select Actions > View group details . c. Optionally, change the group's Access mode. d. Optionally, select or unselect 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 unselect Group permissions . d. Select Save changes .

Duplicate a group

1. Select the check box 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.

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

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 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 do not have this permission cannot 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
- **MAINTENANCE > Tasks:**
 - Decommission
 - Expansion
 - Object existence check
 - Recovery
- **MAINTENANCE > System:**
 - Recovery package
 - Software update
- **SUPPORT > Tools:**
 - Logs

Users who do not have the Maintenance permission can view, but not edit, these pages:

- **MAINTENANCE > Network:**
 - DNS servers
 - Grid Network
 - NTP servers
- **MAINTENANCE > System:**
 - License
- **CONFIGURATION > Security:**
 - Certificates
 - Domain names
- **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 the **SUPPORT > Tools > Metrics** page. This permission also provides access to custom Prometheus metrics queries using the **Metrics** section of the Grid Management API.

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 > Network:**
 - Link cost
- **CONFIGURATION > System:**
 - Display options
 - Grid options
 - Storage options
- **SUPPORT > Alarms (legacy):**
 - Custom events
 - Global alarms
 - Legacy email setup

Storage appliance administrator

This permission provides access to the E-Series SANtricity System Manager on storage appliances through the Grid Manager.

Tenant accounts

This permission provides access to the Tenants page, where you can create, edit, and remove tenant accounts. This permission also allows users to view existing traffic classification policies.

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

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 } }
```

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.

What you'll need

- 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 cannot remove the root user.



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

Access the wizard

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

Enter user credentials

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

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 check box 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 check box 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 .
Deny or allow StorageGRID access	a. Select the check box 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 check box 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 .

Task	Actions menu	Details page
Change groups (local users only)	a. Select the check box 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.

1. Select the check box 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 cannot delete the root user.

1. From the Users page, select the check box 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 cannot sign in to StorageGRID.

How single sign-on works

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

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

Sign in when SSO is enabled


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

Steps

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


The StorageGRID Sign in page appears.

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



The screenshot shows the StorageGRID Sign in page. On the left is the NetApp logo. The main heading is "StorageGRID® Sign in". Below this is a text label "Account ID" followed by a text input field containing "00000000000000000000". Below the input field is a blue link that says "For Grid Manager, leave this field blank." At the bottom right is a "Sign in" button.

- If you have previously accessed either the Grid Manager or the Tenant Manager, you are prompted to select a recent account or to enter an account ID:



The screenshot shows the StorageGRID Sign in page. On the left is the NetApp logo. The main heading is "StorageGRID® Sign in". Below this is a text label "Recent" followed by a dropdown menu showing "S3 tenant". Below the dropdown is a text label "Account ID" followed by a text input field containing "27469746059057031822". Below the input field is a blue link that says "For Grid Manager, leave this field blank." At the bottom right is a "Sign in" button.



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:



Sign in with your organizational account

someone@example.com

Password

Sign in

4. Sign in with your SSO credentials.

If your SSO credentials are correct:

- The identity provider (IdP) provides an authentication response to StorageGRID.
- StorageGRID validates the authentication response.
- 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 do not 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 for using single sign-on

Before enabling single sign-on (SSO) for a StorageGRID system, review the requirements in this section.

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	<ul style="list-style-type: none"> • 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.

- AD FS 3.0, included with Windows Server 2012 R2 update, or higher.

Additional requirements

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

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 through firewalls](#).

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.

What you'll need

- 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** check box is not selected.
 - d. If it is, confirm that any federated groups that might be in use for this tenant account are no longer required, unselect the check box, and 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** check box is selected, uncheck the box and select **Save**.



Edit the tenant

Enter details ————— 2 Select permissions

Select permissions

Select the permissions for this tenant account.

- ☐ Allow platform services ?
- ☐ Use own identity source ?
- ☐ Allow S3 Select ?

The Tenant page appears.

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

Related information

- [Requirements for using 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.

What you'll need

- 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	<ul style="list-style-type: none"> • Active Directory • Azure • PingFederate
Azure	Azure

About this task

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

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

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

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

Access sandbox mode

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

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

Single Sign-on

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

SSO status ⓘ
☒ Disabled
☐ Sandbox Mode
☐ Enabled

Save



If the SSO Status options do not appear, confirm you have configured the identity provider as the federated identity source. See [Requirements for using single sign-on](#).

2. Select **Sandbox Mode**.

The Identity Provider section appears.

Enter identity provider details

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

Active Directory

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



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

- b. Specify which TLS certificate will be used to secure the connection when the identity provider sends SSO configuration information in response to StorageGRID requests.

- **Use operating system CA certificate:** Use the default CA certificate installed on the operating system to secure the connection.
- **Use custom CA certificate:** Use a custom CA certificate to secure the connection.

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

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

- 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 do not anticipate adding more Admin Nodes in the future, enter `SG` or `StorageGRID`.
- If your grid includes more than one Admin Node, include the string `[HOSTNAME]` in the identifier. For example, `SG-[HOSTNAME]`. This generates a table that shows the relying party identifier for each Admin Node in your system, based on the node's hostname.



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

- d. Select **Save**.

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



Azure

- a. Specify which TLS certificate will be used to secure the connection when the identity provider sends SSO configuration information in response to StorageGRID requests.

- **Use operating system CA certificate:** Use the default CA certificate installed on the operating system to secure the connection.
- **Use custom CA certificate:** Use a custom CA certificate to secure the connection.

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


- **Do not use TLS:** Do not use a TLS certificate to secure the connection.
- 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 do not anticipate adding more Admin Nodes in the future, enter `SG` or `StorageGRID`.
 - If your grid includes more than one Admin Node, include the string `[HOSTNAME]` in the identifier. For example, `SG-[HOSTNAME]`. This generates a table that shows an enterprise application name for each Admin Node in your system, based on the node's hostname.



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

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

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



PingFederate

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

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

 - **Do not use TLS:** Do not use a TLS certificate to secure the connection.
- 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 do not anticipate adding more Admin Nodes in the future, enter `SG` or `StorageGRID`.
 - If your grid includes more than one Admin Node, include the string `[HOSTNAME]` in the identifier. For example, `SG-[HOSTNAME]`. This generates a table that 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

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

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

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

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

Test SSO connections

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

Active Directory

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

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


Sandbox mode

Sandbox mode is currently enabled. Use this mode to configure relying party trusts and to confirm that single sign-on (SSO) and single logout (SLO) are correctly configured for the StorageGRID system.

1. Use Active Directory Federation Services (AD FS) to create relying party trusts for StorageGRID. Create one trust for each Admin Node, using the relying party identifier(s) shown below.
2. Go to your identity provider's sign-on page: <https://ad2016.saml.sgws/adfs/ls/idpinitiatedsignon.htm>
3. From this page, sign in to each StorageGRID relying party trust. If the SSO operation is successful, StorageGRID displays a page with a success message. Otherwise, an error message is displayed.

When you have confirmed SSO for each of the relying party trusts and you are ready to enforce the use of SSO for StorageGRID, change the SSO Status to Enabled, and click Save.

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



You are not signed in.

☐ Sign in to this site.

☒ Sign in to one of the following sites:

SG-DC1-ADM1

Sign in

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

✓ Single sign-on authentication and logout test completed successfully.

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

Azure

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

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.

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.

What you'll need

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



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

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

About this task

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

Create a relying party trust using Windows PowerShell

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

Steps

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

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

- For *Admin_Node_Identifier*, enter the Relying Party Identifier for the Admin Node, exactly as it appears on the Single Sign-on page. For example, SG-DC1-ADM1.
- For *Admin_Node_FQDN*, enter the fully qualified domain name for the same Admin Node. (If necessary, you can use the node's IP address instead. However, if you enter an IP address here, be

aware that you must update or recreate this relying party trust if that IP address ever changes.)

3. From Windows Server Manager, select **Tools > AD FS Management**.

The AD FS management tool appears.

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

The list of relying party trusts appears.

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

For example, **ObjectGUID to Name ID**.

- f. For the Attribute Store, select **Active Directory**.
 - g. In the LDAP Attribute column of the Mapping table, type **objectGUID**.
 - h. In the Outgoing Claim Type column of the Mapping table, select **Name ID** from the drop-down list.
 - i. 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 simply enter the values manually.

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

- e. For the Attribute Store, select **Active Directory**.
 - f. In the LDAP Attribute column of the Mapping table, type **objectGUID**.
 - g. In the Outgoing Claim Type column of the Mapping table, select **Name ID** from the drop-down list.
 - h. 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 simply enter the values manually.

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

Create a relying party trust manually

If you choose not to import the data for the relying party 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** check box.
- d. Type the SAML service endpoint URL for the Admin Node:

`https://Admin_Node_FQDN/api/saml-response`

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

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

Admin_Node_Identifier

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

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

The Edit Claim Issuance Policy dialog box appears.



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

6. To start the Claim Rule wizard, 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**.
 - c. For the Attribute Store, select **Active Directory**.
 - d. In the LDAP Attribute column of the Mapping table, type **objectGUID**.
 - e. In the Outgoing Claim Type column of the Mapping table, select **Name ID** from the drop-down list.
 - f. 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 do not have the custom certificate, log in to the Admin Node, go the `/var/local/mgmt-api` directory of the Admin Node, and add the `custom-server.crt` certificate file.

Note: Using the Admin Node's default certificate (`server.crt`) is not recommended. If the Admin Node fails, the default certificate will be regenerated when you recover the node, and you will need to update the relying party trust.

- b. 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.

What you'll need

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



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

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

Access Azure AD

1. Login to the [Azure Portal](#).
2. Navigate to [Azure Active Directory](#).
3. Select [Enterprise applications](#).

Create enterprise applications and save StorageGRID SSO configuration

In order 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.

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.

Repeat these steps for each Admin Node:

1. Sign in to StorageGRID from the Admin Node.
2. Select **CONFIGURATION > Access control > Single sign-on**.
3. Select the button to download the SAML metadata for that Admin Node.
4. 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:

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.

What you'll need

- 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.

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.

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

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

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

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

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 the **objectGUID** attribute and add it.
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 **objectGUID** 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

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, unselect **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** check boxes 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

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.

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.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

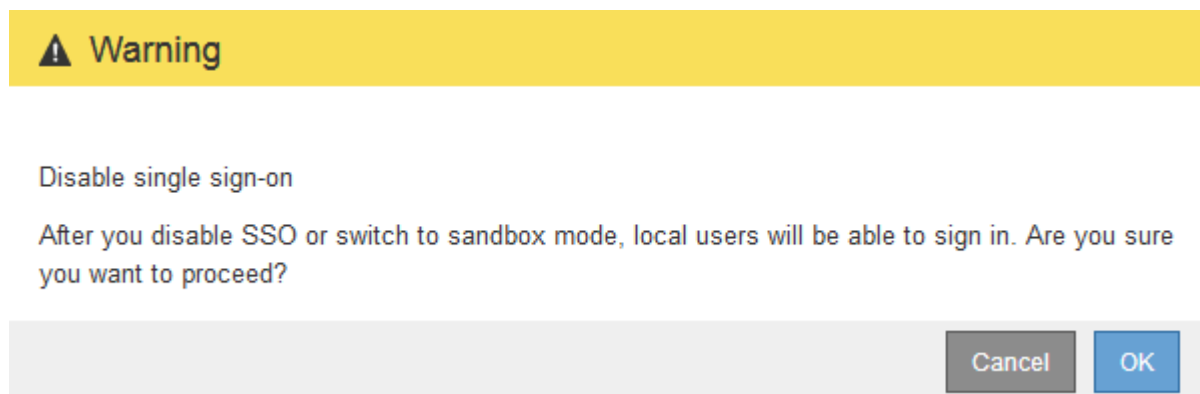
Steps

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.

What you'll need

- 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** check box on the Single Sign-on page in the Grid Manager remains selected, and all existing SSO settings are maintained unless you update them.

Steps

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

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

2. Run the following command: `disable-saml`

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

3. Confirm that you want to disable SSO.

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

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

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

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

Manage security settings

Manage certificates

About security certificates

Security certificates are small data files used to create secure, trusted connections between StorageGRID components and between StorageGRID components and external systems.

StorageGRID uses two types of security certificates:

- **Server certificates** are required when you use HTTPS connections. Server certificates are used to establish secure connections between clients and servers, authenticating the identity of a server to its clients and providing a secure communication path for data. The server and the client each have a copy of the certificate.
- **Client certificates** authenticate a client or user identity to the server, providing more secure authentication than passwords alone. Client certificates do not encrypt data.

When a client connects to the server using HTTPS, the server responds with the server certificate, which contains a public key. The client verifies this certificate by comparing the server signature to the signature on its copy of the certificate. If the signatures match, the client starts a session with the server using the same public key.

StorageGRID functions as the server for some connections (such as the load balancer endpoint) or as the client for other connections (such as the CloudMirror replication service).

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 are not recommended.

- Custom CA certificates do not remove the internal certificates; however, the custom certificates should be the ones specified for verifying server connections.
- All custom certificates must meet the [system hardening guidelines](#) for server certificates.
- 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.

1. From Grid Manager, select **CONFIGURATON > Security > Certificates**.

Certificates

View and manage the certificates that secure HTTPS connections between StorageGRID and external clients, such as S3 or Swift, and external servers, such as a key management server (KMS).

Global

Grid CA

Client

Load balancer endpoints

Tenants

Other

The StorageGRID certificate authority ("grid CA") generates and signs two global certificates during installation. The management interface certificate on Admin Nodes secures the management interface. The S3 and Swift API certificate on Storage and Gateway Nodes secures client access. You should replace each default certificate with your own custom certificate signed by an external certificate authority.

Name	Description	Type ⓘ	Expiration date ⓘ ⌵
Management interface certificate	Secures the connection between client web browsers and the Grid Manager, Tenant Manager, Grid Management API, and Tenant Management API.	Custom	Jun 4th, 2022
S3 and Swift API certificate	Secures the connections between S3 and Swift clients and Storage Nodes or between clients and the deprecated CLB service on Gateway Nodes. You can optionally use this certificate for a load balancer endpoint as well.	Custom	Jun 4th, 2022

2. Select a tab on the Certificates page for information about each certificate category and to access the certificate settings. You can only 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, CLB service (deprecated), 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 cannot 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](#), that you upload or generate, 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:

- [Select an endpoint name to open a browser tab with information about the 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:

- [Identity federation certificates](#)

- [Cloud Storage Pool certificates](#)
- [Key management server \(KMS\) certificates](#)
- [Single sign-on certificates](#)
- [Email alert notification certificates](#)
- [External syslog server 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:

- [View and edit an identity federation certificate](#)
- [Upload key management server \(KMS\) server and client certificates](#)
- [Specify a Cloud Storage Pool certificate for S3, C2S S3, or Azure](#)
- [Manually specify an SSO certificate for relying party trust](#)
- [Specify a certificate for alert email notifications](#)
- [Specify an external syslog server certificate](#)

Security certificate details

Each type of security certificate is described below, with links to articles that contain implementation instructions.

Management interface certificate

Certificate type	Description	Navigation location	Details
Server	<p>Authenticates the connection between client web browsers and the StorageGRID management interface, allowing users to access the Grid Manager and Tenant Manager without security warnings.</p> <p>This certificate also authenticates Grid Management API and Tenant Management API connections.</p> <p>You can use the default certificate created during installation or upload a custom certificate.</p>	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, to the deprecated Connection Load Balancer (CLB) service on a Gateway Node, and 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	<p>Installed on each client, allowing StorageGRID to authenticate external client access.</p> <ul style="list-style-type: none"> • 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	<p>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.</p> <p>You can also use a custom version of the global S3 and Swift API certificate to authenticate connections to the Load Balancer service. If the global certificate is used to authenticate load balancer connections, you do not need to upload or generate a separate certificate for each load balancer endpoint.</p> <p>Note: The certificate used for load balancer authentication is the most used certificate during normal StorageGRID operation.</p>	CONFIGURATION > Network > Load balancer endpoints	<ul style="list-style-type: none"> • Configure load balancer endpoints • Create a load balancer endpoint for FabricPool

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

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

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

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)

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

Email alert notification certificate

Certificate type	Description	Navigation location	Details
Server and client	<p>Authenticates the connection between an SMTP email server and StorageGRID that is used for alert notifications.</p> <ul style="list-style-type: none"> • 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	<p>Authenticates the TLS or RELP/TLS connection between an external syslog server that logs events in StorageGRID.</p> <p>Note: An external syslog server certificate is not required for TCP, RELP/TCP, and UDP connections to an external syslog server.</p>	CONFIGURATION > Monitoring > Audit and syslog server and then select Configure external syslog server	Configure an external syslog server

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).

For more information on how StorageGRID secures client connections for the REST API, see [Use S3](#) or [Use Swift](#).

Configure management interface certificates

You can replace the default management interface certificate with a single custom certificate that allows users to access the Grid Manager and the Tenant Manager without encountering security warnings. You can also revert to the default management interface certificate or generate a new one.

About this task

By default, every Admin Node is issued a certificate signed by the grid CA. These CA signed certificates can be replaced by a single common custom management interface certificate and corresponding private key.

Because a single custom management interface certificate is used for all Admin Nodes, you must specify the certificate as a wildcard or multi-domain certificate if clients need to verify the hostname when connecting to the Grid Manager and Tenant Manager. Define the custom certificate such that it matches all Admin Nodes in the grid.

You need to complete configuration on the server, and depending on the root certificate authority (CA) you are using, users might also need to install the Grid CA certificate in the web browser they will use to access the Grid Manager and the Tenant Manager.



To ensure that operations are not 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.

1. Select **Upload certificate**.
2. 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.
3. 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.
4. 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.

1. Select **Generate certificate**.
2. Specify the certificate information:
 - **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**: X.509 subject or distinguished name (DN) of the certificate owner.
 - **Days valid**: Number of days after creation that the certificate expires.
3. Select **Generate**.
4. 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.

5. Select **Save**.

The custom management interface certificate is used for all subsequent new connections to the Grid Manager, Tenant Manager, Grid Manager API or Tenant Manager API.

5. Refresh the page to ensure the web browser is updated.



After uploading or generating a new certificate, allow up to one day for any related certificate expiration alerts to clear.

6. After you add a custom management interface certificate, the Management interface certificate page displays detailed certificate information for the certificates that are in use. You can download or copy the certificate PEM as required.

Restore the default management interface certificate

You can revert to using the default management interface certificate for Grid Manager and Tenant Manager connections.

Steps

1. Select **CONFIGURATION > Security > Certificates**.
2. On the **Global** tab, select **Management interface certificate**.
3. Select **Use default certificate**.

When you restore the default management interface certificate, the custom server certificate files you configured are deleted and cannot 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.

What you'll need

- 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.

1. 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.

2. 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.

1. 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.

2. Paste the copied certificate into a text editor.
3. 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, the deprecated Connection Load Balancer (CLB) service on Gateway 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 are not 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.

1. Select **Upload certificate**.
2. 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.
3. 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.
4. Select **Save**.

The custom server certificate is used for subsequent new S3 and Swift client connections.

Generate certificate

Generate the server certificate files.

1. Select **Generate certificate**.
2. Specify the certificate information:
 - **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**: X.509 subject or distinguished name (DN) of the certificate owner.
 - **Days valid**: Number of days after creation that the certificate expires.
3. Select **Generate**.
4. 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.

5. Select **Save**.

The custom server certificate is used for subsequent new S3 and Swift client connections.



After uploading or generating a new certificate, allow up to one day for any related certificate expiration alerts to clear.

6. Refresh the page to ensure the web browser is updated.
7. After you add a custom S3 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 and to the deprecated CLB service on Gateway Nodes. However, you cannot 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 cannot 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 and to the deprecated CLB service on Gateway 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](#)
- [Use Swift](#)
- [Configure S3 API 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.

What you'll need

- 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.

1. Select **Download certificate**.
2. 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.

1. Select **Copy certificate PEM**.
2. Paste the copied certificate into a text editor.
3. 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 do not support disabling strict hostname validation, such as ONTAP clients using FabricPool, you can generate or upload a server certificate when you configure the load balancer endpoint.

What you'll need

- You 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](#).



The separate Connection Load Balancer (CLB) service on Gateway Nodes is deprecated and not recommended for use with FabricPool.

Steps

1. Optionally, configure a high availability (HA) group for FabricPool to use.

2. Create an S3 load balancer endpoint for FabricPool to use.

When you create an HTTPS load balancer endpoint, you are prompted to upload your server certificate, certificate private key, and 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 the information about [general security certificate use](#) and [configuring custom server certificates](#).



To ensure that operations are not disrupted by a failed server certificate, the **Expiration of client certificates configured on the Certificates page** alert is triggered when this server certificate is about to expire. As required, you can view when the current certificate expires by selecting **CONFIGURATION > Security > Certificates** and looking at the Expiration date for the client certificate on the Client tab.



If you are using a key management server (KMS) to protect the data on specially configured appliance nodes, see the specific information about [uploading a KMS client certificate](#).

What you'll need

- 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 do not have the original private key, you must create a new one.
- To edit a client certificate:
 - You have the IP address or domain name of the Admin Node.
 - To upload your own certificate or a new certificate, the private key, client certificate, and CA (if used) are available on your local computer.

Add client certificates

Follow the procedure for your scenario to add a client certificate:

- [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 that contains at least 1 and no more than 32 characters.
4. To access Prometheus metrics using your external monitoring tool, select **Allow Prometheus**.
5. In the **Certificate type** section, upload the management interface certificate `.pem` file.
 - a. Select **Upload certificate** and then select **Continue**.
 - b. Upload 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.

6. 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.

- f. Save and test the certificate and private key that you copied from StorageGRID or a local file.

You can now access the Prometheus metrics from StorageGRID with your external monitoring tool.

For information about the metrics, see the [instructions for monitoring StorageGRID](#).

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 that contains at least 1 and no more than 32 characters.
5. To access Prometheus metrics using your external monitoring tool, select **Allow Prometheus**.
6. In the **Certificate type** section, upload the client certificate, private key and CA bundle .pem files:
 - a. Select **Upload certificate** and then select **Continue**.
 - b. Upload 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.

7. 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.

- f. Save and test the certificate and private key that you copied from StorageGRID or a local file.

You can now access the Prometheus metrics from StorageGRID with your external monitoring tool.

For information about the metrics, see the [instructions for monitoring StorageGRID](#).

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 that contains at least 1 and no more than 32 characters.
4. To access Prometheus metrics using your external monitoring tool, select **Allow Prometheus**.
5. In the **Certificate type** section, select **Generate certificate**.
6. Specify the certificate information:
 - **Domain name:** One or more fully qualified domain names of the admin node to include in the certificate. Use an * as a wildcard to represent multiple domain names.
 - **IP:** One or more admin node IP addresses to include in the certificate.
 - **Subject:** X.509 subject or distinguished name (DN) of the certificate owner.
7. Select **Generate**.
8. 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.

9. Select **Create** to save the certificate in the Grid Manager.

The new certificate appears on the Client tab.

10. In the Grid Manager, select **CONFIGURATION > Security > Certificates** and then select the **Global** tab.
11. Select **Management Interface certificate**.
12. Select **Use custom certificate**.
13. 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 **Create** to save the certificate in the Grid Manager.

The new certificate appears on the Client tab.

14. 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 client certificate to both **CA Cert** and **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.

- f. Save and test the certificate and private key that you copied from StorageGRID or a local file.

You can now access the Prometheus metrics from StorageGRID with your external monitoring tool.

For information about the metrics, see the [instructions for monitoring 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 that contains at least 1 and no more than 32 characters.
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.

1. Select **Upload certificate** and then select **Continue**.
2. 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.

3. 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.

1. Select **Generate certificate**.
2. Specify the certificate information:
 - **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:** X.509 subject or distinguished name (DN) of the certificate owner.
 - **Days valid:** Number of days after creation that the certificate expires.
3. Select **Generate**.
4. 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.

5. 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.

1. Select **Download certificate**.
2. 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.

1. Select **Copy certificate PEM**.
2. Paste the copied certificate into a text editor.
3. 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 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.

What is a key management server (KMS)?

A key management server (KMS) is an external, third-party system that provides encryption keys to StorageGRID appliance nodes at the associated StorageGRID site using the Key Management Interoperability Protocol (KMIP).

You can use one or more key management servers to manage the node encryption keys for any StorageGRID appliance nodes that have the **Node Encryption** setting enabled during installation. Using key management servers with these appliance nodes lets you protect your data even if an appliance is removed from the data center. After the appliance volumes are encrypted, you cannot access any data on the appliance unless the node can communicate with the KMS.




StorageGRID does not create or manage the external keys used to encrypt and decrypt appliance nodes. If you plan to use an external key management server to protect StorageGRID data, you must understand how to set up that server, and you must understand how to manage the encryption keys. Performing key management tasks is beyond the scope of these instructions. If you need help, see the documentation for your key management server or contact technical support.

Review StorageGRID encryption methods

StorageGRID provides a number of options for encrypting data. You should review the available methods to determine which methods meet your data-protection requirements.

The table provides a high-level summary of the encryption methods available in StorageGRID.

Encryption option	How it works	Applies to
Key management server (KMS) in Grid Manager	You configure a key management server for the StorageGRID site (CONFIGURATION > Security > Key management server) 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.  Managing encryption keys with a KMS is only supported for Storage Nodes and service appliances.

Encryption option	How it works	Applies to
Drive security in SANtricity System Manager	If the Drive Security feature is enabled for a storage appliance, you can use SANtricity System Manager to create and manage the security key. The key is required to access the data on the secured drives.	<p>Storage appliances that have Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives. All data on the secured drives is protected against physical loss or removal from the data center. Cannot be used with some storage appliances or with any service appliances.</p> <ul style="list-style-type: none"> • SG6000 storage appliances • SG5700 storage appliances • SG5600 storage appliances
Stored Object Encryption grid option	The Stored Object Encryption option can be enabled in the Grid Manager (CONFIGURATION > System > Grid options). When enabled, any new objects that are not encrypted at the bucket level or at the object level are encrypted during ingest.	<p>Newly ingested S3 and Swift object data.</p> <p>Existing stored objects are not encrypted. Object metadata and other sensitive data are not encrypted.</p> <ul style="list-style-type: none"> • Configure stored object encryption
S3 bucket encryption	You issue a PUT Bucket encryption request to enable encryption for the bucket. Any new objects that are not encrypted at the object level are encrypted during ingest.	<p>Newly ingested S3 object data only.</p> <p>Encryption must be specified for the bucket. Existing bucket objects are not encrypted. Object metadata and other sensitive data are not encrypted.</p> <ul style="list-style-type: none"> • Use S3
S3 object server-side encryption (SSE)	You issue an S3 request to store an object and include the <code>x-amz-server-side-encryption</code> request header.	<p>Newly ingested S3 object data only.</p> <p>Encryption must be specified for the object. Object metadata and other sensitive data are not encrypted.</p> <p>StorageGRID manages the keys.</p> <ul style="list-style-type: none"> • Use S3

Encryption option	How it works	Applies to
S3 object server-side encryption with customer-provided keys (SSE-C)	<p>You issue an S3 request to store an object and include three request headers.</p> <ul style="list-style-type: none"> • <code>x-amz-server-side-encryption-customer-algorithm</code> • <code>x-amz-server-side-encryption-customer-key</code> • <code>x-amz-server-side-encryption-customer-key-MD5</code> 	<p>Newly ingested S3 object data only.</p> <p>Encryption must be specified for the object. Object metadata and other sensitive data are not encrypted.</p> <p>Keys are managed outside of StorageGRID.</p> <ul style="list-style-type: none"> • Use S3
External volume or datastore encryption	<p>You use an encryption method outside of StorageGRID to encrypt an entire volume or datastore, if your deployment platform supports it.</p>	<p>All object data, metadata, and system configuration data, assuming every volume or datastore is encrypted.</p> <p>An external encryption method provides tighter control over encryption algorithms and keys. Can be combined with the other methods listed.</p>
Object encryption outside of StorageGRID	<p>You use an encryption method outside of StorageGRID to encrypt object data and metadata before they are ingested into StorageGRID.</p>	<p>Object data and metadata only (system configuration data is not encrypted).</p> <p>An external encryption method provides tighter control over encryption algorithms and keys. Can be combined with the other methods listed.</p> <ul style="list-style-type: none"> • 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 grid option to encrypt all objects when they are ingested.


If only a small portion of your objects require encryption, consider controlling encryption at the bucket or

individual object level instead. Enabling multiple levels of encryption has an additional performance cost.

Overview of KMS and appliance configuration

Before you can use a key management server (KMS) to secure StorageGRID data on appliance nodes, you must complete two configuration tasks: setting up one or more KMS servers and enabling node encryption for the appliance nodes. When these two configuration tasks are complete, the key management process occurs automatically.

The flowchart shows the high-level steps for using a KMS to secure StorageGRID data on appliance nodes.



The flowchart shows KMS setup and appliance setup occurring in parallel; however, you can set up the key

management servers before or after you enable node encryption for new appliance nodes, based on your requirements.

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 cannot enable the **Node Encryption** setting after an appliance is added to the grid, and you cannot use external key management for appliances that do not have node encryption enabled.

2. Run the StorageGRID Appliance Installer. During installation, a random data encryption key (DEK) is assigned to each appliance volume, as follows:
 - The DEKs are used to encrypt the data on each volume. These keys are generated using Linux Unified Key Setup (LUKS) disk encryption in the appliance OS and cannot be changed.
 - Each individual DEK is encrypted by a master key encryption key (KEK). The initial KEK is a temporary key that encrypts the DEKs until the appliance can connect to the KMS.
3. Add the appliance node to StorageGRID.

For details, refer to the following:

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

Key management encryption process (occurs automatically)

Key management encryption includes the following high-level steps that are performed automatically.

1. When you install an appliance that has node encryption enabled into the grid, StorageGRID determines if a

KMS configuration exists for the site that contains the new node.

- If a KMS has already been configured for the site, the appliance receives the KMS configuration.
 - If a KMS has not yet been configured for the site, data on the appliance continues to be encrypted by the temporary KEK until you configure a KMS for the site and the appliance receives the KMS configuration.
2. The appliance uses the KMS configuration to connect to the KMS and request an encryption key.
 3. The KMS sends an encryption key to the appliance. The new key from the KMS replaces the temporary KEK and is now used to encrypt and decrypt the DEKs for the appliance volumes.



Any data that exists before the encrypted appliance node connects to the configured KMS is encrypted with a temporary key. However, the appliance volumes should not be considered protected from removal from the data center until the temporary key is replaced by the KMS encryption key.

4. If the appliance is powered on or rebooted, it reconnects to the KMS to request the key. The key, which is saved in volatile memory, cannot survive a loss of power or a reboot.

Considerations and requirements for using a key management server

Before configuring an external key management server (KMS), you must understand the considerations and requirements.

What are the KMIP requirements?

StorageGRID supports KMIP version 1.4.

[Key Management Interoperability Protocol Specification Version 1.4](#)

Communications between the appliance nodes and the configured KMS use secure TLS connections. StorageGRID supports the following TLS v1.2 ciphers for KMIP:

- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384

You must ensure that each appliance node that uses node encryption has network access to the KMS or KMS cluster you configured for the site.

The network firewall settings must allow each appliance node to communicate through the port used for Key Management Interoperability Protocol (KMIP) communications. The default KMIP port is 5696.

Which appliances are supported?

You can use a key management server (KMS) to manage encryption keys for any StorageGRID appliance in your grid that has the **Node Encryption** setting enabled. This setting can only be enabled during the hardware configuration stage of appliance installation using the StorageGRID Appliance Installer.



You cannot enable node encryption after an appliance is added to the grid, and you cannot use external key management for appliances that do not have node encryption enabled.

You can use the configured KMS for the following StorageGRID appliances and appliance nodes:

Appliance	Node type
SG1000 services appliance	Admin Node or Gateway Node
SG100 services appliance	Admin Node or Gateway Node
SG6000 storage appliance	Storage Node
SG5700 storage appliance	Storage Node
SG5600 storage appliance	Storage Node

You cannot use the configured KMS for software-based (non-appliance) nodes, including the following:

- Nodes deployed as virtual machines (VMs)
- Nodes deployed within 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 cannot use a KMS for non-appliance nodes or for any appliance nodes that did not have the **Node Encryption** setting enabled during installation.



What happens when a key is rotated?

As a security best practice, you should periodically rotate the encryption key used by each configured KMS.

When rotating the encryption key, use the KMS software to rotate from the last used version of the key to a new version of the same key. Do not rotate to an entirely different key.



Never attempt to rotate a key by changing the key name (alias) for the KMS in the Grid Manager. Instead, rotate the key by updating the key version in the KMS software. Use the same key alias for new keys as was used for previous keys. If you change the key alias for a configured KMS, StorageGRID might not be able to decrypt your data.

When the new key version is available:

- It is automatically distributed to the encrypted appliance nodes at the site or sites associated with the KMS. The distribution should occur within an hour of when the key is rotated.
- If the encrypted appliance node is offline when the new key version is distributed, the node will receive the new key as soon as it reboots.
- If the new key version cannot be used to encrypt appliance volumes for any reason, the **KMS encryption key rotation failed** alert is triggered for the appliance node. You might need to contact technical support

for help in resolving this alert.

Can I reuse an appliance node after it has been encrypted?

If you need to install an encrypted appliance into another StorageGRID system, you must first decommission the grid node to move object data to another node. Then, you can use the StorageGRID Appliance Installer to clear the KMS configuration. Clearing the KMS configuration disables the **Node Encryption** setting and removes the association between the appliance node and the KMS configuration for the StorageGRID site.



With no access to the KMS encryption key, any data that remains on the appliance can no longer be accessed and is permanently locked.

Related information

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


Considerations for changing the KMS for a site

Each key management server (KMS) or KMS cluster provides an encryption key to all appliance nodes at a single site or at a group of sites. If you need to change which KMS is used for a site, you might need to copy the encryption key from one KMS to another.


If you change the KMS used for a site, you must ensure that the previously encrypted appliance nodes at that site can be decrypted using the key stored on the new KMS. In some cases, you might need to copy the current version of the encryption key from the original KMS to the new KMS. You must ensure that the KMS has the correct key to decrypt the encrypted appliance nodes at the site.

For example:


1. You initially configure a default KMS that applies to all sites that do not have a dedicated KMS.
2. When the KMS is saved, all appliance nodes that have the **Node Encryption** setting enabled connect to the KMS and request the encryption key. This key is used to encrypt the appliance nodes at all sites. This same key must also be used to decrypt those appliances.



3. You decide to add a site-specific KMS for one site (Data Center 3 in the figure). However, because the appliance nodes are already encrypted, a validation error occurs when you attempt to save the configuration for the site-specific KMS. The error occurs because the site-specific KMS does not have the correct key to decrypt the nodes at that site.



4. To address the issue, you copy the current version of the encryption key from the default KMS to the new KMS. (Technically, you copy the original key to a new key with the same alias. The original key becomes a prior version of the new key.) The site-specific KMS now has the correct key to decrypt the appliance nodes at Data Center 3, so it can be saved in StorageGRID.



Use cases for changing which KMS is used for a site

The table summarizes the required steps for the most common cases for changing the KMS for a site.

Use case for changing a site's KMS	Required steps
You have one or more site-specific KMS entries, and you want to use one of them as the default KMS.	<p>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 do not have a dedicated KMS.</p> <p>Edit a key management server (KMS)</p>
You have a default KMS and you add a new site in an expansion. You do not want to use the default KMS for the new site.	<ol style="list-style-type: none"> 1. If the appliance nodes at the new site have already been encrypted by the default KMS, use the KMS software to copy the current version of the encryption key from the default KMS to a new KMS. 2. Using the Grid Manager, add the new KMS and select the site. <p>Add a key management server (KMS)</p>
You want the KMS for a site to use a different server.	<ol style="list-style-type: none"> 1. If the appliance nodes at the site have already been encrypted by the existing KMS, use the KMS software to copy the current version of the encryption key from the existing KMS to the new KMS. 2. Using the Grid Manager, edit the existing KMS configuration and enter the new host name or IP address. <p>Add a key management server (KMS)</p>

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.

About this task

These instructions apply to Thales CipherTrust Manager k170v, versions 2.0, 2.1, and 2.2. If you have questions about using a different key management server with StorageGRID, contact technical support.

Thales CipherTrust Manager

Steps

1. From the KMS software, create a StorageGRID client for each KMS or KMS cluster you plan to use.

Each KMS manages a single encryption key for the StorageGRID appliances nodes at a single site or at a group of sites.

2. From the KMS software, create an AES encryption key for each KMS or KMS cluster.

The encryption key needs to be exportable.

3. Record the following information for each KMS or KMS cluster.

You need this information when you add the KMS to StorageGRID.

- Host name or IP address for each server.
- KMIP port used by the KMS.
- Key alias for the encryption key in the KMS.



The encryption key must already exist in the KMS. StorageGRID does not create or manage KMS keys.

4. For each KMS or KMS cluster, obtain a server certificate signed by a certificate authority (CA) or a certificate bundle that contains each of the PEM-encoded CA certificate files, concatenated in certificate chain order.

The server certificate allows the external KMS to authenticate itself to StorageGRID.

- The certificate must use the Privacy Enhanced Mail (PEM) Base-64 encoded X.509 format.
- The Subject Alternative Name (SAN) field in each server certificate must include the fully qualified domain name (FQDN) or IP address that StorageGRID will connect to.



When you configure the KMS in StorageGRID, you must enter the same FQDNs or IP addresses in the **Hostname** field.

- The server certificate must match the certificate used by the KMIP interface of the KMS, which typically uses port 5696.
5. Obtain the public client certificate issued to StorageGRID by the external KMS and the private key for the client certificate.

The client certificate allows StorageGRID to authenticate itself to the KMS.

Add a key management server (KMS)

You use the StorageGRID Key Management Server wizard to add each KMS or KMS cluster.

What you'll need

- You 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: Enter KMS Details

In Step 1 (Enter KMS Details) of the Add a Key Management Server wizard, you provide details about the KMS or KMS cluster.

Steps

1. Select **CONFIGURATION > Security > Key management server**.

The Key Management Server page appears with the Configuration Details tab selected.

Key Management Server

If your StorageGRID system includes appliance nodes with node encryption enabled, you can use an external key management server (KMS) to manage the encryption keys that protect your StorageGRID at rest.

Configuration Details [Encrypted Nodes](#)

You can configure more than one KMS (or KMS cluster) to manage the encryption keys for appliance nodes. For example, you can configure one default KMS to manage the keys for all appliance nodes within a group of sites and a second KMS to manage the keys for the appliance nodes at a particular site.

Before adding a KMS:

- Ensure that the KMS is KMIP-compliant.
- Configure StorageGRID as a client in the KMS.
- Enable node encryption for each appliance during appliance installation. You cannot enable node encryption after an appliance is added to the grid and you cannot use a KMS for appliances that do not have node encryption enabled.

For complete instructions, see [administering StorageGRID](#).

[+ Create](#) [Edit](#) [Remove](#)

KMS Display Name ?	Key Name ?	Manages keys for ?	Hostname ?	Certificate Status ?
No key management servers have been configured. Select Create .				

2. Select **Create**.

Step 1 (Enter KMS Details) of the Add a Key Management Server wizard appears.

Add a Key Management Server



Enter information about the external key management server (KMS) and the StorageGRID client you configured in that KMS. If you are configuring a KMS cluster, select + to add a hostname for each server in the cluster.

KMS Display Name	<input type="text"/>
Key Name	<input type="text"/>
Manages keys for	<input type="text" value="-- Choose One --"/>
Port	<input type="text" value="5696"/>
Hostname	<input type="text"/>

+

Cancel

Next

3. Enter the following information for the KMS and the StorageGRID client you configured in that KMS.

Field	Description
KMS Display Name	A descriptive name to help you identify this KMS. Must be between 1 and 64 characters.
Key Name	The exact key alias for the StorageGRID client in the KMS. Must be between 1 and 255 characters.

Field	Description
Manages keys for	<p>The StorageGRID site that will be associated with this KMS. If possible, you should configure any site-specific key management servers before configuring a default KMS that applies to all sites not managed by another KMS.</p> <ul style="list-style-type: none"> • 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 do not have a dedicated KMS and to any sites you add in subsequent expansions. <p>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.</p>
Port	<p>The port the KMS server uses for Key Management Interoperability Protocol (KMIP) communications. Defaults to 5696, which is the KMIP standard port.</p>
Hostname	<p>The fully qualified domain name or IP address for the KMS.</p> <p>Note: The SAN field of the server certificate must include the FQDN or IP address you enter here. Otherwise, StorageGRID will not be able to connect to the KMS or to all servers in a KMS cluster.</p>

4. If you are using a KMS cluster, select the plus sign **+** to add a hostname for each server in the cluster.
5. Select **Next**.

Step 2: Upload Server Certificate

In Step 2 (Upload Server Certificate) of the Add a Key Management Server wizard, you upload the server certificate (or certificate bundle) for the KMS. The server certificate allows the external KMS to authenticate itself to StorageGRID.

Steps

1. From **Step 2 (Upload Server Certificate)**, browse to the location of the saved server certificate or certificate bundle.

Add a Key Management Server



Upload a server certificate signed by the certificate authority (CA) on the external key management server (KMS) or a certificate bundle. The server certificate allows the KMS to authenticate itself to StorageGRID.

Server Certificate ?

Browse

Cancel

Back

Next


2. Upload the certificate file.

The server certificate metadata appears.

Add a Key Management Server



Upload a server certificate signed by the certificate authority (CA) on the external key management server (KMS) or a certificate bundle. The server certificate allows the KMS to authenticate itself to StorageGRID.

Server Certificate  k170vCA.pem

Server Certificate Metadata

```
Server DN: /C=US/ST=MD/L=Belcamp/O=Gemalto/CN=KeySecure Root CA
Serial Number: 71:CD:6D:72:53:B5:6D:0A:8C:69:13:0D:4D:D7:81:0E
Issue DN: /C=US/ST=MD/L=Belcamp/O=Gemalto/CN=KeySecure Root CA
Issued On: 2020-10-15T21:12:45.000Z
Expires On: 2030-10-13T21:12:45.000Z
SHA-1 Fingerprint: EE:E4:6E:17:86:DF:56:B4:F5:AF:A2:3C:BD:56:6B:10:DB:B2:5A:79
```

Cancel

Back

Next



If you uploaded a certificate bundle, the metadata for each certificate appears on its own tab.

3. Select **Next**.

Step 3: Upload Client Certificates

In Step 3 (Upload Client Certificates) of the Add a Key Management Server wizard, you upload the client certificate and the client certificate private key. The client certificate allows StorageGRID to authenticate itself to the KMS.

Steps

1. From **Step 3 (Upload Client Certificates)**, browse to the location of the client certificate.

Add a Key Management Server



Upload the client certificate and the client certificate private key. The client certificate is issued to StorageGRID by the external key management server (KMS), and it allows StorageGRID to authenticate itself to the KMS.

Client Certificate ?

Browse

Client Certificate Private Key ?

Browse

Cancel

Back

Save

2. Upload the client certificate file.

The client certificate metadata appears.

3. Browse to the location of the private key for the client certificate.

4. Upload the private key file.

The metadata for the client certificate and the client certificate private key appear.

Add a Key Management Server



Upload the client certificate and the client certificate private key. The client certificate is issued to StorageGRID by the external key management server (KMS), and it allows StorageGRID to authenticate itself to the KMS.

Client Certificate ?

Browse

k170vClientCert.pem

Server DN: /CN=admin/UID=
Serial Number: 7D:5A:8A:27:02:40:C8:F5:19:A1:28:22:E7:D6:E2:EB
Issue DN: /C=US/ST=MD/L=Belcamp/O=Gemalto/CN=KeySecure Root CA
Issued On: 2020-10-15T23:31:49.000Z
Expires On: 2022-10-15T23:31:49.000Z
SHA-1 Fingerprint: A7:10:AC:39:85:42:80:8F:FF:62:AD:A1:BD:CF:4C:90:F3:E9:36:69

Client Certificate Private Key ?

Browse

k170vClientKey.pem

Cancel

Back

Save

5. Select **Save**.

The connections between the key management server and the appliance nodes are tested. If all connections are valid and the correct key is found on the KMS, the new key management server is added to the table on the Key Management Server page.



Immediately after you add a KMS, the certificate status on the Key Management Server page appears as Unknown. It might take StorageGRID as long as 30 minutes to get the actual status of each certificate. You must refresh your web browser to see the current status.

6. If an error message appears when you select **Save**, review the message details and then select **OK**.

For example, you might receive a 422: Unprocessable Entity error if a connection test failed.

7. If you need to save the current configuration without testing the external connection, select **Force Save**.

Add a Key Management Server



Upload the client certificate and the client certificate private key. The client certificate is issued to StorageGRID by the external key management server (KMS), and it allows StorageGRID to authenticate itself to the KMS.

Client Certificate ?

Browse

k170vClientCert.pem

Server DN: /CN=admin/UID=

Serial Number: 7D:5A:8A:27:02:40:C8:F5:19:A1:28:22:E7:D6:E2:EB

Issue DN: /C=US/ST=MD/L=Belcamp/O=Gemalto/CN=KeySecure Root CA

Issued On: 2020-10-15T23:31:49.000Z

Expires On: 2022-10-15T23:31:49.000Z

SHA-1 Fingerprint: A7:10:AC:39:85:42:80:8F:FF:62:AD:A1:BD:CF:4C:90:F3:E9:36:69

Client Certificate Private Key ?

Browse

k170vClientKey.pem

Select **Force Save** to save this KMS without testing the external connections. If there is an issue with the configuration, you might not be able to reboot any FDE-enabled appliance nodes at the affected site, and you might lose access to your data.

Cancel

Back

Force Save

Save



Selecting **Force Save** saves the KMS configuration, but it does not test the external connection from each appliance to that KMS. If there is an issue with the configuration, you might not be able to reboot appliance nodes that have node encryption enabled at the affected site. You might lose access to your data until the issues are resolved.

8. Review the confirmation warning, and select **OK** if you are sure you want to force save the configuration.

Warning

Confirm force-saving the KMS configuration

Are you sure you want to save this KMS without testing the external connections?

If there is an issue with the configuration, you might not be able to reboot any appliance nodes with node encryption enabled at the affected site, and you might lose access to your data.

Cancel

OK

The KMS configuration is saved but the connection to the KMS is not tested.

View KMS details

You can view information about each key management server (KMS) in your StorageGRID system, including the current status of the server and client certificates.

Steps

1. Select **CONFIGURATION > Security > Key management server**.

The Key Management Server page appears. The Configuration Details tab shows any key management servers that are configured.

Key Management Server

If your StorageGRID system includes appliance nodes with node encryption enabled, you can use an external key management server (KMS) to manage the encryption keys that protect your StorageGRID at rest.

Configuration Details

Encrypted Nodes

You can configure more than one KMS (or KMS cluster) to manage the encryption keys for appliance nodes. For example, you can configure one default KMS to manage the keys for all appliance nodes within a group of sites and a second KMS to manage the keys for the appliance nodes at a particular site.

Before adding a KMS:

- Ensure that the KMS is KMIP-compliant.
- Configure StorageGRID as a client in the KMS.
- Enable node encryption for each appliance during appliance installation. You cannot enable node encryption after an appliance is added to the grid and you cannot use a KMS for appliances that do not have node encryption enabled.

For complete instructions, see [administering StorageGRID](#).

+ Create Edit Remove

KMS Display Name ?	Key Name ?	Manages keys for ?	Hostname ?	Certificate Status ?
Default KMS	test	Sites not managed by another KMS (default KMS)	10.96.99.164	✓ All certificates are valid

2. Review the information in the table for each KMS.

Field	Description
KMS Display Name	The descriptive name of the KMS.
Key Name	The key alias for the StorageGRID client in the KMS.
Manages keys for	The StorageGRID site associated with the KMS. This field displays the name of a specific StorageGRID site or Sites not managed by another KMS (default KMS) .

Field	Description
Hostname	<p>The fully qualified domain name or IP address of the KMS.</p> <p>If there is a cluster of two key management servers, the fully qualified domain name or IP address of both servers are listed. If there are more than two key management servers in a cluster, the fully qualified domain name or IP address of the first KMS is listed along with the number of additional key management servers in the cluster.</p> <p>For example: 10.10.10.10 and 10.10.10.11 or 10.10.10.10 and 2 others.</p> <p>To view all hostnames in a cluster, select a KMS and then select Edit.</p>
Certificate Status	<p>Current state of the server certificate, optional CA certificate, and the client certificate: valid, expired, nearing expiration, or unknown.</p> <p>Note: It might take StorageGRID as long as 30 minutes to get updates to the certificate status. You must refresh your web browser to see the current values.</p>

- If the Certificate Status is Unknown, wait up to 30 minutes and then refresh your web browser.



Immediately after you add a KMS, the certificate status on the Key Management Server page appears as Unknown. It might take StorageGRID as long as 30 minutes to get the actual status of each certificate. You must refresh your web browser to see the actual status.

- If the Certificate Status column indicates that a certificate has expired or is nearing expiration, address the issue as soon as possible.

See the recommended actions for the **KMS CA certificate expiration**, **KMS client certificate expiration**, and **KMS server certificate expiration** alerts in the instructions for [monitoring and troubleshooting StorageGRID](#).



You must address any certificate issues as soon as possible to maintain data access.

View encrypted nodes

You can view information about the appliance nodes in your StorageGRID system that have the **Node Encryption** setting enabled.

Steps

- 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.

Key Management Server

If your StorageGRID system includes appliance nodes with node encryption enabled, you can use an external key management server (KMS) to manage the encryption keys that protect your StorageGRID at rest.

Configuration Details

Encrypted Nodes

You can configure more than one KMS (or KMS cluster) to manage the encryption keys for appliance nodes. For example, you can configure one default KMS to manage the keys for all appliance nodes within a group of sites and a second KMS to manage the keys for the appliance nodes at a particular site.

Before adding a KMS:

- Ensure that the KMS is KMIP-compliant.
- Configure StorageGRID as a client in the KMS.
- Enable node encryption for each appliance during appliance installation. You cannot enable node encryption after an appliance is added to the grid and you cannot use a KMS for appliances that do not have node encryption enabled.

For complete instructions, see [administering StorageGRID](#).

+ Create

Edit

Remove

KMS Display Name ?	Key Name ?	Manages keys for ?	Hostname ?	Certificate Status ?
Default KMS	test	Sites not managed by another KMS (default KMS)	10.96.99.164	✓ All certificates are valid

2. From the top of the page, select the **Encrypted Nodes** tab.

Key Management Server

If your StorageGRID system includes appliance nodes with Full Disk Encryption (FDE) enabled, you can use an external key management server (KMS) to manage the encryption keys that protect your StorageGRID data at rest.

Configuration Details

Encrypted Nodes

You can configure more than one KMS (or KMS cluster) to manage the encryption keys for appliance nodes. For example, you can configure one default KMS to manage the keys for all appliance nodes within a group of sites and a second KMS to manage the keys for the appliance nodes at a particular site.

The Encrypted Nodes tab lists the appliance nodes in your StorageGRID system that have the **Node Encryption** setting enabled.

Configuration Details

Encrypted Nodes

Review the KMS status for all appliance nodes that have node encryption enabled. Address any issues immediately to ensure your data is fully protected. If no KMS exists for a site, select Configuration Details and add a KMS.

Nodes with Encryption Enabled

Node Name	Node Type	Site	KMS Display Name ?	Key UID ?	Status ?
SGA-010-096-104-67	Storage Node	Data Center 1	Default KMS	41b0...5c57	✓ Connected to KMS (2021-03-12 10:59:32 MST)

3. Review the information in the table for each appliance node.

Column	Description
Node Name	The name of the appliance node.
Node Type	The type of node: Storage, Admin, or Gateway.

Column	Description
Site	The name of the StorageGRID site where the node is installed.
KMS Display Name	<p>The descriptive name of the KMS used for the node.</p> <p>If no KMS is listed, select the Configuration Details tab to add a KMS.</p> <p>Add a key management server (KMS)</p>
Key UID	<p>The unique ID of the encryption key used to encrypt and decrypt data on the appliance node. To view an entire key UID, hover your cursor over the cell.</p> <p>A dash (--) indicates the key UID is unknown, possibly because of a connection issue between the appliance node and the KMS.</p>
Status	<p>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.</p> <p>Note: You must refresh your web browser to see the new values.</p>

- If the Status column indicates a KMS issue, address the issue immediately.

During normal KMS operations, the status will be **Connected to KMS**. If a node is disconnected from the grid, the node connection state is shown (Administratively Down or Unknown).

Other status messages correspond to StorageGRID alerts with the same names:

- KMS configuration failed to load
- KMS connectivity error
- KMS encryption key name not found
- KMS encryption key rotation failed
- KMS key failed to decrypt an appliance volume
- KMS is not configured

See the recommended actions for these alerts in the instructions for [monitoring and troubleshooting StorageGRID](#).



You must address any issues immediately to ensure that your data is fully protected.

Edit a key management server (KMS)

You might need to edit the configuration of a key management server, for example, if a certificate is about to expire.

What you'll need

- You have reviewed the [considerations and requirements for using a key management server](#).
- 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.

Key Management Server

If your StorageGRID system includes appliance nodes with node encryption enabled, you can use an external key management server (KMS) to manage the encryption keys that protect your StorageGRID at rest.

Configuration Details

Encrypted Nodes

You can configure more than one KMS (or KMS cluster) to manage the encryption keys for appliance nodes. For example, you can configure one default KMS to manage the keys for all appliance nodes within a group of sites and a second KMS to manage the keys for the appliance nodes at a particular site.

Before adding a KMS:

- Ensure that the KMS is KMIP-compliant.
- Configure StorageGRID as a client in the KMS.
- Enable node encryption for each appliance during appliance installation. You cannot enable node encryption after an appliance is added to the grid and you cannot use a KMS for appliances that do not have node encryption enabled.

For complete instructions, see [administering StorageGRID](#).

+ Create


Edit


Remove

KMS Display Name ?	Key Name ?	Manages keys for ?	Hostname ?	Certificate Status ?
<input checked="" type="radio"/> Default KMS	test	Sites not managed by another KMS (default KMS)	10.96.99.164	✓ All certificates are valid

2. Select the KMS you want to edit, and select **Edit**.
3. Optionally, update the details in **Step 1 (Enter KMS Details)** of the Edit a Key Management Server wizard.

Field	Description
KMS Display Name	A descriptive name to help you identify this KMS. Must be between 1 and 64 characters.

Field	Description
Key Name	<p>The exact key alias for the StorageGRID client in the KMS. Must be between 1 and 255 characters.</p> <p>You only need to edit the key name in rare cases. For example, you must edit the key name if the alias is renamed in the KMS or if all versions of the previous key have been copied to the version history of the new alias.</p> <div>  <p>Never attempt to rotate a key by changing the key name (alias) for the KMS. Instead, rotate the key by updating the key version in the KMS software. StorageGRID requires all previously used key versions (as well as any future ones) to be accessible from the KMS with the same key alias. If you change the key alias for a configured KMS, StorageGRID might not be able to decrypt your data.</p> <p>Considerations and requirements for using a key management server</p> </div>
Manages keys for	<p>If you are editing a site-specific KMS and you do not already have a default KMS, optionally select 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 do not have a dedicated KMS and to any sites added in an expansion.</p> <p>Note: If you are editing a site-specific KMS, you cannot select another site. If you are editing the default KMS, you cannot select a specific site.</p>
Port	The port the KMS server uses for Key Management Interoperability Protocol (KMIP) communications. Defaults to 5696, which is the KMIP standard port.
Hostname	<p>The fully qualified domain name or IP address for the KMS.</p> <p>Note: The SAN field of the server certificate must include the FQDN or IP address you enter here. Otherwise, StorageGRID will not be able to connect to the KMS or to all servers in a KMS cluster.</p>

- If you are configuring a KMS cluster, select the plus sign  to add a hostname for each server in the cluster.
- Select **Next**.

Step 2 (Upload Server Certificate) of the Edit a Key Management Server wizard appears.

- If you need to replace the server certificate, select **Browse** and upload the new file.
- Select **Next**.

Step 3 (Upload Client Certificates) of the Edit a Key Management Server wizard appears.

- If you need to replace the client certificate and the client certificate private key, select **Browse** and upload the new files.

9. Select **Save**.

The connections between the key management server and all node-encrypted appliance nodes at the affected sites are tested. If all node connections are valid and the correct key is found on the KMS, the key management server is added to the table on the Key Management Server page.

10. If an error message appears, review the message details, and select **OK**.

For example, you might receive a 422: Unprocessable Entity error if the site you selected for this KMS is already managed by another KMS, or if a connection test failed.

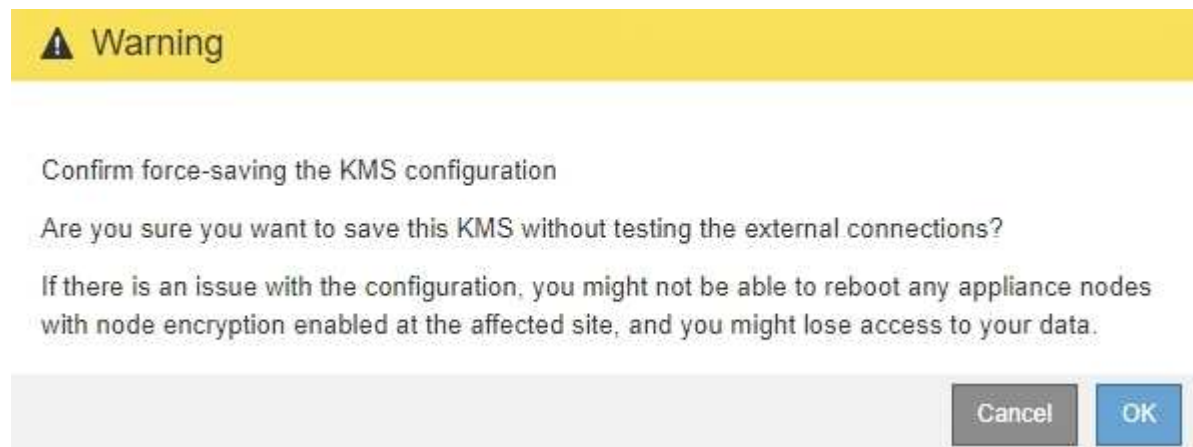
11. If you need to save the current configuration before resolving the connection errors, select **Force Save**.



Selecting **Force Save** saves the KMS configuration, but it does not test the external connection from each appliance to that KMS. If there is an issue with the configuration, you might not be able to reboot appliance nodes that have node encryption enabled at the affected site. You might lose access to your data until the issues are resolved.

The KMS configuration is saved.

12. Review the confirmation warning, and select **OK** if you are sure you want to force save the configuration.



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.

What you'll need

- 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.

Key Management Server

If your StorageGRID system includes appliance nodes with node encryption enabled, you can use an external key management server (KMS) to manage the encryption keys that protect your StorageGRID at rest.

Configuration Details

Encrypted Nodes

You can configure more than one KMS (or KMS cluster) to manage the encryption keys for appliance nodes. For example, you can configure one default KMS to manage the keys for all appliance nodes within a group of sites and a second KMS to manage the keys for the appliance nodes at a particular site.

Before adding a KMS:

- Ensure that the KMS is KMIP-compliant.
- Configure StorageGRID as a client in the KMS.
- Enable node encryption for each appliance during appliance installation. You cannot enable node encryption after an appliance is added to the grid and you cannot use a KMS for appliances that do not have node encryption enabled.

For complete instructions, see [administering StorageGRID](#).

[+ Create](#) [✎ Edit](#) [🗑 Remove](#)

KMS Display Name ?	Key Name ?	Manages keys for ?	Hostname ?	Certificate Status ?
<input checked="" type="radio"/> Default KMS	test	Sites not managed by another KMS (default KMS)	10.96.99.164	✓ All certificates are valid

2. Select the radio button for the KMS you want to remove, and select **Remove**.
3. Review the considerations in the warning dialog.

Warning

Delete KMS Configuration

You can only remove a KMS in these cases:

- You are removing a site-specific KMS for a site that has no appliance nodes with node encryption enabled.
- You are removing the default KMS, but a site-specific KMS already exists for each site with node encryption.

Are you sure you want to delete the Default KMS KMS configuration?

Cancel

OK

4. Select **OK**.

The KMS configuration is removed.

Manage proxy settings

Configure Storage proxy settings

If you are using platform services or Cloud Storage Pools, you can configure a non-transparent proxy between Storage Nodes and the external S3 endpoints. For example, you might need a non-transparent proxy to allow platform services messages to be sent to external endpoints, such as an endpoint on the internet.

What you'll need

- You 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**.

The Storage Proxy Settings page appears. By default, **Storage** is selected in the sidebar menu.



2. Select the **Enable Storage Proxy** check box.

The fields for configuring a Storage proxy appear.

Storage Proxy Settings

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

Enable Storage Proxy ☒

Protocol ☒ HTTP ☐ SOCKS5

Hostname

Port (optional)

3. Select the protocol for the non-transparent Storage proxy.
4. Enter the hostname or IP address of the proxy server.
5. Optionally, enter the port used to connect to the proxy server.

You can leave this field blank if you use the default port for the protocol: 80 for HTTP or 1080 for SOCKS5.

6. 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.

After you finish

If you need to disable a Storage proxy, deselect the **Enable Storage Proxy** check box, and select **Save**.

Related information

- [Network and ports for platform services](#)
- [Manage objects with ILM](#)

Configure Admin proxy settings

If you send AutoSupport messages using HTTP or HTTPS (see [Configure AutoSupport](#)), you can configure a non-transparent proxy server between Admin Nodes and technical support (AutoSupport).

What you'll need

- 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 Admin Proxy Settings page appears. By default, **Storage** is selected in the sidebar menu.

2. From the sidebar menu, select **Admin**.



3. Select the **Enable Admin Proxy** check box.

Admin Proxy Settings

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

Enable Admin Proxy ☒

Hostname

Port

Username (optional)

Password (optional)

4. Enter the hostname or IP address of the proxy server.

5. Enter the port used to connect to the proxy server.

6. Optionally, enter the proxy username.

Leave this field blank if your proxy server does not require a username.

7. Optionally, enter the proxy password.

Leave this field blank if your proxy server does not require a password.

8. 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 proxy, deselect the **Enable Admin Proxy** check box, and select **Save**.

Manage untrusted Client Networks

Manage untrusted Client Networks: Overview

If you are using a Client Network, you can help secure StorageGRID from hostile attacks by accepting inbound client traffic only on explicitly configured endpoints.

By default, the Client Network on each grid node is *trusted*. That is, by default, StorageGRID trusts inbound connections to each grid node on all available external ports (see the information about external communications in the [Networking guidelines](#)).

You can reduce the threat of hostile attacks on your StorageGRID system by specifying that the Client Network on each node be *untrusted*. If a node's Client Network is untrusted, the node only accepts inbound connections on ports explicitly configured as load balancer endpoints. See [Configure load balancer endpoints](#).

Example 1: Gateway Node only accepts HTTPS S3 requests

Suppose you want a Gateway Node to refuse all inbound traffic on the Client Network except for HTTPS S3 requests. You would perform these general steps:

1. From the Load Balancer Endpoints page, configure a load balancer endpoint for S3 over HTTPS on port 443.
2. From the Untrusted Client Networks page, specify that the Client Network on the Gateway Node is untrusted.

After you save your configuration, all inbound traffic on the Gateway Node's Client Network is dropped except for HTTPS S3 requests on port 443 and ICMP echo (ping) requests.

Example 2: Storage Node sends S3 platform services requests

Suppose you want to enable outbound S3 platform service traffic from a Storage Node, but you want to prevent any inbound connections to that Storage Node on the Client Network. You would perform this general step:

- From the Untrusted Client Networks page, indicate that the Client Network on the Storage Node is untrusted.

After you save your configuration, the Storage Node no longer accepts any incoming traffic on the Client Network, but it continues to allow outbound requests to Amazon Web Services.

Specify node's Client Network is untrusted

If you are using a Client Network, you can specify whether each node's Client Network is trusted or untrusted. You can also specify the default setting for new nodes added in an expansion.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the Root access permission.
- If you want an Admin Node or Gateway Node to accept inbound traffic only on explicitly configured endpoints, you have defined the load balancer endpoints.



Existing client connections might fail if load balancer endpoints have not been configured.

Steps

1. Select **CONFIGURATION > Security > Untrusted Client Networks**.

The Untrusted Client Networks page lists all nodes in your StorageGRID system. The Unavailable Reason column includes an entry if the Client Network on the node must be trusted.

Untrusted Client Networks

If you are using a Client Network, you can specify whether a node trusts inbound traffic from the Client Network. If the Client Network is untrusted, the node only accepts inbound traffic on ports configured as [load balancer endpoints](#).

Set New Node Default

This setting applies to new nodes expanded into the grid.

New Node Client Network ☒ Trusted
Default ☐ Untrusted

Select Untrusted Client Network Nodes

Select nodes that should have untrusted Client Network enforcement.

<input type="checkbox"/>	Node Name	Unavailable Reason
<input type="checkbox"/>	DC1-ADM1	
<input type="checkbox"/>	DC1-G1	
<input type="checkbox"/>	DC1-S1	
<input type="checkbox"/>	DC1-S2	
<input type="checkbox"/>	DC1-S3	
<input type="checkbox"/>	DC1-S4	

Client Network untrusted on 0 nodes.

Save

- In the **Set New Node Default** section, specify what the default setting should be when new nodes are added to the grid in an expansion procedure.
 - Trusted:** When a node is added in an expansion, its Client Network is trusted.
 - Untrusted:** When a node is added in an expansion, its Client Network is untrusted. As required, you can return to this page to change the setting for a specific new node.



This setting does not affect the existing nodes in your StorageGRID system.

- In the **Select Untrusted Client Network Nodes** section, select the nodes that should allow client connections only on explicitly configured load balancer endpoints.

You can select or unselect the check box in the title to select or unselect all nodes.

- Select **Save**.

The new firewall rules are immediately added and enforced. Existing client connections might fail if load balancer endpoints have not been configured.

Manage tenants

Manage tenants

As a grid administrator, you create and manage the tenant accounts that S3 and Swift

clients use to store and retrieve objects, monitor storage usage, and manage the actions that clients are able to perform using your StorageGRID system.

What are tenant accounts?

Tenant accounts allow client applications that use the Simple Storage Service (S3) REST API or the Swift REST API to store and retrieve objects on StorageGRID.

Each tenant account supports the use of a single protocol, which you specify when you create the account. To store and retrieve objects to a StorageGRID system with both protocols, you must create two tenant accounts: one for S3 buckets and objects, and one for Swift containers and objects. Each tenant account has its own account ID, authorized groups and users, buckets or containers, and objects.

Optionally, you can create additional tenant accounts if you want to segregate the objects stored on your system by different entities. For example, you might set up multiple tenant accounts in either of these use cases:

- **Enterprise use case:** If you are administering a StorageGRID system in an enterprise application, you might want to segregate the grid's object storage by the different departments in your organization. In this case, you could create tenant accounts for the Marketing department, the Customer Support department, the Human Resources department, and so on.



If you use the S3 client protocol, you can simply use S3 buckets and bucket policies to segregate objects between the departments in an enterprise. You do not need to use tenant accounts. See the instructions for implementing S3 client applications for more information.

- **Service provider use case:** If you are administering a StorageGRID system as a service provider, you can segregate the grid's object storage by the different entities that will lease the storage on your grid. In this case, you would create tenant accounts for Company A, Company B, Company C, and so on.

Create and configure tenant accounts

When you create a tenant account, you specify the following information:

- Display name for the tenant account.
- Which client protocol will be used by the tenant account (S3 or Swift).
- For S3 tenant accounts: Whether the tenant account has permission to use platform services with S3 buckets. If you permit tenant accounts to use platform services, you must ensure that the grid is configured to support their use. See “Managing platform services.”
- Optionally, a storage quota for the tenant account—the maximum number of gigabytes, terabytes, or petabytes available for the tenant's objects. If the quota is exceeded, the tenant cannot create new objects.



A tenant's storage quota represents a logical amount (object size), not a physical amount (size on disk).

- If identity federation is enabled for the StorageGRID system, which federated group has Root access permission to configure the tenant account.
- If single sign-on (SSO) is not in use for the StorageGRID system, whether the tenant account will use its own identity source or share the grid's identity source, and the initial password for the tenant's local root user.

After a tenant account is created, you can perform the following tasks:

- **Manage platform services for the grid:** If you enable platform services for tenant accounts, ensure that you understand how platform services messages are delivered and the networking requirements that the use of platform services place on your StorageGRID deployment.
- **Monitor a tenant account's storage usage:** After tenants begin using their accounts, you can use Grid Manager to monitor how much storage each tenant consumes.



A tenant's storage usage values might become out of date if nodes are isolated from other nodes in the grid. The totals will be updated when network connectivity is restored.

If you have set quotas for tenants, you can enable the **Tenant quota usage high** alert to determine if tenants are consuming their quotas. If enabled, this alert is triggered when a tenant has used 90% of its quota. For more information, see the alerts reference in the instructions for monitoring and troubleshooting StorageGRID.

- **Configure client operations:** You can configure if some types of client operations are forbidden.

Configure S3 tenants

After an S3 tenant account is created, tenant users can access the Tenant Manager to perform tasks such as the following:

- Setting up identity federation (unless the identity source is shared with the grid) and creating local groups and users
- Managing S3 access keys
- Creating and managing S3 buckets
- Monitoring storage usage
- Using platform services (if enabled)



S3 tenant users can create and manage S3 access key and buckets with the Tenant Manager, but they must use an S3 client application to ingest and manage objects.

Configure Swift tenants

After a Swift tenant account is created, the tenant's root user can access the Tenant Manager to perform tasks such as the following:

- Setting up identity federation (unless the identity source is shared with the grid), and creating local groups and users
- Monitoring storage usage



Swift users must have the Root access permission to access the Tenant Manager. However, the Root access permission does not allow users to authenticate into the Swift REST API to create containers and ingest objects. Users must have the Swift Administrator permission to authenticate into the Swift REST API.

Related information

[Use a tenant account](#)

Create tenant account

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

When you create a tenant account, you specify a name, a client protocol, and optionally a storage quota. If single sign-on (SSO) is enabled for StorageGRID, you also specify which federated group has Root access permission to configure the tenant account. If StorageGRID is not using single sign-on, you must also specify whether the tenant account will use its own identity source and configure the initial password for the tenant's local root user.

The Grid Manager provides a wizard that takes you through the steps for creating a tenant account. The steps 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.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.
- 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 do not need to assign any Grid Manager permissions to this admin group. See the [instructions for managing admin groups](#).

Steps

1. Select **TENANTS**.
2. Select **Create** and enter the following information for the tenant:
 - a. **Name**: Enter a name for the tenant account. Tenant names do not have to be unique. When the tenant account is created, it receives a unique, numeric account ID.
 - b. **Description** (optional): Enter a description that helps you identify the tenant.
 - c. **Client type**: Select the client type of either **S3** or **Swift**.
 - d. **Storage quota** (optional): If you want this tenant to have a storage quota, enter a numerical value for the quota and select the correct units (GB, TB, or PB).

Create a tenant

1 Enter details

2 Select permissions

3 Define root access

Enter tenant details

Name [?](#)

Description (optional) [?](#)

Client type [?](#)

☒ S3 ☐ Swift

Storage quota (optional) [?](#)

GB [v](#)

[Cancel](#)[Continue](#)

3. Select **Continue** and configure the S3 or Swift tenant.

S3 tenant

Select the appropriate permissions for the tenant. Some of these permissions have additional requirements. For details, see the online help for each permission.

- Allow platform services
- Use own identity source (selectable only if SSO is not being used)
- Allow S3 select (see [Manage S3 Select for tenant accounts](#))

Swift tenant

If the tenant will use its own identity source, select **Use own identity source** (selectable only if SSO is not being used).

4. Select **Continue** and define root access for the tenant account.

Identity federation not configured

- a. Enter a password for the local root user.
- b. Select **Create tenant**.

SSO enabled

When SSO is enabled for StorageGRID, the tenant must use the identity source that was configured for the Grid Manager. No local users can sign in. You specify which federated group has Root access permission to configure the tenant account.

- a. Select an existing federated group from the Grid Manager to have the initial Root access permission for the tenant.



If you have adequate permissions, the existing federated groups from the Grid Manager are listed when you select the field. Otherwise, enter the group's unique name.

- b. Select **Create tenant**.

SSO not enabled

- a. Complete the steps described in the table depending on whether the tenant will manage its own groups and users or use the identity source that was configured for the Grid Manager.

If the tenant will...	Do this...
Manage its own groups and users	<ol style="list-style-type: none">i. Select Use own identity source. Note: If this check box is selected and you want to use identity federation for tenant groups and users, the tenant must configure its own identity source. See the instructions for using tenant accounts.ii. Specify a password for the tenant's local root user, then select Create tenant.iii. Select Sign in as root to configure the tenant, or select Finish to configure the tenant later.
Use the groups and users configured for the Grid Manager	<ol style="list-style-type: none">i. Do either or both of the following:<ul style="list-style-type: none">◦ Select an existing federated group from the Grid Manager that should have the initial Root access permission for the tenant. Note: If you have adequate permissions, the existing federated groups from the Grid Manager are listed when you select the field. Otherwise, enter the group's unique name.◦ Specify a password for the tenant's local root user.ii. Select Create tenant.

5. To sign in to the tenant now:

- If you are accessing the Grid Manager on a restricted port, 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
- If you are accessing the Grid Manager on port 443 but you did not set a password for the local root user, in the Tenants table of the Grid Manager, select **Sign in**, and enter the credentials for a user in the Root access federated group.
 - If you are accessing the Grid Manager on port 443 and you set a password for the local root user:
 - a. Select **Sign in as root** to configure the tenant now.

When you sign in, links appear for configuring buckets or containers, identity federation, groups, and users.


Create a tenant

Enter details — Select permissions — Define root access







The tenant Tenant02 was created.

If you're ready to configure the tenant, select **Sign in as root**.

Sign in as root  Signed in

You can now access the Tenant Manager to configure these settings:

- **Buckets**  : Create and manage buckets.
- **Identity federation**  : Configure an external identity source to use federated groups.
- **Groups**  : Manage groups and assign permissions.
- **Users**  : Manage local users and assign users to groups.

Finish

- b. 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](#).

c. Otherwise, select **Finish** to access the tenant later.

6. To access the tenant later:

If you are using...	Do one of these...
Port 443	<ul style="list-style-type: none">• From the Grid Manager, select TENANTS, and select Sign in to the right of the tenant name.• Enter the tenant's URL in a web browser: <code>https://FQDN_or_Admin_Node_IP/?accountId=20-digit-account-id/</code><ul style="list-style-type: none">◦ <i>FQDN_or_Admin_Node_IP</i> is a fully qualified domain name or the IP address of an Admin Node◦ <i>20-digit-account-id</i> is the tenant's unique account ID
A restricted port	<ul style="list-style-type: none">• From the Grid Manager, select TENANTS, and select Restricted.• Enter the tenant's URL in a web browser: <code>https://FQDN_or_Admin_Node_IP:port/?accountId=20-digit-account-id</code><ul style="list-style-type: none">◦ <i>FQDN_or_Admin_Node_IP</i> is a fully qualified domain name or the IP address of an Admin Node◦ <i>port</i> is the tenant-only restricted port◦ <i>20-digit-account-id</i> is the tenant's unique account ID

Related information

- [Control access through firewalls](#)
- [Manage platform services for S3 tenant accounts](#)

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.

What you'll need

- 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 cannot sign in to the tenant account. To perform root user tasks, users must belong to a federated group that has the Root access permission for the tenant.

Steps

1. Select **TENANTS**.

Tenants

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

Create

Export to CSV

Actions

Search tenants by name or ID

Displaying 5 results

<input type="checkbox"/>	Name	Logical space used	Quota utilization	Quota	Object count	Sign in/Copy URL
<input type="checkbox"/>	Tenant 01	2.00 GB	<div><div></div></div> 10%	20.00 GB	100	→ 📄
<input type="checkbox"/>	Tenant 02	85.00 GB	<div><div></div></div> 85%	100.00 GB	500	→ 📄
<input type="checkbox"/>	Tenant 03	500.00 TB	<div><div></div></div> 50%	1.00 PB	10,000	→ 📄
<input type="checkbox"/>	Tenant 04	475.00 TB	<div><div></div></div> 95%	500.00 TB	50,000	→ 📄
<input type="checkbox"/>	Tenant 05	5.00 GB	—	—	500	→ 📄

2. Select the tenant account you want to edit.

The Actions button becomes enabled.

3. From the **Actions** drop-down, select **Change root password**.
4. Enter the new password for the tenant account.
5. Select **Save**.

Edit tenant account

You can edit a tenant account to change the display name, change the identity source setting, allow or disallow platform services, or enter a storage quota.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

Steps

1. Select **TENANTS**.

Tenants

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

Create

Export to CSV

Actions

Search tenants by name or ID

Displaying 5 results

<input type="checkbox"/>	Name	Logical space used	Quota utilization	Quota	Object count	Sign in/Copy URL
<input type="checkbox"/>	Tenant 01	2.00 GB	<div><div></div></div> 10%	20.00 GB	100	→ 📄
<input type="checkbox"/>	Tenant 02	85.00 GB	<div><div></div></div> 85%	100.00 GB	500	→ 📄
<input type="checkbox"/>	Tenant 03	500.00 TB	<div><div></div></div> 50%	1.00 PB	10,000	→ 📄
<input type="checkbox"/>	Tenant 04	475.00 TB	<div><div></div></div> 95%	500.00 TB	50,000	→ 📄
<input type="checkbox"/>	Tenant 05	5.00 GB	—	—	500	→ 📄

2. Select the tenant account you want to edit.

Use the search box to search for a tenant account by name or tenant ID.

3. From the Actions drop-down, select **Edit**.

This example is for a grid that does not use single sign-on (SSO). This tenant account has not configured its own identity source.

Edit the tenant

1 Enter details

Select permissions

Enter tenant details

Name ?

Tenant 01

Description (optional) ?

Description

Client type ?

☒ S3 ☐ Swift

Storage quota (optional) ?

GB

Cancel

Continue

4. Change the values for these fields as required:

- **Name**
- **Description**
- **Client type**
- **Storage quota**

5. Select **Continue**.

6. Select or unselect 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.
- Change the setting of the **Uses own identity source** check box to determine whether the tenant account will use its own identity source or the identity source that was configured for the Grid Manager.

If the **Uses own identity source** check box is:

- Disabled and checked, the tenant has already enabled its own identity source. A tenant must disable its identity source before it can use the identity source that was configured for the Grid Manager.
- Disabled and unchecked, SSO is enabled for the StorageGRID system. The tenant must use the

identity source that was configured for the Grid Manager.

- Enable or disable **S3 Select** as needed. See [Manage S3 Select for tenant accounts](#).

7. Select **Save**.

Related information

- [Manage platform services for S3 tenant accounts](#)
- [Use a tenant account](#)

Delete tenant account

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

What you'll need

- You must be signed in to the Grid Manager using a [supported web browser](#).
- You must have specific access permissions.
- You must have removed all buckets (S3), containers (Swift), and objects associated with the tenant account.

Steps

1. Select **TENANTS**.
2. Select the tenant account you want to delete.

Use the search box to search for a tenant account by name or tenant ID.

3. From the **Actions** drop-down, select **Delete**.
4. Select **OK**.

Manage platform services

Manage platform services for S3 tenant accounts

If you enable platform services for S3 tenant accounts, you must configure your grid so that tenants can access the external resources necessary to use these services.

What are platform services?

Platform services include CloudMirror replication, event notifications, and the search integration service.

These services allow tenants to use the following functionality with their S3 buckets:

- **CloudMirror replication:** The StorageGRID CloudMirror replication service is used to mirror specific objects from a StorageGRID bucket to a specified external destination.

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



CloudMirror replication is not supported if the source bucket has S3 Object Lock enabled.

- **Notifications:** Per-bucket event notifications are used to send notifications about specific actions performed on objects to a specified external Amazon Simple Notification Service™ (SNS).

For example, you could configure alerts to be sent to administrators about each object added to a bucket, where the objects represent log files associated with a critical system event.



Although event notification can be configured on a bucket with S3 Object Lock enabled, the S3 Object Lock metadata (including Retain Until Date and Legal Hold status) of the objects will not be included in the notification messages.

- **Search integration service:** The search integration service is used to send S3 object metadata to a specified Elasticsearch index where the metadata can be searched or analyzed using the external service.

For example, you could configure your buckets to send S3 object metadata to a remote Elasticsearch service. You could then use Elasticsearch to perform searches across buckets, and perform sophisticated analyses of patterns present in your object metadata.



Although Elasticsearch integration can be configured on a bucket with S3 Object Lock enabled, the S3 Object Lock metadata (including Retain Until Date and Legal Hold status) of the objects will not be included in the notification messages.

Platform services give tenants the ability to use external storage resources, notification services, and search or analysis services with their data. Because the target location for platform services is typically external to your StorageGRID deployment, you must decide if you want to permit tenants to use these services. If you do, you must enable the use of platform services when you create or edit tenant accounts. You must also configure your network such that the platform services messages that tenants generate can reach their destinations.

Recommendations for using platform services

Before using platform services, 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 cannot 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 are not unfairly penalized.

Related information

- [Use a tenant account](#)
- [Configure Storage proxy settings](#)
- [Monitor and troubleshoot](#)

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 Simple Notification Service (SNS) messages
- A locally-hosted S3 bucket on the same or another instance of StorageGRID
- An external endpoint, such as an endpoint on Amazon Web Services.

To ensure that platform services messages can be delivered, you must configure the network or networks containing the ADC Storage Nodes. You must ensure that the following ports can be used to send platform services messages to the destination endpoints.

By default, platform services messages are sent on the following ports:

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

Tenants can specify a different port when they create or edit an endpoint.



If a StorageGRID deployment is used as the destination for CloudMirror replication, replication messages might be received on a port other than 80 or 443. Ensure that the port being used for S3 by the destination StorageGRID deployment is specified in the endpoint.

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


Related information

- [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, a Simple Notification Service topic, or an Elasticsearch cluster hosted locally or on AWS. Each endpoint includes both the location of the external resource and the credentials needed to access that resource.

When a tenant creates an endpoint, the StorageGRID system validates that the endpoint exists and that it can be reached using the credentials that were specified. The connection to the endpoint is validated from one node at each site.

If endpoint validation fails, an error message explains why endpoint validation failed. The tenant user should resolve the issue, then try creating the endpoint again.




Endpoint creation will fail if platform services are not enabled for the tenant account.

Issues with existing endpoints

If an error occurs when StorageGRID tries to reach an existing endpoint, a message is displayed on the Dashboard in the Tenant Manager.



One or more endpoints have experienced an error and might not be functioning properly. Go to the [Endpoints](#) page to view the error details. The last error occurred 2 hours ago.

Tenant users can go to the Endpoints page to review the most recent error message for each endpoint and to determine how long ago the error occurred. The **Last error** column displays the most recent error message for each endpoint and indicates how long ago the error occurred. Errors that include the  icon occurred within the past 7 days.

Platform services endpoints

A platform services endpoint stores the information StorageGRID needs to use an external resource as a target for a platform service (CloudMirror replication, notifications, or search integration). You must configure an endpoint for each platform service you plan to use.















One or more endpoints have experienced an error. Select the endpoint for more details about the error. Meanwhile, the platform service request will be retried automatically.

5 endpoints

Create endpoint

Delete endpoint

<input type="checkbox"/>	Display name  	Last error  	Type  	URI  	URN  
<input type="checkbox"/>	my-endpoint-2	 2 hours ago	Search	http://10.96.104.30:9200	urn:sgws:es:::mydomain/sveloso/_doc
<input type="checkbox"/>	my-endpoint-3	 3 days ago	Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example1
<input type="checkbox"/>	my-endpoint-5	12 days ago	Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example3
<input type="checkbox"/>	my-endpoint-4		Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example2
<input type="checkbox"/>	my-endpoint-1		S3 Bucket	http://10.96.104.167:10443	urn:sgws:s3:::bucket1



Some error messages in the **Last error** column might include a logID in parentheses. A grid administrator or technical support can use this ID to locate more detailed information about the error in the bycast.log.

Issues related to proxy servers

If you have configured a Storage proxy between Storage Nodes and platform service endpoints, errors might occur if your proxy service does not allow messages from StorageGRID. To resolve these issues, check the

settings of your proxy server to ensure that platform service-related messages are not blocked.

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 cannot be delivered.

If StorageGRID encounters a recoverable error, the platform service request will be retried until it succeeds.

Other errors are unrecoverable. For example, an unrecoverable error occurs if the endpoint is deleted.

If StorageGRID encounters an unrecoverable endpoint error, the Total Events (SMTT) 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 cannot be delivered

If the destination encounters an issue that prevents it from accepting platform services messages, the client operation on the bucket succeeds, but the platform services message is not delivered. For example, this error might happen if credentials are updated on the destination such that StorageGRID can no longer authenticate to the destination service.

If platform services messages cannot be delivered because of an unrecoverable error, the Total Events (SMTT) 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.

[Network](#)[Storage](#)[Objects](#)[ILM](#)[Platform services](#)[Load balancer](#)[1 hour](#)[1 day](#)[1 week](#)[1 month](#)[Custom](#)

Platform services unavailable alert

The **Platform services unavailable** alert indicates that no platform service operations can be performed at a site because too few Storage Nodes with the RSM service are running or available.

The RSM service ensures platform service requests are sent to their respective endpoints.

To resolve this alert, determine which Storage Nodes at the site include the RSM service. (The RSM service is present on Storage Nodes that also include the ADC service.) Then, ensure that a simple majority of those Storage Nodes are running and available.



If more than one Storage Node that contains the RSM service fails at a site, you lose any pending platform service requests for that site.

Additional troubleshooting guidance for platform services endpoints

For additional information about troubleshooting platform services endpoints, see the instructions for [using a tenant account](#).

Related information

- [Monitor and troubleshoot](#)
- [Configure Storage proxy settings](#)

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

StorageGRID requires the following for S3 Select queries:

- The object you want to query is in CSV format, or is a GZIP or BZIP2 compressed file containing a CSV formatted file.
- Tenants must be granted S3 Select ability by the grid administrator. Select **Allow S3 Select** when [creating a tenant](#) or [editing a tenant](#).
- The `SelectObjectContent` request must be sent to a [StorageGRID load balancer endpoint](#). The Admin and Gateway Nodes used by the endpoint must be SG100 or SG1000 appliance nodes or VMware-based software nodes.

Note the following limitations:

- Bare-metal load balancer nodes are not supported.
- Queries cannot be sent directly to Storage Nodes.
- Queries sent through the deprecated CLB service are not supported.



`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 S3 and Swift client connections

About S3 and Swift client connections

As a grid administrator, you manage the configuration options that control how S3 and Swift tenants can connect client applications to your StorageGRID system to store and retrieve data. There are a number of different options to meet different client and tenant requirements.

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

- The Load Balancer service on Admin Nodes or Gateway Nodes, or optionally, the virtual IP address of a high availability (HA) group of Admin Nodes or Gateway Nodes
- The CLB service on Gateway Nodes, or optionally, the virtual IP address of a high availability group of Gateway Nodes



The CLB service is deprecated. Clients configured before the StorageGRID 11.3 release can continue to use the CLB service on Gateway Nodes. All other client applications that depend on StorageGRID to provide load balancing should connect using the Load Balancer service.

- Storage Nodes, with or without an external load balancer

You can optionally configure the following features on your StorageGRID system:

- **VLAN interfaces:** You can create virtual LAN (VLAN) interfaces on Admin Nodes and Gateway Nodes to isolate and partition client and tenant traffic for security, flexibility, and performance. After creating a VLAN interface, you add it to a high availability (HA) group.
- **High availability groups:** You can create an HA group of the interfaces for Gateway Nodes or Admin Nodes to create an active-backup configuration, or you can use round-robin DNS or a third-party load balancer and multiple HA groups to achieve an active-active configuration. Client connections are made using the virtual IP addresses of HA groups.
- **Load Balancer service:** You can enable clients to use the Load Balancer service by creating load balancer endpoints for client connections. When creating a load balancer endpoint, you specify a port number, whether the endpoint accepts HTTP or HTTPS connections, the type of client (S3 or Swift) that will use the endpoint, and the certificate to be used for HTTPS connections (if applicable).
- **Untrusted Client Network:** You can make the Client Network more secure by configuring it as untrusted. When the Client Network is untrusted, clients can only connect using load balancer endpoints.

You can also enable the use of HTTP for clients that connect to StorageGRID either directly to Storage Nodes or using the CLB service (deprecated), and you can configure S3 API endpoint domain names for S3 clients.

Summary: IP addresses and ports for client connections

Client applications can connect to StorageGRID using the IP address of a grid node and the port number of a service on that node. If high availability (HA) groups are configured, client applications can connect using the virtual IP address of the HA group.

About this task

This table summarizes the different ways that clients can connect to StorageGRID and the IP addresses and ports that are used for each type of connection. The instructions describe how to find this information in the

Grid Manager if load balancer endpoints and high availability (HA) groups are already configured.

Where connection is made	Service that client connects to	IP address	Port
HA group	Load Balancer	Virtual IP address of an HA group	<ul style="list-style-type: none">• Load balancer endpoint port
HA group	CLB Note: The CLB service is deprecated.	Virtual IP address of an HA group	Default S3 ports: <ul style="list-style-type: none">• HTTPS: 8082• HTTP: 8084 Default Swift ports: <ul style="list-style-type: none">• HTTPS:8083• HTTP:8085
Admin Node	Load Balancer	IP address of the Admin Node	<ul style="list-style-type: none">• Load balancer endpoint port
Gateway Node	Load Balancer	IP address of the Gateway Node	<ul style="list-style-type: none">• Load balancer endpoint port
Gateway Node	CLB Note: The CLB service is deprecated.	IP address of the Gateway Node Note: By default, HTTP ports for CLB and LDR are not enabled.	Default S3 ports: <ul style="list-style-type: none">• HTTPS: 8082• HTTP: 8084 Default Swift ports: <ul style="list-style-type: none">• HTTPS:8083• HTTP:8085
Storage Node	LDR	IP address of Storage Node	Default S3 ports: <ul style="list-style-type: none">• HTTPS: 18082• HTTP: 18084 Default Swift ports: <ul style="list-style-type: none">• HTTPS: 18083• HTTP:18085

Examples

To connect an S3 client to the Load Balancer endpoint of an HA group of Gateway Nodes, use a URL structured as shown below:

- `https://VIP-of-HA-group:LB-endpoint-port`

For example, if the virtual IP address of the HA group is 192.0.2.5 and the port number of an S3 Load Balancer endpoint is 10443, then an S3 client could use the following URL to connect to StorageGRID:

- `https://192.0.2.5:10443`

To connect a Swift client to the Load Balancer endpoint of an HA group of Gateway Nodes, use a URL structured as shown below:

- `https://VIP-of-HA-group:LB-endpoint-port`

For example, if the virtual IP address of the HA group is 192.0.2.6 and the port number of a Swift Load Balancer endpoint is 10444, then a Swift client could use the following URL to connect to StorageGRID:

- `https://192.0.2.6:10444`

It is possible to configure a DNS name for the IP address that clients use to connect to StorageGRID. Contact your local network administrator.

Steps

1. Sign in to the Grid Manager using a [supported 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**.

The Load Balancer Endpoints page appears, showing the list of endpoints that have already been configured.

- b. Select an endpoint, and select **Edit endpoint**.

The Edit Endpoint window opens and displays additional details about the endpoint.

- c. Confirm that the endpoint you have selected is configured for use with the correct protocol (S3 or Swift), then select **Cancel**.
- d. Note the port number for the endpoint that you want to use for a client connection.



If the port number is 80 or 443, the endpoint is configured only on Gateway Nodes, since those ports are reserved on Admin Nodes. All other ports are configured on both Gateway Nodes and Admin Nodes.

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 cannot 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 or CentOS (before installing the node):** [Create node configuration files](#)
 - **Ubuntu or Debian (before installing the node):** [Create node configuration files](#)
 - **RHEL, CentOS, Ubuntu, or Debian (after installing the node):** [Linux: Add trunk or access interfaces to a node](#)

Create a VLAN interface

What you'll need

- 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

1. Select **CONFIGURATION > Network > VLAN interfaces**.
2. Select **Create**.


Enter details for the VLAN interfaces

1. Specify the ID of the VLAN in your network. You can enter any value between 1 and 4094.

VLAN IDs do not 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 cannot share the same interface on a node.

If you specify an ID that has already been used, a message appears. You can continue creating another VLAN interface for the same VLAN ID, or you can select **Cancel** and then edit the existing ID.

2. Optionally, enter a short description for the VLAN interface.



VLAN details

VLAN ID ⓘ

203

Description (optional) ⓘ

VLAN for S3 tenants. Uses Admin and Gateway Nodes at site 1.

60/64

Cancel Continue

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 cannot be used as parent interfaces and are not shown.

1. Select one or more parent interfaces to attach this VLAN to.

For example, you might want to attach a VLAN to the Client Network (eth2) interface for a Gateway Node and an Admin Node.

Parent interfaces

Select one or more parent interfaces for this VLAN interface. You can only select one parent interface on each node for each VLAN interface.

Search...


	Site ?	Node name ?	Interface ?	Description ?	Node type ?	Attached VLANs ?
<input type="checkbox"/>	Data Center 2	DC2-ADM1	eth0	Grid Network	Non-primary Admin	—
<input checked="" type="checkbox"/>	Data Center 2	DC2-ADM1	eth2	Client Network	Non-primary Admin	—
<input type="checkbox"/>	Data Center 1	DC1-G1	eth0	Grid Network	Gateway	—
<input checked="" type="checkbox"/>	Data Center 1	DC1-G1	eth2	Client Network	Gateway	—
<input type="checkbox"/>	Data Center 1	DC1-ADM1	eth0	Grid Network	Primary Admin	—

2 interfaces are selected.

PreviousContinue

2. Select **Continue**.

Confirm the settings

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 check box 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 unselect the check boxes 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 check box 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 high availability 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 SG100 or SG1000 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 cannot 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 are not 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 cannot be hosted on this node because the High Availability (keepalived) service has been stopped manually.

- **Fault:** The HA group cannot 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 failure of the CLB service (deprecated) or 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 and the CLB service (deprecated).

Purpose of HA group	Add nodes of this type to the HA group
Access to Grid Manager	<ul style="list-style-type: none"> • Primary Admin Node (Primary) • Non-primary Admin Nodes <p>Note: The primary Admin Node must be the Primary interface. Some maintenance procedures can only be performed from the primary Admin Node.</p>
Access to Tenant Manager only	<ul style="list-style-type: none"> • Primary or non-primary Admin Nodes
S3 or Swift client access — Load Balancer service	<ul style="list-style-type: none"> • Admin Nodes • Gateway Nodes

Purpose of HA group	Add nodes of this type to the HA group
S3 client access for S3 Select	<ul style="list-style-type: none"> • SG100 or SG1000 appliances • VMware-based software nodes <p>Note: HA groups are recommended when using S3 Select, but not required.</p>
S3 or Swift client access — CLB service	<ul style="list-style-type: none"> • Gateway Nodes
Note: The CLB service is deprecated.	

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 cannot be performed when the primary Admin Node is unavailable. During failover, you can use the Grid Manager to monitor your StorageGRID system.

Limitations of using HA groups with the CLB service

The failure of the CLB service does not trigger failover within the HA group.






The CLB service is deprecated.

Configuration options for HA groups

The following diagrams provide examples of different ways you can configure HA groups. Each option has advantages and disadvantages.

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	<ul style="list-style-type: none"> Managed by StorageGRID with no external dependencies. Fast failover. 	<ul style="list-style-type: none"> Only one node in an HA group is active. At least one node per HA group will be idle.
DNS Round Robin	<ul style="list-style-type: none"> Increased aggregate throughput. No idle hosts. 	<ul style="list-style-type: none"> Slow failover, which could depend on client behavior. Requires configuration of hardware outside of StorageGRID. Needs a customer-implemented health check.

Configuration	Advantages	Disadvantages
Active-Active HA	<ul style="list-style-type: none"> Traffic is distributed across multiple HA groups. High aggregate throughput that scales with the number of HA groups. Fast failover. 	<ul style="list-style-type: none"> 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.

What you'll need

- 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 or CentOS (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

- Select **CONFIGURATION > Network > High availability groups**.
- Select **Create**.

Enter details for the HA group

- Provide a unique name for the HA group.

×

Create a high availability group

1 Enter details

2 Add interfaces

3 Prioritize interfaces

4 Enter IP addresses

Enter details for the HA group

HA group name

Description (optional)

2. Optionally, enter a description for the HA group.
3. Select **Continue**.

Add interfaces to the HA group

1. Select one or more interfaces to add to this HA group.

Use the column headers to sort the rows, or enter a search term to locate interfaces more quickly.

Add interfaces to the HA group

Select one or more interfaces for this HA group. You can select only one interface for each node.

Total interface count: 4

	Node	Interface	Site	IPv4 subnet	Node type
<input type="checkbox"/>	DC1-ADM1-104-96	eth0	DC1	10.96.104.0/22	Primary Admin Node
<input type="checkbox"/>	DC1-ADM1-104-96	eth2	DC1	—	Primary Admin Node
<input type="checkbox"/>	DC2-ADM1-104-103	eth0	DC2	10.96.104.0/22	Admin Node
<input type="checkbox"/>	DC2-ADM1-104-103	eth2	DC2	—	Admin Node


0 interfaces selected



After creating a VLAN interface, wait up to 5 minutes for the new interface to appear in the table.

Guidelines for selecting interfaces

- You must select at least one interface.
- You can select only one interface for a node.
- If the HA group is for HA protection of Admin Node services, which include the Grid Manager and the Tenant Manager, select interfaces on Admin Nodes only.
- If the HA group is for HA protection of S3 or Swift client traffic, select interfaces on Admin Nodes, Gateway Nodes, or both.
- If the HA group is for HA protection of the deprecated CLB service, select interfaces on Gateway Nodes only.
- 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 cannot provide HA protection of Admin Node services. Similarly, a backup Admin Node cannot perform all of the maintenance procedures that the primary Admin Node can provide.
- If you cannot select an interface, its check box is disabled. The tool tip provides more information.



- You cannot select an interface if its subnet value or gateway conflicts with another selected interface.
- You cannot select a configured interface if it does not have a static IP address.

2. Select **Continue**.

Determine the priority order

1. Determine the Primary interface and any Backup (failover) interfaces for this HA group.

Drag and drop rows to change the values in the **Priority order** column.

Determine the priority order

Determine the primary interface and the backup (failover) interfaces for this HA group. Drag and drop rows or select the arrows.

Priority order 	Node	Interface 	Node type 
1 (Primary interface)	 DC1-ADM1-104-96 	eth2	Primary Admin Node
2	 DC2-ADM1-104-103 	eth2	Admin Node



If the HA group provides access to the Grid Manager, you must select an interface on the primary Admin Node to be the Primary interface. Some maintenance procedures can only be performed from the primary Admin Node.

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 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.

2. Select **Continue**.

Enter IP addresses

1. In the **Subnet CIDR** field, specify the VIP subnet in CIDR notation—an IPv4 address followed by a slash and the subnet length (0-32).

The network address must not have any host bits set. For example, 192.16.0.0/22.



If you use a 32-bit prefix, the VIP network address also serves as the gateway address and the VIP address.

Enter details for the HA group

Subnet CIDR ?

Specify the subnet in CIDR notation. The optional gateway IP and all VIPs must be in this subnet.

IPv4 address followed by a slash and the subnet length (0-32)

Gateway IP address (optional) ?

Optionally specify the IP address of the gateway, which must be in the subnet. If the subnet address length is 32, the gateway IP address is automatically set to the subnet IP.

Virtual IP address ?

Specify at least 1 and no more than 10 virtual IPs for the HA group. All virtual IPs must be in the same subnet. If the subnet length is 32, only one VIP is allowed, which is automatically set to the subnet/gateway IP.

[Add another IP address](#)

- 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.

- Enter one or more **virtual IP addresses** for the HA group. You can add up to 10 IP addresses. All VIPs must be within the VIP subnet.

You must provide at least one IPv4 address. Optionally, you can specify additional IPv4 and IPv6 addresses.

- Select **Create HA group** and select **Finish**.

The HA Group is created, and you can now use the configured virtual IP addresses.



Wait up to 15 minutes for changes to an HA group to be applied to all nodes.

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.


High availability groups

[Learn more about HA groups](#)

You can group the network interfaces of multiple Admin and Gateway Nodes into a high availability (HA) group. If the active interface in the group fails, a backup interface can manage the workload.

Each HA group provides access to the shared services on the selected nodes. Select Gateway Nodes, Admin Nodes, or both for load balancing. Select Admin Nodes for management services. All interfaces in a group must be in the same subnet. You assign one or more virtual IP addresses (VIPs) to each group. Clients use these VIPs to connect to StorageGRID.

 You cannot select an interface if it has a DHCP-assigned IP address.

 Wait up to 15 minutes for changes to an HA group to be applied to all nodes.

Create

Actions ▾

Search...

🔍

Total HA groups count: 2

<input type="checkbox"/>	Name ? ▴ ▾	Description ? ▴ ▾	Virtual IP address ? ▴ ▾	Interfaces (in priority order) ? ▴ ▾
<input type="checkbox"/>	FabricPool	Use for FabricPool client access	10.96.104.5 10.96.104.6	DC1-ADM1-104-96:eth2 (active) DC2-ADM1-104-103:eth2
<input type="checkbox"/>	S3 Clients	use for S3 client access	10.96.104.10	DC1-ADM1-104-96:eth0 DC2-ADM1-104-103:eth0

← Previous 1 Next →

2. Select the check box 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 unselect the check boxes 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 and drop 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**.



Wait up to 15 minutes for changes to an HA group to be applied to all nodes.

Remove a high availability group

You can remove one or more high availability (HA) groups at a time. However, you cannot remove an HA group if it is bound to one or more load balancer endpoints.

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. Select the check box for each HA group you want to remove. Then, select **Actions > Remove HA group**.
3. 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

Manage load balancing: Overview

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

You can load balance client workloads in the following ways:

- Use the Load Balancer service, which is installed on Admin Nodes and Gateway Nodes. The Load Balancer service provides Layer 7 load balancing and performs TLS termination of client requests, inspects the requests, and establishes new secure connections to the Storage Nodes. This is the recommended load balancing mechanism.

See [How load balancing works - Load Balancer service](#).

- Use the deprecated Connection Load Balancer (CLB) service, which is installed on Gateway Nodes only. The CLB service provides Layer 4 load balancing and supports link costs.

See [How load balancing works - CLB service \(deprecated\)](#).

- Integrate a third-party load balancer. Contact your NetApp account representative for details.

How load balancing works - Load Balancer service

The Load Balancer service distributes incoming network connections from client applications to Storage Nodes. To enable load balancing, you must configure load balancer endpoints using the Grid Manager.

You can configure load balancer endpoints only for Admin Nodes or Gateway Nodes, since these node types contain the Load Balancer service. You cannot configure endpoints for Storage Nodes or Archive Nodes.

Each load balancer endpoint specifies a port, a network protocol (HTTP or HTTPS), a client type (S3 or Swift), and a binding mode. HTTPS endpoints require a server certificate. Binding modes allow you to restrict the accessibility of endpoint ports to:

- The virtual IP addresses (VIPs) of specific high availability (HA) groups
- Specific network interfaces of specific Admin and Gateway Nodes

Port considerations

Clients can access any of the endpoints you configure on any node running the Load Balancer service, with two exceptions: ports 80 and 443 are reserved on Admin Nodes, so endpoints configured on these ports support load balancing operations only on Gateway Nodes.

If you have remapped any ports, you cannot use the same ports to configure load balancer endpoints. You can create endpoints using remapped ports, but those endpoints will be remapped to the original CLB ports and service, not the Load Balancer service. Follow the steps in [Remove port remaps](#).



The CLB service is deprecated.

CPU availability

The Load Balancer service on each Admin Node and Gateway Node operates independently when forwarding S3 or Swift traffic to the Storage Nodes. Through a weighting process, the Load Balancer service routes more requests to Storage Nodes with higher CPU availability. Node CPU load information is updated every few minutes, but weighting might be updated more frequently. All Storage Nodes are assigned a minimal base weight value, even if a node reports 100% utilization or fails to report its utilization.

In some cases, information about CPU availability is limited to the site where the Load Balancer service is located.

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.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the Root access permission.
- 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 SG100 or SG1000 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](#).

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`. See [Configure S3 API endpoint domain names](#).

Create a load balancer endpoint

Each load balancer endpoint specifies a port, a client type (S3 or Swift), and a network protocol (HTTP or HTTPS).

Access the wizard

1. Select **CONFIGURATION > Network > Load balancer endpoints**.
2. Select **Create**.

Enter endpoint details

1. Enter details for the endpoint.

×

Create a load balancer endpoint

1 Enter endpoint details

2 Select binding mode

3 Attach certificate

Endpoint details

Name ?

Port ?

Enter an unused port or accept the suggested port.

10443

Client type ?

Select the type of client application that will use this endpoint.

☒ S3
 ☐ Swift

Network protocol ?

Select the network protocol clients will use with this endpoint. If you select HTTPS, attach the security certificate before saving the endpoint.

☐ HTTPS (recommended)
 ☒ HTTP

Cancel

Continue

Field	Description
Name	A descriptive name for the endpoint, which will appear in the table on the Load balancer endpoints page.
Port	<p>The port clients will use to connect to the Load Balancer service on Admin Nodes and Gateway Nodes.</p> <p>Accept the suggested port number or enter any external port that is not used by another grid service. Enter a value between 1 and 65535.</p> <p>If you enter 80 or 443, the endpoint is configured only on Gateway Nodes. These ports are reserved on Admin Nodes.</p> <p>See the Networking guidelines for information about external ports.</p>
Client type	The type of client application that will use this endpoint, either S3 or Swift .

Field	Description
Network protocol	<p>The network protocol that clients will use when connecting to this endpoint.</p> <ul style="list-style-type: none"> • 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.

2. Select **Continue**.

Select the binding mode

1. Select a binding mode for the endpoint to control how the endpoint is accessed.

Option	Description
Global (default)	<p>Clients can access the endpoint using a fully qualified domain name (FQDN), the IP address of any Gateway Node or Admin Node, or the virtual IP address of any HA group on any network.</p> <p>Use the Global setting (default) unless you need to restrict the accessibility of this endpoint.</p>
Node interfaces	<p>Clients must use the IP address of a selected node and network interface to access this endpoint.</p>
Virtual IPs of HA groups	<p>Clients must use a virtual IP address of an HA group to access this endpoint.</p> <p>Endpoints with this binding mode can all use the same port number, as long as the HA groups you select for the endpoints do not overlap.</p> <p>Endpoints with this mode can all use the same port number as long as the interfaces you select for the endpoints do not overlap.</p>



If you use the same port for more than one endpoint, an endpoint using **Virtual IPs of HA groups** mode overrides an endpoint using **Node interfaces** mode, which overrides an endpoint using **Global** mode.

2. 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.

Binding mode ?

Select a binding mode if you plan to monitor or limit the use of this endpoint with a traffic classification policy.

The binding mode controls how the endpoint is accessed—using any IP address or using specific IP addresses and network interfaces.

☐ Global ☒ Node interfaces ☐ Virtual IPs of HA groups

If you use the same port for more than one endpoint, an endpoint bound to HA groups overrides an endpoint bound to Node interfaces, which overrides a Global endpoint. If this behavior does not meet your requirements, consider using a different port number for each endpoint.

Search...

Total interface count: 3

<input type="checkbox"/>	Node	Node interface	Site	IP address	Node type
<input type="checkbox"/>	DC1-ADM1	eth0	Data Center 1	172.16.3.246 and 2 more	Primary Admin Node
<input type="checkbox"/>	DC1-ADM1	eth1	Data Center 1	10.224.3.246 and 5 more	Primary Admin Node
<input type="checkbox"/>	DC1-ADM1	eth2	Data Center 1	47.47.3.246 and 3 more	Primary Admin Node

3. If you selected **Virtual IPs of HA groups**, select one or more HA groups.

Binding mode ?

Select a binding mode if you plan to monitor or limit the use of this endpoint with a traffic classification policy.

The binding mode controls how the endpoint is accessed—using any IP address or using specific IP addresses and network interfaces.

☐ Global ☐ Node interfaces ☒ Virtual IPs of HA groups

If you use the same port for more than one endpoint, an endpoint bound to HA groups overrides an endpoint bound to Node interfaces, which overrides a Global endpoint. If this behavior does not meet your requirements, consider using a different port number for each endpoint.

Search...

Q

Total interface count: 2

<input type="checkbox"/>	Name ?	Description ?	Virtual IP address ?	Interfaces (in priority order) ?
<input type="checkbox"/>	FabricPool	Use for FabricPool client access	10.96.104.5 10.96.104.6	DC1-ADM1-104-96:eth2 (active) DC2-ADM1-104-103:eth2
<input type="checkbox"/>	S3 Clients	use for S3 client access	10.96.104.10	DC1-ADM1-104-96:eth0 DC2-ADM1-104-103:eth0

4. If you are creating an **HTTP** endpoint, you do not 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

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 cannot 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](#).

2. If you are not 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 and the endpoint.

Generate certificate

a. Select **Generate certificate**.

b. Specify the certificate information:

- **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**: X.509 subject or distinguished name (DN) of the certificate owner.
- **Days valid**: Number of days after creation that the certificate expires.

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 and this endpoint.

After you finish

1. If you use a domain name system (DNS), ensure that the DNS includes a record to associate the StorageGRID fully qualified domain name to each IP address that clients will use to make connections.

The IP address you enter in the DNS record depends on whether you are using an HA group of load-balancing nodes:

- If you have configured an HA group, clients will connect to the virtual IP addresses of that HA group.
- If you are not using an HA group, clients will connect to the StorageGRID Load Balancer service using the IP address of any Gateway Node or Admin Node.

You must also ensure that the DNS record references all required endpoint domain names, including any wildcard names.

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 also change an endpoint's name or binding mode and update any associated certificates.

You cannot change the service type (S3 or Swift), the port, or the protocol (HTTP or HTTPS).

- To view basic information for all load balancer endpoints, review the table on the Load balancer endpoints page.

<div>Create Actions Search... Total endpoints count: 1</div>					
<input type="checkbox"/>	Name ?	Port ?	Network protocol ?	Binding mode ?	Certificate expiration ?
<input type="checkbox"/>	FabricPool endpoint	10443	HTTPS	Global	Oct 19th, 2022

- To view all details about a specific endpoint, including certificate metadata, select the endpoint's name in the table.

FabricPool endpoint

Port: 10443
 Client type: S3
 Network protocol: HTTPS
 Binding mode: Global
 Endpoint ID: c2b6feb3-c567-449d-b717-4fed98c4a411

[Remove](#)

Binding Mode

[Certificate](#)

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 or the details page for a specific endpoint.



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	<ol style="list-style-type: none"> Select the check box for the endpoint. Select Actions > Edit endpoint name. Enter the new name. Select Save. 	<ol style="list-style-type: none"> Select the endpoint name to display the details. Select the edit icon . Enter the new name. Select Save.
Edit endpoint binding mode	<ol style="list-style-type: none"> Select the check box for the endpoint. Select Actions > Edit endpoint binding mode. Update the binding mode as required. Select Save changes. 	<ol style="list-style-type: none"> Select the endpoint name to display the details. Select Edit binding mode. Update the binding mode as required. Select Save changes.

Task	Actions menu	Details page
Edit endpoint certificate	<ol style="list-style-type: none"> Select the check box for the endpoint. Select Actions > Edit endpoint certificate. Upload or generate a new custom certificate or begin using the global S3 and Swift certificate, as required. Select Save changes. 	<ol style="list-style-type: none"> Select the endpoint name to display the details. Select the Certificate tab. Select Edit certificate. Upload or generate a new custom certificate or begin using the global S3 and Swift certificate, as required. 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.

- To remove one or more endpoints:
 - From the Load balancer page, select the check box for each endpoint you want to remove.
 - Select **Actions > Remove**.
 - Select **OK**.
- To remove one endpoint from the details page:
 - From the Load balancer page, select the endpoint name.
 - Select **Remove** on the details page.
 - Select **OK**.

How load balancing works - CLB service (deprecated)

The Connection Load Balancer (CLB) service on Gateway Nodes is deprecated. The Load Balancer service is now the recommended load balancing mechanism.


The CLB service uses Layer 4 load balancing to distribute incoming TCP network connections from client applications to the optimal Storage Node based on availability, system load, and the administrator-configured link cost. When the optimal Storage Node is chosen, the CLB service establishes a two-way network connection and forwards the traffic to and from the chosen node. The CLB does not consider the Grid Network configuration when directing incoming network connections.

To view information about the CLB service, select **SUPPORT > Tools > Grid topology**, and then expand a Gateway Node until you can select **CLB** and the options below it.

Grid Topology
StorageGRID Webscale Deployment
Data Center 1
DC1-ADM1-98-160
DC1-G1-98-161
SSM
CLB
HTTP
Events
Resources
DC1-G1-98-162
DC1-S2-98-163
DC1-S3-98-164
DC1-ARC1-98-165
Data Center 2
Data Center 3









OverviewAlarmsReportsConfiguration

Main



Overview: Summary - DC1-G1-98-161
Updated: 2015-10-27 16:23:33 PDT

Storage Capacity

Storage Nodes Installed:	N/A	
Storage Nodes Readable:	N/A	
Storage Nodes Writable:	N/A	
Installed Storage Capacity:	N/A	
Used Storage Capacity:	N/A	
Used Storage Capacity for Data:	N/A	
Used Storage Capacity for Metadata:	N/A	
Usable Storage Capacity:	N/A	

If you choose to use the CLB service, you should consider configuring link costs for your StorageGRID system.

- [What link costs are](#)
- [Update link costs](#)

Configure S3 API endpoint domain names

To support S3 virtual hosted-style requests, you must use the Grid Manager to configure the list of endpoint domain names that S3 clients connect to.

What you'll need

- You 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.



Do not make any changes to the domain name configuration when a grid upgrade is in progress.

About this task

To enable clients to use S3 endpoint domain names, you must do all of the following:

- Use the Grid Manager to add the S3 endpoint domain names to the StorageGRID system.
- Ensure that the certificate the client uses for HTTPS connections to StorageGRID is signed for all domain names that the client requires.

For example, if the endpoint is `s3.company.com`, you must ensure that the certificate used for HTTPS connections includes the `s3.company.com` endpoint and the endpoint's wildcard Subject Alternative Name (SAN): `*.s3.company.com`.

- Configure the DNS server used by the client. Include DNS records for the IP addresses that clients use to make connections, and ensure that the records reference all required endpoint domain names, including any wildcard names.



Clients can connect to StorageGRID using the IP address of a Gateway Node, an Admin Node, or a Storage Node, or by connecting to the virtual IP address of a high availability group. You should understand how client applications connect to the grid so you include the correct IP addresses in the DNS records.

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 endpoint domain names.
- Clients that connect to a load balancer endpoint, directly to a Storage Node, or directly to the deprecated CLB service on a Gateway Node can customize the global S3 and Swift API certificate to include all required endpoint domain names.

Steps

1. Select **CONFIGURATION > Network > Domain names**.

The Endpoint Domain Names page appears.

Endpoint Domain Names

Virtual Hosted-Style Requests

Enable support of S3 virtual hosted-style requests by specifying API endpoint domain names. Support is disabled if this list is empty. Examples: s3.example.com, s3.example.co.uk, s3-east.example.com

Endpoint 1	<input type="text" value="s3.example.com"/>	✕
Endpoint 2	<input type="text"/>	+ ✕

[Save](#)

2. Enter the list of S3 API endpoint domain names in the **Endpoint** fields. Use the **+** icon to add additional fields.

If this list is empty, support for S3 virtual hosted-style requests is disabled.

3. Select **Save**.
4. Ensure that the server certificates that clients use match the required 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, directly to Storage Nodes, or to the CLB service on Gateway 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.

Related information

- [Use S3](#)

- [View IP addresses](#)
- [Configure high availability groups](#)
- [Configure S3 and Swift API certificates](#)
- [Configure load balancer endpoints](#)

Enable HTTP for client communications

By default, client applications use the HTTPS network protocol for all connections to Storage Nodes or to the deprecated CLB service on Gateway Nodes. You can optionally enable HTTP for these connections, for example, when testing a non-production grid.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

About this task

Complete this task only if S3 and Swift clients need to make HTTP connections directly to Storage Nodes or to the deprecated CLB service on Gateway Nodes.

You do not need to complete this task for clients that only use HTTPS connections or for clients that connect to the Load Balancer service (because you can configure each Load Balancer endpoint to use either HTTP or HTTPS). See the information on configuring load balancer endpoints for more information.

See [Summary: IP addresses and ports for client connections](#) to learn which ports S3 and Swift clients use when connecting to Storage Nodes or to the deprecated CLB service using HTTP or HTTPS



Be careful when enabling HTTP for a production grid because requests will be sent unencrypted.

Steps

1. Select **CONFIGURATION > System > Grid options**.
2. In the Network Options section, select the **Enable HTTP Connection** check box.

Network Options

Prevent Client Modification

Enable HTTP Connection ☒

Network Transfer Encryption AES128-SHA AES256-SHA

3. Select **Save**.

Related information

- [Configure load balancer endpoints](#)
- [Use S3](#)

- [Use Swift](#)

Control which client operations are permitted

You can select the Prevent Client Modification grid option to deny specific HTTP client operations.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

About this task

Prevent Client Modification is a system wide setting. When the Prevent Client Modification option is selected, the following requests are denied:

• S3 REST API

- Delete Bucket requests
- Any requests to modify an existing object's data, user-defined metadata, or S3 object tagging



This setting does not apply to buckets with versioning enabled. Versioning already prevents modifications to object data, user-defined metadata, and object tagging.

• Swift REST API

- Delete Container requests
- Requests to modify any existing object. For example, the following operations are denied: Put Overwrite, Delete, Metadata Update, and so on.

Steps

1. Select **CONFIGURATION > System > Grid options**.
2. In the Network Options section, select the **Prevent Client Modification** check box.

Network Options

A screenshot of the 'Network Options' configuration page. The 'Prevent Client Modification' checkbox is checked and highlighted with a yellow rounded rectangle. Below it, the 'Enable HTTP Connection' checkbox is unchecked. At the bottom, the 'Network Transfer Encryption' section shows two radio button options: 'AES128-SHA' (unselected) and 'AES256-SHA' (selected). Each option has a small blue question mark icon to its right.

Prevent Client Modification ☒

Enable HTTP Connection ☐

Network Transfer Encryption ☐ AES128-SHA ☒ AES256-SHA

3. Select **Save**.

Manage networks and connections

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 grid node requires a dedicated network interface, IP address, subnet mask, and gateway for each network it is assigned to.
- A grid node cannot have more than one interface on a network.
- A single gateway, per network, per grid node is supported, and it must be on the same subnet as the node. You can implement more complex routing in the gateway, if required.
- On each node, each network maps to a specific network interface.

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

- If the node is connected to a StorageGRID appliance, specific ports are used for each network. For details,

see the installation instructions for your appliance.

- The default route is generated automatically, per node. If eth2 is enabled, then 0.0.0.0/0 uses the Client Network on eth2. If eth2 is not enabled, then 0.0.0.0/0 uses the Grid Network on eth0.
- The Client Network does not become operational until the grid node has joined the grid
- The Admin Network can be configured during grid node deployment to allow access to the installation user interface before the grid is fully installed.

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](#)
- **RHEL or CentOS (before installing the node):** [Create node configuration files](#)
- **Ubuntu or Debian (before installing the node):** [Create node configuration files](#)
- **RHEL, CentOS, 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 into the grid node at the command line and perform various maintenance procedures.

What you'll need

You are signed in to the Grid Manager using a [supported web browser](#).

About this task

For information on changing IP addresses, see [Recover and maintain](#).

Steps

1. Select **NODES > *grid node* > Overview**.
2. Select **Show more** to the right of the IP Addresses title.


The IP addresses for that grid node are listed in a table.

[Overview](#) [Hardware](#) [Network](#) [Storage](#) [Objects](#) [ILM](#) [Tasks](#)Node information [?](#)

Name: DC2-SGA-010-096-106-021

Type: Storage Node

ID: f0890e03-4c72-401f-ae92-245511a38e51

Connection state:  **Connected**

Storage used:

Object data	<div><div></div></div>	7%	?
Object metadata	<div><div></div></div>	5%	?

Software version: 11.6.0 (build 20210915.1941.afce2d9)

IP addresses: 10.96.106.21 - eth0 (Grid Network)

[Hide additional IP addresses](#) [^](#)

Interface ⌵	IP address ⌵
eth0 (Grid Network)	10.96.106.21
eth0 (Grid Network)	fe80::2a0:98ff:fe64:6582
hic2	10.96.106.21
hic4	10.96.106.21
mtc2	169.254.0.1

Alerts

Alert name ⌵	Severity ? ⌵	Time triggered ⌵	Current values
ILM placement unachievable 🔗	 Major	2 hours ago ?	
A placement instruction in an ILM rule cannot be achieved for certain objects.			

Supported ciphers for outgoing TLS connections

The StorageGRID system supports a limited set of cipher suites for Transport Layer Security (TLS) connections to the external systems used for identity federation and Cloud Storage Pools.

Supported versions of TLS

StorageGRID supports TLS 1.2 and TLS 1.3 for connections to external systems used for identity federation and Cloud Storage Pools.

The TLS ciphers that are supported for use with external systems have been selected to ensure compatibility with a range of external systems. The list is larger than the list of ciphers that are supported for use with S3 or Swift client applications.



TLS configuration options such as protocol versions, ciphers, key exchange algorithms, and MAC algorithms are not configurable in StorageGRID. Contact your NetApp account representative if you have specific requests about these settings.

Supported TLS 1.2 cipher suites

The following TLS 1.2 cipher suites are supported:

- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305
- TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305
- TLS_RSA_WITH_AES_128_GCM_SHA256
- TLS_RSA_WITH_AES_256_GCM_SHA384

Supported TLS 1.3 cipher suites

The following TLS 1.3 cipher suites are supported:

- TLS_AES_256_GCM_SHA384
- TLS_CHACHA20_POLY1305_SHA256
- TLS_AES_128_GCM_SHA256

Change network transfer encryption

The StorageGRID system uses Transport Layer Security (TLS) to protect internal control traffic between grid nodes. The Network Transfer Encryption option sets the algorithm used by TLS to encrypt control traffic between grid nodes. This setting does not affect data encryption.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

About this task

By default, network transfer encryption uses the AES256-SHA algorithm. Control traffic can also be encrypted using the AES128-SHA algorithm.

Steps

1. Select **CONFIGURATION > System > Grid options**.
2. In the Network Options section, change Network Transfer Encryption to **AES128-SHA** or **AES256-SHA** (default).

Network Options



Prevent Client Modification ☐

Enable HTTP Connection ☐

Network Transfer Encryption ☒ AES128-SHA ☒ AES256-SHA

3. Select **Save**.

Manage traffic classification policies

Manage traffic classification policies

To enhance your quality-of-service (QoS) offerings, you can create traffic classification policies to identify and monitor different types of network traffic. These policies can assist with traffic limiting and monitoring.

Traffic classification policies are applied to endpoints on the StorageGRID Load Balancer service for Gateway Nodes and Admin Nodes. To create traffic classification policies, you must have already created load balancer endpoints.

Matching rules

Each traffic classification policy contains one or more matching rules to identify the network traffic related to one or more of the following entities:

- Buckets
- Tenants
- Subnets (IPv4 subnets containing the client)
- Endpoints (load balancer endpoints)

StorageGRID monitors traffic that matches any rule within the policy according to the objectives of the rule. Any traffic that matches any rule for a policy is handled by that policy. Conversely, you can set rules to match all traffic except a specified entity.

Traffic limiting

Optionally, you can set limits for a policy based on the following parameters:

- Aggregate Bandwidth In
- Aggregate Bandwidth Out
- Concurrent Read Requests
- Concurrent Write Requests
- Per-Request Bandwidth In
- Per-Request Bandwidth Out
- Read Request Rate

- Write Requests Rate

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 cannot limit both types of bandwidth at the same time. Aggregate bandwidth limits might impose an additional minor performance impact on non-limited traffic.

For aggregate or per-request bandwidth limits, the requests stream in or out at the rate you set. StorageGRID can only enforce one speed, so the most specific policy match, by matcher type, is the one enforced. For all other limit types, client requests are delayed by 250 milliseconds and receive a 503 Slow Down response for requests that exceed any matching policy limit.

In the Grid Manager, you can view traffic charts and verify that the policies are enforcing the traffic limits you expect.

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.

Traffic classification limits are implemented per load balancer. If traffic is distributed simultaneously across multiple load balancers, the total maximum rates are a multiple of the rate limits you specify.

The following example shows three tiers of an SLA. You can create traffic classification policies to achieve the performance objectives of each SLA tier.

Service Level Tier	Capacity	Data Protection	Performance	Cost
Gold	1 PB storage allowed	3 copy ILM rule	25 K requests/sec 5 GB/sec (40 Gbps) bandwidth	\$\$\$ per month
Silver	250 TB storage allowed	2 copy ILM rule	10 K requests/sec 1.25 GB/sec (10 Gbps) bandwidth	\$\$ per month
Bronze	100 TB storage allowed	2 copy ILM rule	5 K requests/sec 1 GB/sec (8 Gbps) bandwidth	\$ per month

Create traffic classification policies

You create traffic classification policies if you want to monitor, and optionally limit, network traffic by bucket, tenant, IP subnet, or load balancer endpoint. Optionally, you can set limits for a policy based on bandwidth, the number of concurrent requests, or the request rate.

What you'll need

- You 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**.

The Traffic Classification Policies page appears.

Traffic Classification Policies

Traffic classification policies can be used to identify network traffic for metrics reporting and optional traffic limiting.

Create

Edit

Remove

Metrics

Name	Description	ID
No policies found.		

2. Select **Create**.

The Create Traffic Classification Policy dialog box appears.

Create Traffic Classification Policy

Policy

Name ⓘ

Description

Matching Rules

Traffic that matches any rule is included in the policy.

+ Create

✎ Edit

✕ Remove

Type	Inverse Match	Match Value
------	---------------	-------------

No matching rules found.

Limits (Optional)

+ Create

✎ Edit

✕ Remove

Type	Value	Units
------	-------	-------

No limits found.

Cancel

Save

3. In the **Name** field, enter a name for the policy.

Enter a descriptive name so you can recognize the policy.

4. Optionally, add a description for the policy in the **Description** field.

For example, describe what this traffic classification policy applies to and what it will limit.

5. Create one or more matching rules for the policy.



Matching rules control which entities will be affected by this traffic classification policy. For example, select Tenant if you want this policy to apply to the network traffic for a specific tenant. Or select Endpoint if you want this policy to apply to the network traffic on a specific load balancer endpoint.


- a. Select **Create** in the **Matching Rules** section.


The Create Matching Rule dialog box appears.



Create Matching Rule

Matching Rules

Type  -- Choose One -- 

Match Value  Choose type before providing match value

Inverse Match  ☐

b. From the **Type** drop-down, select the type of entity to be included in the matching rule.

c. In the **Match Value** field, enter a match value based on the type of entity you chose.

- **Bucket:** Enter a bucket name.
- **Bucket Regex:** Enter a regular expression that will be used to match a set of bucket names.

The regular expression is unanchored. Use the ^ anchor to match at the beginning of the bucket name, and use the \$ anchor to match at the end of the name.

- **CIDR:** Enter an IPv4 subnet, in CIDR notation, that matches the desired subnet.
- **Endpoint:** Select an endpoint from the list of existing endpoints. These are the load balancer endpoints you defined on the Load Balancer Endpoints page. See [Configure load balancer endpoints](#).
- **Tenant:** Select a tenant from the list of existing tenants. Tenant matching is based on the ownership of the bucket being accessed. Anonymous access to a bucket matches the tenant that owns the bucket.

d. If you want to match all network traffic *except* traffic consistent with the Type and Match Value just defined, select the **Inverse** check box. Otherwise, leave the check box unselected.

For example, if you want this policy to apply to all but one of the load balancer endpoints, specify the load balancer endpoint to be excluded, and select **Inverse**.



For a policy containing multiple matchers where at least one is an inverse matcher, be careful not to create a policy that matches all requests.

e. Select **Apply**.

The rule is created and is listed in the Matching Rules table.

+ Create
Edit
Remove

Type	Inverse Match	Match Value
Bucket Regex	✓	control-ld+

Displaying 1 matching rule.


Limits (Optional)

+ Create
Edit
Remove


Type	Value	Units
No limits found.		

Cancel
Save

- f. Repeat these steps for each rule you want to create for the policy.

 Traffic that matches any rule is handled by the policy.

6. Optionally, create limits for the policy.

 Even if you do not create limits, StorageGRID collects metrics so that you can monitor network traffic that matches the policy.

- a. Select **Create** in the **Limits** section.

The Create Limit dialog box appears.

Create Limit

Limits (Optional)

Type
?
-- Choose One --
Aggregate rate limits in use. Per-request rate limits are not available.
Value
?

Cancel
Apply

- b. From the **Type** drop-down, select the type of limit you want to apply to the policy.

In the following list, **In** refers to traffic from S3 or Swift clients to the StorageGRID load balancer, and **Out** refers to traffic from the load balancer to S3 or Swift clients.

- Aggregate Bandwidth In
- Aggregate Bandwidth Out
- Concurrent Read Requests
- Concurrent Write Requests
- Per-Request Bandwidth In
- Per-Request Bandwidth Out
- Read Request Rate
- Write Requests Rate



You can create policies to limit aggregate bandwidth or to limit per-request bandwidth. However, StorageGRID cannot limit both types of bandwidth at the same time. Aggregate bandwidth limits might impose an additional minor performance impact on non-limited traffic.

For bandwidth limits, StorageGRID applies the policy that best matches the type of limit set. For example, if you have a policy that limits traffic in only one direction, then traffic in the opposite direction will be unlimited, even if there is traffic that matches additional policies that have bandwidth limits. StorageGRID implements “best” matches for bandwidth limits in the following order:

- Exact IP address (/32 mask)
- Exact bucket name
- Bucket regex
- Tenant
- Endpoint
- Non-exact CIDR matches (not /32)
- Inverse matches

c. In the **Value** field, enter a numerical value for the type of limit you chose.

The expected units are shown when you select a limit.

d. Select **Apply**.

The limit is created and is listed in the Limits table.

+ Create
Edit
Remove

Type	Inverse Match	Match Value
<input checked="" type="radio"/> Bucket Regex	✓	control-ld+

Displaying 1 matching rule.

Limits (Optional)

+ Create
Edit
Remove

Type	Value	Units
<input checked="" type="radio"/> Aggregate Bandwidth Out	10000000000	Bytes/Second

Displaying 1 limit.

Cancel
Save

e. Repeat these steps for each limit you want to add to the policy.

For example, if you want to create a 40 Gbps bandwidth limit for an SLA tier, create an Aggregate Bandwidth In limit and an Aggregate Bandwidth Out limit and set each one to 40 Gbps.



To convert megabytes per second to gigabits per second, multiply by eight. For example, 125 MB/s is equivalent to 1,000 Mbps or 1 Gbps.

7. When you are finished creating rules and limits, select **Save**.

The policy is saved and is listed in the Traffic Classification Policies table.

Traffic Classification Policies

Traffic classification policies can be used to identify network traffic for metrics reporting and optional traffic limiting.

+ Create Edit Remove Metrics		
Name	Description	ID
<input type="radio"/> ERP Traffic Control	Manage ERP traffic into the grid	cd9afbc7-b85e-4208-b6f8-7e8a79e2c574
<input checked="" type="radio"/> Fabric Pools	Monitor Fabric Pools	223b0cbb-6968-4646-b32d-7665bddc894b

Displaying 2 traffic classification policies.

S3 and Swift client traffic is now handled according to the traffic classification policies. You can view traffic charts and verify that the policies are enforcing the traffic limits you expect. See [View network traffic metrics](#).

Edit a traffic classification policy

You can edit a traffic classification policy to change its name or description, or to create, edit, or delete any rules or limits for the policy.

What you'll need

- You 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 the table.

Traffic Classification Policies

Traffic classification policies can be used to identify network traffic for metrics reporting and optional traffic limiting.

+ Create

Edit

Remove

Metrics

	Name	Description	ID
<input type="radio"/>	ERP Traffic Control	Manage ERP traffic into the grid	cd9afbc7-b85e-4208-b6f8-7e8a79e2c574
<input checked="" type="radio"/>	Fabric Pools	Monitor Fabric Pools	223b0cbb-6968-4646-b32d-7665bddc894b


Displaying 2 traffic classification policies.

2. Select the radio button to the left of the policy you want to edit.
3. Select **Edit**.

The Edit Traffic Classification Policy dialog box appears.

Edit Traffic Classification Policy "Fabric Pools"

Policy

Name 

Fabric Pools

Description (optional)

Monitor Fabric Pools

Matching Rules

Traffic that matches any rule is included in the policy.

 Create

 Edit

 Remove

	Type	Inverse Match	Match Value
<input checked="" type="radio"/>	CIDR		10.10.152.0/24

Displaying 1 matching rule.

Limits (Optional)

 Create

 Edit

 Remove

	Type	Value	Units
--	------	-------	-------

No limits found.

Cancel

Save

4. Create, edit, or remove matching rules and limits as needed.
 - a. To create a matching rule or limit, select **Create**, and follow the instructions for creating a rule or creating a limit.
 - b. To edit a matching rule or limit, select the radio button for the rule or limit, select **Edit** in the **Matching Rules** section or the **Limits** section, and follow the instructions for creating a rule or creating a limit.
 - c. To remove a matching rule or limit, select the radio button for the rule or limit, and select **Remove**. Then, select **OK** to confirm that you want to remove the rule or limit.
5. When you are finished creating or editing a rule or a limit, select **Apply**.
6. When you are finished editing the policy, select **Save**.

The changes you made to the policy are saved, and network traffic is now handled according to the traffic classification policies. You can view traffic charts and verify that the policies are enforcing the traffic limits you expect.

Delete a traffic classification policy

If you no longer need a traffic classification policy, you can delete it.

What you'll need

- 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 the table.

Traffic Classification Policies

Traffic classification policies can be used to identify network traffic for metrics reporting and optional traffic limiting.

+ Create

✎ Edit

✕ Remove

📊 Metrics

	Name	Description	ID
<input type="radio"/>	ERP Traffic Control	Manage ERP traffic into the grid	cd9afbc7-b85e-4208-b6f8-7e8a79e2c574
<input checked="" type="radio"/>	Fabric Pools	Monitor Fabric Pools	223b0cbb-6968-4646-b32d-7665bddc894b

Displaying 2 traffic classification policies.

2. Select the radio button to the left of the policy you want to delete.
3. Select **Remove**.

A Warning dialog box appears.

⚠ Warning

Delete Policy

Are you sure you want to delete the policy "Fabric Pools"?

Cancel

OK

4. Select **OK** 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.

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 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 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.

Traffic Classification Policies

Traffic classification policies can be used to identify network traffic for metrics reporting and optional traffic limiting.

+ Create

Edit

Remove

Metrics

Name	Description	ID
<input type="radio"/> ERP Traffic Control	Manage ERP traffic into the grid	cd9afbc7-b85e-4208-b6f8-7e8a79e2c574
<input checked="" type="radio"/> Fabric Pools	Monitor Fabric Pools	223b0cbb-6968-4646-b32d-7665bddc894b

Displaying 2 traffic classification policies.

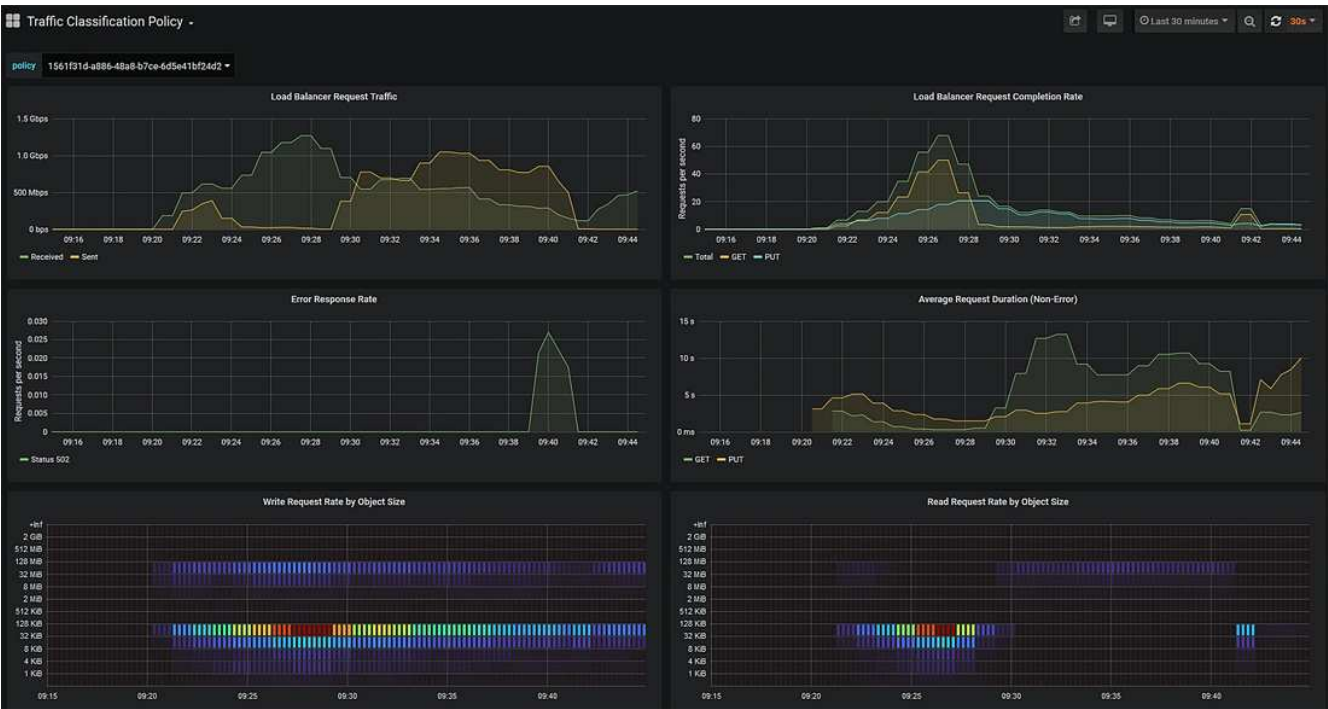


The **Create**, **Edit**, and **Remove** buttons are disabled if you have the Tenant Accounts permission but you do not have the Root access permission.

- 2. Select the radio button to the left of the policy you want to view metrics for.
- 3. Select **Metrics**.

A new browser window opens, and the Traffic Classification Policy graphs appear. The graphs display metrics only for the traffic that matches the selected policy.


You can select other policies to view by using the **policy** pull-down.




The following graphs are included on the web page.

- Load Balancer Request Traffic: This graph provides a 3-minute moving average of the throughput of data transmitted between load balancer endpoints and the clients making the requests, in bits per second.
- Load Balancer Request Completion Rate: This graph provides a 3-minute moving average of the number of completed requests per second, broken down by request type (GET, PUT, HEAD, and DELETE). This value is updated when the headers of a new request have been validated.
- Error Response Rate: This graph provides a 3-minute moving average of the number of error responses returned to clients per second, broken down by the error response code.
- Average Request Duration (Non-Error): This graph provides a 3-minute moving average of request durations, broken down by request type (GET, PUT, HEAD, and DELETE). Each request duration starts when a request header is parsed by the Load Balancer service and ends when the complete response body is returned to the client.
- Write Request Rate by Object Size: This heatmap provides a 3-minute moving average of the rate at which write requests are completed based on object size. In this context, write requests refer only to PUT requests.
- Read Request Rate by Object Size: This heatmap provides a 3-minute moving average of the rate at which read requests are completed based on object size. In this context, read requests refer only to GET requests. The colors in the heatmap indicate the relative frequency of an object size within an individual graph. The cooler colors (for example, purple and blue) indicate lower relative rates, and the warmer colors (for example, orange and red) indicate higher relative rates.

4. Hover the cursor over a line graph to see a pop-up of values on a specific part of the graph.



5. Hover the cursor over a heatmap to see a pop-up that shows the date and time of the sample, object sizes that are aggregated into the count, and the number of requests per second during that time period.



6. Use the **Policy** pull-down in the upper left to select a different policy.

The graphs for the selected policy appear.

7. Alternatively, access the graphs from the **SUPPORT** menu.

- a. Select **SUPPORT > Tools > Metrics**.
- b. In the **Grafana** section of the page, select **Traffic Classification Policy**.
- c. Select the policy from the pull-down on the upper left of the page.

Traffic classification policies are identified by their ID. Policy IDs are listed on the Traffic Classification Policies page.

8. Analyze the graphs to determine how often the policy is limiting traffic and whether you need to adjust the policy.

Related information

[Monitor and troubleshoot](#)


Manage link costs

What link costs are

Link costs let you prioritize which data center site provides a requested service when two or more data center sites exist. You can adjust link costs to reflect latency between sites.

- Link costs are used to prioritize which object copy is used to fulfill object retrievals.
- Link costs are used by the Grid Management API and the Tenant Management API to determine which internal StorageGRID services to use.
- Link costs are used by the deprecated Connection Load Balancer (CLB) service on Gateway Nodes to direct client connections. See [How load balancing works - CLB service](#).

The diagram shows a three site grid that has link costs configured between sites:



- The CLB service on Gateway Nodes equally distribute client connections to all Storage Nodes at the same data center site and to any data center sites with a link cost of 0.

In the example, a Gateway Node at data center site 1 (DC1) equally distributes client connections to Storage Nodes at DC1 and to Storage Nodes at DC2. A Gateway Node at DC3 sends client connections only to Storage Nodes at DC3.

- When retrieving an object that exists as multiple replicated copies, StorageGRID retrieves the copy at the data center that has the lowest link cost.

In the example, if a client application at DC2 retrieves an object that is stored both at DC1 and DC3, the object is retrieved from DC1, because the link cost from DC1 to D2 is 0, which is lower than the link cost from DC3 to DC2 (25).

Link costs are arbitrary relative numbers with no specific unit of measure. For example, a link cost of 50 is used less preferentially than a link cost of 25. The table shows commonly used link costs.

Link	Link cost	Notes
Between physical data center sites	25 (default)	Data centers connected by a WAN link.
Between logical data center sites at the same physical location	0	Logical data centers in the same physical building or campus connected by a LAN.

Update link costs

You can update the link costs between data center sites to reflect latency between sites.


What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).

- You have the Grid Topology Page Configuration permission.

Steps


1. Select **CONFIGURATION** > **Network** > **Link cost**.



The screenshot shows the 'Link Cost' configuration page. At the top, there's a header with a 'Link Cost' title and a timestamp 'Updated: 2021-03-29 12:28:41 EDT'. Below this is a section for 'Site Names (1 - 2 of 2)' with a table containing two rows: '10 Data Center 1' and '20 Data Center 2'. Each row has an 'Actions' column with a pencil icon. Below the table are controls for 'Show 50 Records Per Page', a 'Refresh' button, and pagination links 'Previous', '« 1 » Next'. The main section is titled 'Link Costs' and contains a table with columns 'Link Source', 'Link Destination', and 'Actions'. The 'Link Source' column has a dropdown menu, and the 'Link Destination' column has input fields with values '10' and '20'. An 'Apply Changes' button with a right-pointing arrow is at the bottom right.

2. Select a site under **Link Source** and enter a cost value between 0 and 100 under **Link Destination**.

You cannot change the link cost if the source is the same as the destination.

To cancel changes, select  **Revert**.

3. Select **Apply Changes**.

Use AutoSupport

What is AutoSupport?

The AutoSupport feature enables your StorageGRID system to send health and status messages to technical support.

Using AutoSupport can significantly speed problem determination and resolution. Technical support can also monitor the storage needs of your system and help you determine if you need to add new nodes or sites. Optionally, you can configure AutoSupport messages to be sent to one additional destination.

Information included in AutoSupport messages

AutoSupport messages include information such as the following:

- StorageGRID software version
- Operating system version
- System-level and location-level attribute information
- Recent alerts and alarms (legacy system)

- Current status of all grid tasks, including historical data
- Admin Node database usage
- Number of lost or missing objects
- Grid configuration settings
- NMS entities
- Active ILM policy
- Provisioned grid specification file
- Diagnostic metrics

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

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

If you close the message, it will not appear again until your browser cache is cleared, even if AutoSupport remains disabled.

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.

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

[Active IQ Digital Advisor Documentation](#)

Protocols for sending AutoSupport messages

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

- HTTPS
- HTTP
- SMTP

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

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

AutoSupport options

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

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

Related information

[NetApp Support](#)

Configure AutoSupport

You can enable the AutoSupport feature and the individual AutoSupport options when you first install StorageGRID, or you can enable them later.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the Root access or Other grid configuration permission.
- If you will use the HTTPS or HTTP protocol for sending AutoSupport messages, you have provided outbound internet access to the primary Admin Node, either directly or using a proxy server (inbound connections not required).
- If you will use the HTTPS or HTTP protocol and you want to use a proxy server, you have [configured an Admin proxy server](#).
- If you will use SMTP as the protocol for AutoSupport messages, you have configured an SMTP mail server. The same mail server configuration is used for alarm email notifications (legacy system).

Specify the protocol for AutoSupport messages

You can use any of the following protocols for sending AutoSupport messages:

- **HTTPS:** This is the default and recommended setting for new installations. The HTTPS protocol uses port 443. If you want to enable the AutoSupport on Demand feature, you must use the HTTPS protocol.
- **HTTP:** This protocol is not secure, unless it is used in a trusted environment where the proxy server converts to HTTPS when sending data over the internet. The HTTP protocol uses port 80.
- **SMTP:** Use this option if you want AutoSupport messages to be emailed. If you use SMTP as the protocol for AutoSupport messages, you must configure an SMTP mail server on the Legacy Email Setup page (**SUPPORT > Alarms (legacy) > Legacy email setup**).



SMTP was the only protocol available for AutoSupport messages before the StorageGRID 11.2 release. If you installed an earlier version of StorageGRID initially, SMTP might be the selected protocol.

The protocol you set is used for sending all types of AutoSupport messages.

Steps

1. Select **SUPPORT > Tools > AutoSupport**.

The AutoSupport page appears, and the **Settings** tab is selected.

The screenshot shows the 'AutoSupport' settings page. At the top, there's a title 'AutoSupport' and a descriptive paragraph. Below this are two tabs: 'Settings' (selected) and 'Results'. The main content is divided into sections: 'Protocol Details' with radio buttons for 'HTTPS' (selected), 'HTTP', and 'SMTP'; 'NetApp Support Certificate Validation' with a dropdown menu set to 'Use NetApp support certificate'; 'AutoSupport Details' with checkboxes for 'Enable Weekly AutoSupport' (checked), 'Enable Event-Triggered AutoSupport' (checked), and 'Enable AutoSupport on Demand' (unchecked); 'Software Updates' with a checked checkbox for 'Check for software updates'; and 'Additional AutoSupport Destination' with an unchecked checkbox for 'Enable Additional AutoSupport Destination'. At the bottom right are 'Save' and 'Send User-Triggered AutoSupport' buttons.

2. Select the protocol you want to use to send AutoSupport messages.
3. If you selected **HTTPS**, select whether to use a TLS certificate to secure the connection to the NetApp Support server.
 - **Use NetApp support certificate** (default): Certificate validation ensures that the transmission of AutoSupport messages is secure. The NetApp support certificate is already installed with the StorageGRID software.
 - **Do not verify certificate**: Select this 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 messages are sent using the selected protocol.

Disable weekly AutoSupport messages

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

To determine when the weekly AutoSupport message will be sent, go to the **AutoSupport > Results** tab. In **Weekly AutoSupport** section, look at the value for **Next Scheduled Time**.

AutoSupport

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

Settings

Results

Weekly AutoSupport

Next Scheduled Time ⓘ 2021-09-14 21:10:00 MDT

Most Recent Result ⓘ Idle (NetApp Support)

Last Successful Time ⓘ N/A (NetApp Support)

You can disable the automatic sending of weekly AutoSupport messages at any time.

Steps

1. Select **SUPPORT > Tools > AutoSupport**.
2. Unselect the **Enable Weekly AutoSupport** check box.
3. Select **Save**.

Disable event-triggered AutoSupport messages

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

You can disable event-triggered AutoSupport messages at any time.



Event-triggered AutoSupport messages are also suppressed when you suppress email notifications system wide. (Select **CONFIGURATION > System > Display options**. Then, select **Notification Suppress All**.)

Steps

1. Select **SUPPORT > Tools > AutoSupport**.
2. Unselect the **Enable Event-Triggered AutoSupport** check box.
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 messages automatically. Technical support can also set the polling time interval for AutoSupport on Demand queries.

Technical support cannot enable or disable AutoSupport on Demand.

Steps

1. Select **SUPPORT > Tools > AutoSupport**.
2. Select the **HTTPS** for the protocol.
3. Select the **Enable Weekly AutoSupport** check box.
4. Select the **Enable AutoSupport on Demand** check box.
5. Select **Save**.

AutoSupport on Demand is enabled, and technical support can send AutoSupport on Demand requests to StorageGRID.

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**.
2. Unselect the **Check for software updates** check box.
3. Select **Save**.

Add an additional AutoSupport destination

When you enable AutoSupport, health and status messages are sent to NetApp support. You can specify one additional destinations for all AutoSupport messages.

To verify or change the protocol used to send AutoSupport messages, see the instructions to [Specify the protocol for AutoSupport messages](#).




You cannot use the SMTP protocol to send AutoSupport messages to an additional destination.


Steps


1. Select **SUPPORT > Tools > AutoSupport**.
2. Select **Enable additional AutoSupport destination**.


The Additional AutoSupport Destination fields appear.

Additional AutoSupport Destination

Enable Additional AutoSupport Destination  ☒

Hostname 

Port 

Certificate Validation 

You are not using a TLS certificate to secure the connection to the additional AutoSupport destination.

Save

Send User-Triggered AutoSupport


3. Enter the server hostname or IP address of an additional AutoSupport destination server.





You can enter only one additional destination.


4. Enter the port used to connect to an additional AutoSupport destination server (default is port 80 for HTTP or port 443 for HTTPS).
5. To send your AutoSupport messages with certificate validation, select **Use custom CA bundle** in the **Certificate Validation** drop-down. Then, do one of the following:
 - Use an editing tool to copy and paste all the contents of each of the PEM-encoded CA certificate files into the **CA bundle** field, concatenated in certificate chain order. You must include `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` in your selection.


Additional AutoSupport Destination

Enable Additional AutoSupport Destination  ☒

Hostname 

Port 

Certificate Validation 

CA Bundle 

```
-----BEGIN CERTIFICATE-----  
abcdefghijklmnopqrstuvwxyz123456780ABCDEFGHIJKL  
123456/7890ABCDEFabcdefghijklmnopqrstuvwxyzABCD  
-----END CERTIFICATE-----
```

- Select **Browse**, navigate to the file containing the certificates, and then select **Open** to upload the file. Certificate validation ensures that the transmission of AutoSupport messages is secure.

6. To send your AutoSupport messages without certificate validation, select **Do not verify certificate** in the

Certificate Validation drop-down.

Select this choice only when you have a good reason not to use certificate validation, such as when there is a temporary problem with a certificate.

A caution message appears: "You are not using a TLS certificate to secure the connection to the additional AutoSupport destination."

7. Select **Save**.

All future weekly, event-triggered, and user-triggered AutoSupport messages will be sent to the additional destination.

Manually trigger an AutoSupport message

To assist technical support in troubleshooting issues with your StorageGRID system, you can manually trigger an AutoSupport message to be sent.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the Root access or Other grid configuration permission.

Steps

1. Select **SUPPORT > Tools > AutoSupport**.

The AutoSupport page appears with the **Settings** tab selected.

2. Select **Send User-Triggered AutoSupport**.

StorageGRID attempts to send an AutoSupport message to technical support. If the attempt is successful, the **Most Recent Result** and **Last Successful Time** values on the **Results** tab are updated. If there is a problem, the **Most Recent Result** value updates to "Failed," and StorageGRID does not try to send the AutoSupport message again.



After sending an User-triggered AutoSupport message, refresh the AutoSupport page in your browser after 1 minute to access the most recent results.

Troubleshoot AutoSupport messages

If an attempt to send an AutoSupport message fails, the StorageGRID system takes different actions depending on the type of AutoSupport message. You can check the status of AutoSupport messages by selecting **SUPPORT > Tools > AutoSupport > Results**.



Event-triggered AutoSupport messages are suppressed when you suppress email notifications system wide. (Select **CONFIGURATION > System > Display options**. Then, select **Notification Suppress All**.)

When the AutoSupport message fails to send, "Failed" appears on the **Results** tab of the **AutoSupport** page.

AutoSupport

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

Settings

Results

Weekly AutoSupport

Next Scheduled Time ? 2020-12-11 23:30:00 EST

Most Recent Result ? Idle (NetApp Support)

Last Successful Time ? N/A (NetApp Support)

Event-Triggered AutoSupport

Most Recent Result ? N/A (NetApp Support)

Last Successful Time ? N/A (NetApp Support)

User-Triggered AutoSupport

Most Recent Result ? Failed (NetApp Support)

Last Successful Time ? N/A (NetApp Support)

AutoSupport On Demand

AutoSupport On Demand messages are only sent to NetApp Support.

Most Recent Result ? N/A (NetApp Support)

Last Successful Time ? N/A (NetApp Support)

Weekly AutoSupport message failure

If a weekly AutoSupport message fails to send, the StorageGRID system takes the following actions:

1. Updates the Most Recent Result attribute to Retrying.
2. Attempts to resend the AutoSupport message 15 times every four minutes for one hour.
3. After one hour of send failures, updates the Most Recent Result attribute to Failed.
4. Attempts to send an AutoSupport message again at the next scheduled time.
5. Maintains the regular AutoSupport schedule if the message fails because the NMS service is unavailable, and if a message is sent before seven days pass.
6. When the NMS service is available again, sends an AutoSupport message immediately if a message has not been sent for seven days or more.

User-triggered or event-triggered AutoSupport message failure

If a user-triggered or an event-triggered AutoSupport message fails to send, the StorageGRID system takes the following actions:

1. Displays an error message if the error is known. For example, if a user selects the SMTP protocol without providing correct email configuration settings, the following error is displayed: AutoSupport messages cannot be sent using SMTP protocol due to incorrect settings on the E-mail Server page.
2. Does not attempt to send the message again.
3. Logs the error in `nms.log`.

If a failure occurs and SMTP is the selected protocol, verify that the StorageGRID system's email server is correctly configured and that your email server is running (**SUPPORT > Alarms (legacy) > > Legacy Email Setup**). The following error message might appear on the AutoSupport page: AutoSupport messages cannot be sent using SMTP protocol due to incorrect settings on the E-mail Server page.

Learn how to configure email server settings in the [monitor and troubleshoot instructions](#).

Correct an AutoSupport message failure

If a failure occurs and SMTP is the selected protocol, verify that the StorageGRID system's email server is correctly configured and that your email server is running. The following error message might appear on the AutoSupport page: AutoSupport messages cannot be sent using SMTP protocol due to incorrect settings on the E-mail Server page.

Send E-Series AutoSupport messages through StorageGRID

You can send E-Series SANtricity System Manager AutoSupport messages to technical support through a StorageGRID Admin Node rather than the storage appliance management port.

What you'll need

- You are signed into the Grid Manager using a [supported web browser](#).
- You have the Storage Appliance Administrator permission or Root access permission.



You must have SANtricity firmware 8.70 (11.7) or higher to access SANtricity System Manager using the Grid Manager.

About this task

E-Series AutoSupport messages contain details of the storage hardware and are more specific than other AutoSupport messages sent by the StorageGRID system.

Configure a special proxy server address in SANtricity System Manager to cause the AutoSupport messages to be transmitted through a StorageGRID Admin Node without the use of the appliance's management port. AutoSupport messages transmitted in this way respect the Preferred Sender and Admin proxy settings which may have been configured in the Grid Manager.

If you want to configure the Admin proxy server in Grid Manager, see [Configure Admin proxy settings](#).




This procedure is only for configuring a StorageGRID proxy server for E-Series AutoSupport messages. For additional details on E-Series AutoSupport configuration, 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.

[Support Resources](#)

[Diagnostics](#)

AutoSupport

AutoSupport operations

AutoSupport status: **Enabled** 

[Enable/Disable AutoSupport Features](#)

AutoSupport proactively monitors the health of your storage array and automatically sends support data ("dispatches") to the support team.

[Configure AutoSupport Delivery Method](#)

Connect to the support team via HTTPS, HTTP or Mail (SMTP) server delivery methods.

[Schedule AutoSupport Dispatches](#)

AutoSupport dispatches are sent daily at 03:06 PM UTC and weekly at 07:39 AM UTC on Thursday.

[Send AutoSupport Dispatch](#)

Automatically sends the support team a dispatch to troubleshoot system issues without waiting for periodic dispatches.

[View AutoSupport Log](#)

The AutoSupport log provides information about status, dispatch history, and errors encountered during delivery of AutoSupport dispatches.

[Enable AutoSupport Maintenance Window](#)

Enable AutoSupport Maintenance window to allow maintenance activities to be performed on the storage array without generating support cases.

[Disable AutoSupport Maintenance Window](#)

Disable AutoSupport Maintenance window to allow the storage array to generate support cases on component failures and other destructive actions.

5. Select **Configure AutoSupport Delivery Method**.

The Configure AutoSupport Delivery Method page appears.

6. Select **HTTPS** for the delivery method.



The certificate that enables the HTTPS protocol is pre-installed.

7. Select **via Proxy server**.

8. Enter `tunnel-host` for the **Host address**.

`tunnel-host` is the special address to use an Admin Node to send E-Series AutoSupport messages.

9. Enter `10225` for the **Port number**.

`10225` is the port number on the StorageGRID proxy server that receives AutoSupport messages from the E-Series controller in the appliance.

10. Select **Test Configuration** to test the routing and configuration of your AutoSupport proxy server.

If correct, a message in a green banner appears: "Your AutoSupport configuration has been verified."

If the test fails, an error message appears in a red banner. Check your StorageGRID DNS settings and

networking, ensure the preferred sender Admin Node can connect to the NetApp support site, and try the test again.

11. Select **Save**.

The configuration is saved, and a confirmation message appears: “AutoSupport delivery method has been configured.”

Manage Storage Nodes

About managing Storage Nodes

Storage Nodes provide disk storage capacity and services. Managing Storage Nodes entails the following:

- Managing storage options
- Understanding what storage volume watermarks are and how you can use watermark overrides to control when Storage Nodes become read-only
- Monitoring and managing the space used for object metadata
- Configuring global settings for stored objects
- Applying Storage Node configuration settings
- Managing full Storage Nodes

What is a Storage Node?

Storage Nodes manage and store object data and metadata. Each StorageGRID system must have at least three Storage Nodes. If you have multiple sites, each site within your StorageGRID system must also have three Storage Nodes.

A Storage Node includes the services and processes required to store, move, verify, and retrieve object data and metadata on disk. You can view detailed information about the Storage Nodes on the **NODES** page.

What is the ADC service?

The Administrative Domain Controller (ADC) service authenticates grid nodes and their connections with each other. The ADC service is hosted on each of the first three Storage Nodes at a site.

The ADC service maintains topology information including the location and availability of services. When a grid node requires information from another grid node or an action to be performed by another grid node, it contacts an ADC service to find the best grid node to process its request. In addition, the ADC service retains a copy of the StorageGRID deployment’s configuration bundles, allowing any grid node to retrieve current configuration information. You can view ADC information for a Storage Node on the Grid Topology page (**SUPPORT > Grid topology**).

To facilitate distributed and islanded operations, each ADC service synchronizes certificates, configuration bundles, and information about services and topology with the other ADC services in the StorageGRID system.

In general, all grid nodes maintain a connection to at least one ADC service. This ensures that grid nodes are always accessing the latest information. When grid nodes connect, they cache other grid nodes’ certificates, enabling systems to continue functioning with known grid nodes even when an ADC service is unavailable. New grid nodes can only establish connections by using an ADC service.

The connection of each grid node lets the ADC service gather topology information. This grid node information includes the CPU load, available disk space (if it has storage), supported services, and the grid node’s site ID. Other services ask the ADC service for topology information through topology queries. The ADC service responds to each query with the latest information received from the StorageGRID system.


What is the DDS service?

Hosted by a Storage Node, the Distributed Data Store (DDS) service interfaces with the Cassandra database to perform background tasks on the object metadata stored in the StorageGRID system.

Object counts

The DDS service tracks the total number of objects ingested into the StorageGRID system as well as the total number of objects ingested through each of the system’s supported interfaces (S3 or Swift).

You can see the Total Objects count on the Nodes page > Objects tab for any Storage Node.



Queries

You can identify the average time that it takes to run a query against the metadata store through the specific DDS service, the total number of successful queries, and the total number of queries that failed because of a timeout issue.

You might want to review query information to monitor the health of the metadata store, Cassandra, which impacts the system's ingest and retrieval performance. For example, if the latency for an average query is slow and the number of failed queries due to timeouts is high, the metadata store might be encountering a higher load or performing another operation.

You can also view the total number of queries that failed because of consistency failures. Consistency level failures result from an insufficient number of available metadata stores at the time a query is performed through the specific DDS service.

You can use the Diagnostics page to obtain additional information on the current state of your grid. See [Run diagnostics](#).

Consistency guarantees and controls

StorageGRID guarantees read-after-write consistency for newly created objects. Any GET operation following a successfully completed PUT operation will be able to read the newly written data. Overwrites of existing objects, metadata updates, and deletes remain eventually consistent.

What is the LDR service?

Hosted by each Storage Node, the Local Distribution Router (LDR) service handles content transport for the StorageGRID system. Content transport encompasses many tasks including data storage, routing, and request handling. The LDR service does the majority of the StorageGRID system's hard work by handling data transfer loads and data traffic functions.

The LDR service handles the following tasks:

- Queries
- Information lifecycle management (ILM) activity
- Object deletion
- Object data storage
- Object data transfers from another LDR service (Storage Node)
- Data storage management
- Protocol interfaces (S3 and Swift)

The LDR service also manages the mapping of S3 and Swift objects to the unique "content handles" (UUIDs) that the StorageGRID system assigns to each ingested object.

Queries

LDR queries include queries for object location during retrieve and archive operations. You can identify the average time that it takes to run a query, the total number of successful queries, and the total number of queries that failed because of a timeout issue.

You can review query information to monitor the health of the metadata store, which impacts the system's ingest and retrieval performance. For example, if the latency for an average query is slow and the number of failed queries due to timeouts is high, the metadata store might be encountering a higher load or performing another operation.

You can also view the total number of queries that failed because of consistency failures. Consistency level failures result from an insufficient number of available metadata stores at the time a query is performed through the specific LDR service.

You can use the Diagnostics page to obtain additional information on the current state of your grid. See [Run diagnostics](#).

















ILM activity

Information lifecycle management (ILM) metrics allow you to monitor the rate at which objects are evaluated for ILM implementation. You can view these metrics on the Dashboard or at **NODES > Storage Node > ILM**.

Object stores

The underlying data storage of an LDR service is divided into a fixed number of object stores (also known as storage volumes). Each object store is a separate mount point.

You can see the object stores for a Storage Node on the Nodes page > Storage tab.

Object stores						
ID 	Size 	Available 	Replicated data 	EC data 	Object data (%) 	Health 
0000	107.32 GB	96.44 GB 	124.60 KB 	0 bytes 	0.00%	No Errors
0001	107.32 GB	107.18 GB 	0 bytes 	0 bytes 	0.00%	No Errors
0002	107.32 GB	107.18 GB 	0 bytes 	0 bytes 	0.00%	No Errors

The object stores in a Storage Node are identified by a hexadecimal number from 0000 to 002F, which is known as the volume ID. Space is reserved in the first object store (volume 0) for object metadata in a Cassandra database; any remaining space on that volume is used for object data. All other object stores are used exclusively for object data, which includes replicated copies and erasure-coded fragments.

To ensure even space usage for replicated copies, object data for a given object is stored to one object store based on available storage space. When one or more object stores fill to capacity, the remaining object stores continue to store objects until there is no more room on the Storage Node.

Metadata protection

Object metadata is information related to or a description of an object; for example, object modification time, or storage location. StorageGRID stores object metadata in a Cassandra database, which interfaces with the LDR service.

To ensure redundancy and thus protection against loss, three copies of object metadata are maintained at each site. The copies are evenly distributed across all Storage Nodes at each site. This replication is non-configurable and performed automatically.

[Manage object metadata storage](#)

Manage Storage options


Storage options include the object segmentation settings, the current values for storage volume watermarks, and the Metadata Reserved Space setting. You can also view the S3 and Swift ports used by the deprecated CLB service on Gateway Nodes and by the LDR service on Storage Nodes.

For information about port assignments, see [Summary: IP addresses and ports for client connections](#).

Storage Options

Overview

Configuration



Storage Options Overview

Updated: 2021-11-23 11:01:41 MST

Object Segmentation

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1 GB

Storage Watermarks

Description	Settings
Storage Volume Read-Write Watermark Override	0 B
Storage Volume Soft Read-Only Watermark Override	0 B
Storage Volume Hard Read-Only Watermark Override	0 B
Metadata Reserved Space	3,000 GB


Ports

Description	Settings
CLB S3 Port	8082
CLB Swift Port	8083
LDR S3 Port	18082
LDR Swift Port	18083

What is object segmentation?

Object segmentation is the process of splitting up an object into a collection of smaller fixed-size objects in order to optimize storage and resources usage for large objects. S3 multi-part upload also creates segmented objects, with an object representing each part.

When an object is ingested into the StorageGRID system, the LDR service splits the object into segments, and creates a segment container that lists the header information of all segments as content.



On retrieval of a segment container, the LDR service assembles the original object from its segments and returns the object to the client.

The container and segments are not necessarily stored on the same Storage Node. Container and segments can be stored on any Storage Node 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 IO 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.

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 read-only mode. It determines when the node can become read-write again. When the free space on any one storage volume in a Storage Node is greater than that volume's Read-Write Watermark, the node automatically transitions back to the read-write state.

For example, suppose the Storage Node has transitioned to read-only mode. Also suppose that each volume has a Read-Write Watermark of 30 GB. As soon as the free space for any volume increases to 30 GB, the node becomes read-write again.

The Read-Write Watermark is always larger than both the Soft Read-Only Watermark and the Hard Read-Only Watermark.

View storage volume watermarks

You can view the current watermark settings and the system-optimized values. If optimized watermarks are not being used, you can determine if you can or should adjust the settings.

What you'll need

- You have completed the upgrade to StorageGRID 11.6.
- 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 **CONFIGURATION > System > Storage options**.
2. In the Storage Watermarks section, look at the settings for the three storage volume watermark overrides.

Storage Options Overview
Updated: 2021-11-22 13:57:51 MST

Object Segmentation

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1 GB

Storage Watermarks

Description	Settings
Storage Volume Read-Write Watermark Override	0 B
Storage Volume Soft Read-Only Watermark Override	0 B
Storage Volume Hard Read-Only Watermark Override	0 B
Metadata Reserved Space	3,000 GB

Ports

Description	Settings
CLB S3 Port	8082
CLB Swift Port	8083
LDR S3 Port	18082
LDR Swift Port	18083

- If the watermark overrides are **0**, 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. You should not update these values. As required, you can optionally [View optimized storage watermarks](#).

- If the watermark overrides are non-0 values, 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.

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. The three copies of object metadata are evenly distributed across all Storage Nodes at each site.

This figure represents the Storage Nodes at two sites. Each site has the same amount of object metadata, which is equally distributed across the Storage Nodes at that site.



Where is object metadata stored?

This figure represents the storage volumes for a single Storage Node.



As shown in the figure, StorageGRID reserves space for object metadata on storage volume 0 of each Storage Node. It uses the reserved space to store object metadata and to perform essential database operations. Any remaining space on storage volume 0 and all other storage volumes in the Storage Node are used exclusively for object data (replicated copies and erasure-coded fragments).

The amount of space that is reserved for object metadata on a particular Storage Node depends on a number of factors, which are described below.

Metadata Reserved Space setting

The *Metadata Reserved Space* is a system-wide setting that represents the amount of space that will be reserved for metadata on volume 0 of every Storage Node. As shown in the table, the default value of this setting for StorageGRID 11.6 is based the following:


- The software version you were using when you initially installed StorageGRID.

- The amount of RAM on each Storage Node.

Version used for initial StorageGRID installation	Amount of RAM on Storage Nodes	Default Metadata Reserved Space setting for StorageGRID 11.6
11.5/11.6	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)

To view the Metadata Reserved Space setting for your StorageGRID system:

1. Select **CONFIGURATION > System > Storage options**.
2. In the Storage Watermarks table, locate **Metadata Reserved Space**.



Storage Options Overview
Updated: 2021-12-10 13:53:01 MST

Object Segmentation

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1 GB

Storage Watermarks

Description	Settings
Storage Volume Read-Write Watermark Override	0 B
Storage Volume Soft Read-Only Watermark Override	0 B
Storage Volume Hard Read-Only Watermark Override	0 B
Metadata Reserved Space	8,000 GB

In the screenshot, the **Metadata Reserved Space** value is 8,000 GB (8 TB). This is the default setting for a new StorageGRID 11.6 installation in which each Storage Node has 128 GB or more of RAM.


Actual reserved space for metadata

In contrast to the system-wide Metadata Reserved Space setting, the *actual reserved space* for object metadata is determined for each Storage Node. For any given Storage Node, the actual reserved space for metadata depends on the size of volume 0 for the node and the system-wide **Metadata Reserved Space** setting.

Size of volume 0 for the node	Actual reserved space for metadata
Less than 500 GB (non production use)	10% of volume 0
500 GB or more	The smaller of these values: <ul style="list-style-type: none">• Volume 0• Metadata Reserved Space setting

To view the actual reserved space for metadata on a particular Storage Node:

1. From the Grid Manager, select **NODES > Storage Node**.
2. Select the **Storage** tab.
3. Hover your cursor over the Storage Used — Object Metadata chart and locate the **Actual reserved** value.



In the screenshot, the **Actual reserved** value is 8 TB. This screenshot is for a large Storage Node in a new StorageGRID 11.6 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

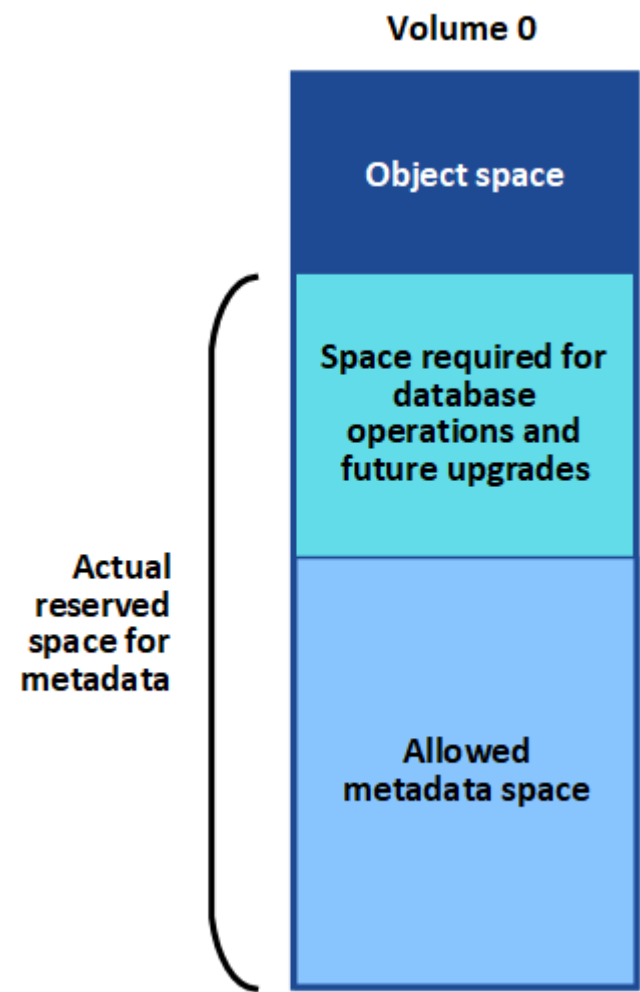
Suppose you install a new StorageGRID system using version 11.6. For this example, assume that each Storage Node has more than 128 GB of RAM and that volume 0 of Storage Node 1 (SN1) is 6 TB. Based on these values:

- The system-wide **Metadata Reserved Space** is set to 8 TB. (This is the default value for a new StorageGRID 11.6 installation if each Storage Node has more than 128 GB RAM.)

- The actual reserved space for metadata for SN1 is 6 TB. (The entire volume is reserved because volume 0 is smaller than the **Metadata Reserved Space** setting.)

Allowed metadata space

Each Storage Node's actual reserved space for metadata is subdivided into the space available for object metadata (the *allowed metadata space*) and the space required for essential database operations (such as compaction and repair) and future hardware and software upgrades. The allowed metadata space governs overall object capacity.



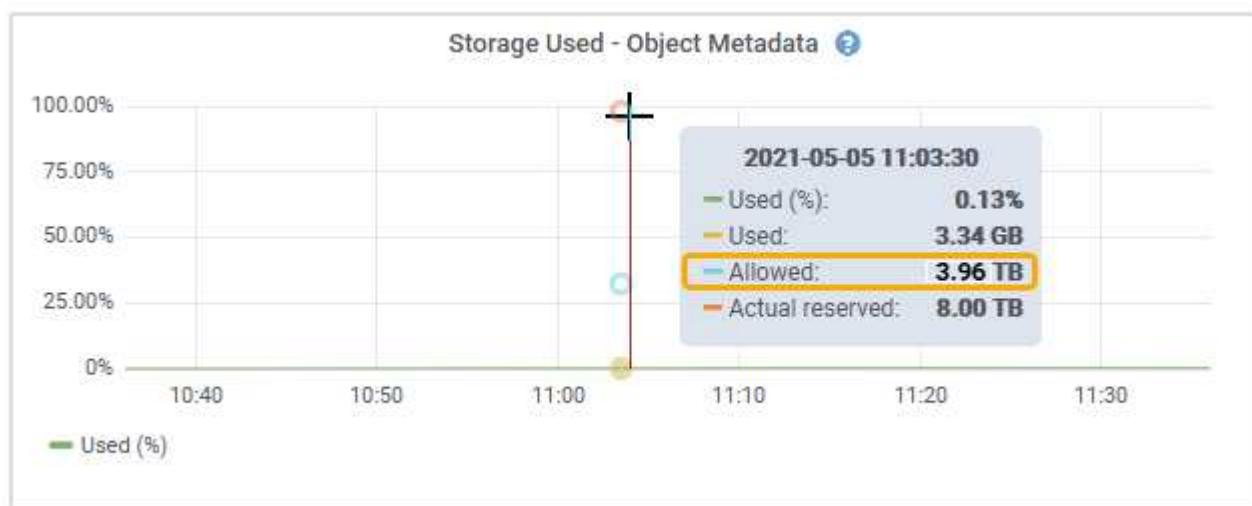
The following table shows how StorageGRID calculates the **allowed metadata space** for different Storage Nodes, based on the amount of memory for the node and the actual reserved space for metadata.

Amount of memory on Storage Node	
< 128 GB	>= 128 GB

Actual reserved space for metadata	<= 4 TB	60% of actual reserved space for metadata, up to a maximum of 1.32 TB	60% of actual reserved space for metadata, up to a maximum of 1.98 TB
	> 4 TB	(Actual reserved space for metadata – 1 TB) × 60%, up to a maximum of 1.32 TB	(Actual reserved space for metadata – 1 TB) × 60%, up to a maximum of 3.96 TB

To view the allowed metadata space for a Storage Node:

1. From the Grid Manager, select **NODES**.
2. Select the Storage Node.
3. Select the **Storage** tab.
4. Hover 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 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 cannot be used.



Similarly, because StorageGRID maintains all object metadata for a StorageGRID system at each site, the overall metadata capacity of a StorageGRID system is determined by the object metadata capacity of the smallest site.

And because object metadata capacity controls the maximum object count, when one node runs out of metadata capacity, the grid is effectively full.

Related information

- To learn how to monitor the object metadata capacity for each Storage Node, go to [Monitor and troubleshoot](#).
- To increase the object metadata capacity for your system, add new Storage Nodes. Go to [Expand your grid](#).

Configure global settings for stored objects

Configure stored object compression

You can use the Compress Stored Objects grid option to reduce the size of objects stored in StorageGRID, so that objects consume less storage.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

About this task

The Compress Stored Objects grid option is disabled by default. If you enable this option, StorageGRID attempts to compress each object when saving it, using lossless compression.



If you change this setting, it will take about one minute for the new setting to be applied. The configured value is cached for performance and scaling.

Before enabling this option, be aware of the following:

- You should not enable compression unless you know that the data being stored is compressible.
- Applications that save objects to StorageGRID might compress objects before saving them. If a client application has already compressed an object before saving it to StorageGRID, enabling Compress Stored Objects will not further reduce an object's size.
- Do not enable compression if you are using NetApp FabricPool with StorageGRID.
- If the Compress Stored Objects grid option is enabled, S3 and Swift client applications should avoid performing GET Object operations that specify a range of bytes be returned. These “range read” operations are inefficient because StorageGRID must effectively uncompress the objects to access the requested bytes. GET Object operations that request a small range of bytes from a very large object are especially inefficient; for example, it is inefficient to read a 10 MB range from a 50 GB compressed object.

If ranges are read from compressed objects, client requests can time out.



If you need to compress objects and your client application must use range reads, increase the read timeout for the application.

Steps

1. Select **CONFIGURATION > System > Grid options**.
2. In the Stored Object Options section, select the **Compress Stored Objects** check box.

Stored Object Options

Compress Stored Objects ? 

Stored Object Encryption ? ☒ None ☐ AES-128 ☐ AES-256

Stored Object Hashing ? ☒ SHA-1 ☐ SHA-256

3. Select **Save**.

Configure stored object encryption

You can encrypt stored objects if you want to ensure that data cannot be retrieved in a readable form if an object store is compromised. By default, objects are not encrypted.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

About this task

Stored object encryption enables the encryption of all object data as it is ingested through S3 or Swift. When you enable the setting, all newly ingested objects are encrypted but no change is made to existing stored objects. If you disable encryption, currently encrypted objects remain encrypted but newly ingested objects are not encrypted.



If you change this setting, it will take about one minute for the new setting to be applied. The configured value is cached for performance and scaling.


Stored objects can be encrypted using the AES-128 or AES-256 encryption algorithm.

The Stored Object Encryption setting applies only to S3 objects that have not been encrypted by bucket-level or object-level encryption.

Steps

1. Select **CONFIGURATION > System > Grid options**.
2. In the Stored Object Options section, change Stored Object Encryption to **None** (default), **AES-128**, or **AES-256**.

Stored Object Options

Compress Stored Objects ? 

Stored Object Encryption ? ☒ None ☐ AES-128 ☐ AES-256

Stored Object Hashing ? ☒ SHA-1 ☐ SHA-256

3. Select **Save**.

Configure stored object hashing

The Stored Object Hashing option specifies the hashing algorithm used to verify object integrity.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

About this task

By default, object data is hashed using the SHA-1 algorithm. The SHA-256 algorithm requires additional CPU resources and is generally not recommended for integrity verification.



If you change this setting, it will take about one minute for the new setting to be applied. The configured value is cached for performance and scaling.

Steps

1. Select **CONFIGURATION > System > Grid options**.
2. In the Stored Object Options section, change Stored Object Hashing to **SHA-1** (default) or **SHA-256**.

Stored Object Options

Compress Stored Objects ? ☐

Stored Object Encryption ? ☒ None ☐ AES-128 ☐ AES-256

Stored Object Hashing ? ☒ SHA-1 ☐ SHA-256

3. Select **Save**.

Storage Node configuration settings

Each Storage Node uses a number of configuration settings and counters. You might need to view the current settings or reset counters to clear alarms (legacy system).



Except when specifically instructed in documentation, you should consult with technical support before modifying any Storage Node configuration settings. As required, you can reset event counters to clear legacy alarms.

To access a Storage Node's configuration settings and counters:

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **site > Storage Node**.
3. Expand the Storage Node and select the service or component.

4. Select the **Configuration** tab.

The following tables summarize Storage Node configuration settings.

LDR

Attribute Name	Code	Description
HTTP State	HSTE	<p>The current state of the HTTP protocol for S3, Swift, and other internal StorageGRID traffic:</p> <ul style="list-style-type: none">• 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	<ul style="list-style-type: none">• 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	<p>A user-configurable setting for the desired state of the storage component. The LDR service reads this value and attempts to match the status indicated by this attribute. The value is persistent across restarts.</p> <p>For example, you can use this setting to force storage to become read-only even when there is ample available storage space. This can be useful for troubleshooting.</p> <p>The attribute can take one of the following values:</p> <ul style="list-style-type: none"> • Offline: When the desired state is Offline, the LDR service takes the LDR > Storage component offline. • Read-only: When the desired state is Read-only, the LDR service moves the storage state to read-only and stops accepting new content. Note that content might continue to be saved to the Storage Node for a short time until open sessions are closed. • Online: Leave the value at Online during normal system operations. The Storage State — Current of the storage component will be dynamically set by the service based on the condition of the LDR service, such as the amount of available object storage space. If space is low, the component becomes Read-only.
Health Check Timeout	SHCT	The time limit in seconds within which a health check test must complete in order for a storage volume to be considered healthy. Only change this value when directed to do so by Support.

LDR > Verification

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 on configuring the background verification rate.

Attribute Name	Code	Description
Reset Corrupt Objects Count	VCCR	Reset the counter for corrupt replicated object data found during background verification. This option can be used to clear the Corrupt Objects Detected (OCOR) alarm condition. For details, see the instructions for monitoring and troubleshooting StorageGRID.
Delete Quarantined Objects	OQRT	<p>Delete corrupt objects from the quarantine directory, reset the count of quarantined objects to zero, and clear the Quarantined Objects Detected (OQRT) alarm. This option is used after corrupt objects have been automatically restored by the StorageGRID system.</p> <p>If a Lost Objects alarm is triggered, technical support might want to access the quarantined objects. In some cases, quarantined objects might be useful for data recovery or for debugging the underlying issues that caused the corrupt object copies.</p>

LDR > Erasure Coding

Attribute Name	Code	Description
Reset Writes Failure Count	RSWF	Reset the counter for write failures of erasure-coded object data to the Storage Node.
Reset Reads Failure Count	RSRF	Reset the counter for read failures of erasure-coded object data from the Storage Node.
Reset Deletes Failure Count	RSDF	Reset the counter for delete failures of erasure-coded object data from the Storage Node.
Reset Corrupt Copies Detected Count	RSCC	Reset the counter for the number of corrupt copies of erasure-coded object data on the Storage Node.
Reset Corrupt Fragments Detected Count	RSCD	Reset the counter for corrupt fragments of erasure-coded object data on the Storage Node.
Reset Missing Fragments Detected Count	RSMD	Reset the counter for missing fragments of erasure-coded object data on the Storage Node. Use only after 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	<p>Select to disable inbound replication as part of a maintenance or testing procedure. Leave unchecked during normal operation.</p> <p>When inbound replication is disabled, objects can be retrieved from the Storage Node for copying to other locations in the StorageGRID system, but objects cannot be copied to this Storage Node from other locations: the LDR service is read-only.</p>
Disable Outbound Replication	DSOR	<p>Select to disable outbound replication (including content requests for HTTP retrievals) as part of a maintenance or testing procedure. Leave unchecked during normal operation.</p> <p>When outbound replication is disabled, objects can be copied to this Storage Node, but objects cannot be retrieved from the Storage Node to be copied to other locations in the StorageGRID system. The LDR service is write-only.</p>

Related information

[Monitor and troubleshoot](#)

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

What an Admin Node is

Admin Nodes provide management services such as system configuration, monitoring, and logging. Each grid must have one primary Admin Node and might have any number of non-primary Admin Nodes for redundancy.

When you sign in to the Grid Manager or the Tenant Manager, you are connecting to an Admin Node. You can connect to any Admin Node, and each Admin Node displays a similar view of the StorageGRID system. However, maintenance procedures must be performed using the primary Admin Node.

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

Admin Nodes host the following services:

- AMS service
- CMN service
- NMS service
- Prometheus service
- Load Balancer and High Availability services (to support S3 and Swift client traffic)

Admin Nodes also support the Management Application Program Interface (mgmt-api) to process requests from the Grid Management API and the Tenant Management API. See [Use the Grid Management API](#).

What the AMS service is

The Audit Management System (AMS) service tracks system activity and events.

What the CMN service is

The Configuration Management Node (CMN) service manages system-wide configurations of connectivity and protocol features needed by all services. In addition, the CMN service is used to run and monitor grid tasks. There is only one CMN service per StorageGRID deployment. The Admin Node that hosts the CMN service is known as the primary Admin Node.

What the NMS service is

The Network Management System (NMS) service powers the monitoring, reporting, and configuration options that are displayed through the Grid Manager, the StorageGRID system's browser-based interface.


What the Prometheus service is

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

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 messages are still sent. However, having multiple Admin Nodes does not provide failover protection except for notifications and AutoSupport messages. In particular, alarm acknowledgments made from one Admin Node are not copied to other Admin Nodes.



There are two options for continuing to view and configure the StorageGRID system if an Admin Node fails:

- Web clients can reconnect to any other available Admin Node.
- If a system administrator has configured a high availability group of Admin Nodes, web clients can continue to access the Grid Manager or the Tenant Manager using the virtual IP address of the HA group. See [Manage high availability groups](#).



When using an HA group, access is interrupted if the Master Admin Node fails. Users must sign in again after the virtual IP address of the HA group fails over to another Admin Node in the group.

Some maintenance tasks can only be performed using the primary Admin Node. If the primary Admin Node fails, it must be recovered before the StorageGRID system is fully functional again.


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.

What you'll need

- 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 **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.

Select a preferred sender

If your StorageGRID deployment includes multiple Admin Nodes, you can select which Admin Node should be the preferred sender of notifications. By default, the primary Admin Node is selected, but any Admin Node can be the preferred sender.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

About this task

The **CONFIGURATION > System > Display options** page shows which Admin Node is currently selected to be the preferred sender. The primary Admin Node is selected by default.

Under normal system operations, only the preferred sender sends the following notifications:

- AutoSupport messages
- SNMP notifications
- Alert emails
- Alarm emails (legacy system)

However, all other Admin Nodes (standby senders) monitor the preferred sender. If a problem is detected, a standby sender can also send these notifications.

Both the preferred sender and a standby sender might send notifications in these cases:

- If Admin Nodes become “islanded” from each other, both the preferred sender and the standby senders will attempt to send notifications, and multiple copies of notifications might be received.
- After a standby sender detects problems with the preferred sender and starts sending notifications, the

preferred sender might regain its ability to send notifications. If this occurs, duplicate notifications might be sent. The standby sender will stop sending notifications when it no longer detects errors on the preferred sender.



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

Steps

1. Select **CONFIGURATION > System > Display options**.
2. From the Display Options menu, select **Options**.
3. Select the Admin Node you want to set as the preferred sender from the drop-down list.



Display Options

Updated: 2017-08-30 16:31:10 MDT

Current Sender	ADMIN-DC1-ADM1
Preferred Sender	ADMIN-DC1-ADM1
GUI Inactivity Timeout	900
Notification Suppress All	<input type="checkbox"/>

Apply Changes

4. Select **Apply Changes**.

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


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
Updated: 2009-03-09 10:12:17 PDT

NMS Interface Engine Status: Connected 


Connected Services: 15 


E-mail Notification Events


E-mail Notifications Status: No Errors 

E-mail Notifications Queued: 0 

Database Connection Pool

Maximum Supported Capacity: 100 

Remaining Capacity: 95 % 


Active Connections: 5 

Notifications are processed through the email notifications queue and are sent to the mail server one after another in the order they are triggered. If there is a problem (for example, a network connection error) and the mail server is unavailable when the attempt is made to send the notification, a best effort attempt to resend the notification to the mail server continues for a period of 60 seconds. If the notification is not sent to the mail server after 60 seconds, the notification is dropped from the notifications queue and an attempt to send the next notification in the queue is made. Because notifications can be dropped from the notifications queue without being sent, it is possible that an alarm can be triggered without a notification being sent. In the event that a notification is dropped from the queue without being sent, the MINS (E-mail Notification Status) Minor alarm is triggered.

How Admin Nodes show acknowledged alarms (legacy system)

When you acknowledge an alarm on one Admin Node, the acknowledged alarm is not copied to any other Admin Node. Because acknowledgments are not copied to other Admin Nodes, the Grid Topology tree might not look the same for each Admin Node.

This difference can be useful when connecting web clients. Web clients can have different views of the StorageGRID system based on the administrator needs.



Note that notifications are sent from the Admin Node where the acknowledgment occurs.

Configure audit client access

The Admin Node, through the Audit Management System (AMS) service, logs all audited system events to a log file available through the audit share, which is added to each Admin Node at installation. For easy access to audit logs, you can configure client access to audit shares for both CIFS and NFS.

The StorageGRID system uses positive acknowledgment to prevent loss of audit messages before they are written to the log file. A message remains queued at a service until the AMS service or an intermediate audit relay service has acknowledged control of it.

For more information, see [Review audit logs](#).



Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release. If you have the option to use CIFS or NFS, choose NFS.

Configure audit clients for CIFS

The procedure used to configure an audit client depends on the authentication method: Windows Workgroup or Windows Active Directory (AD). When added, the audit share is automatically enabled as a read-only share.



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

Configure audit clients for Workgroup

Perform this procedure for each Admin Node in a StorageGRID deployment from which you want to retrieve audit messages.

What you'll need

- You have the `Passwords.txt` file with the root/admin account password (available in the SAID package).
- You have the `Configuration.txt` file (available in the SAID package).

About this task

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

Steps

1. Log in to the primary Admin Node:
 - a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
 - b. Enter the password listed in the `Passwords.txt` file.
 - c. Enter the following command to switch to root: `su -`
 - d. Enter the password listed in the `Passwords.txt` file.

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

2. Confirm that all services have a state of Running or Verified: `storagegrid-status`

If all services are not Running or Verified, resolve issues before continuing.

3. Return to the command line, press **Ctrl+C**.
4. Start the CIFS configuration utility: `config_cifs.rb`

Shares	Authentication	Config	

add-audit-share	set-authentication	validate-config	
enable-disable-share	set-netbios-name	help	
add-user-to-share	join-domain	exit	
remove-user-from-share	add-password-server		
modify-group	remove-password-server		
	add-wins-server		
	remove-wins-server		

5. Set the authentication for the Windows Workgroup:

If authentication has already been set, an advisory message appears. If authentication has already been set, go to the next step.

- a. Enter: `set-authentication`
- b. When prompted for Windows Workgroup or Active Directory installation, enter: `workgroup`
- c. When prompted, enter a name of the Workgroup: `workgroup_name`
- d. When prompted, create a meaningful NetBIOS name: `netbios_name`

or

Press **Enter** to use the Admin Node's hostname as the NetBIOS name.

The script restarts the Samba server and changes are applied. This should take less than one minute. After setting authentication, add an audit client.

- e. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

6. Add an audit client:

- a. Enter: `add-audit-share`



The share is automatically added as read-only.

- b. When prompted, add a user or group: `user`

- c. When prompted, enter the audit user name: *audit_user_name*
- d. When prompted, enter a password for the audit user: *password*
- e. When prompted, re-enter the same password to confirm it: *password*
- f. When prompted, press **Enter**.

The CIFS configuration utility is displayed.



There is no need to enter a directory. The audit directory name is predefined.

- 7. If more than one user or group is permitted to access the audit share, add the additional users:

- a. Enter: `add-user-to-share`

A numbered list of enabled shares is displayed.

- b. When prompted, enter the number of the audit-export share: *share_number*

- c. When prompted, add a user or group: *user*

or group

- d. When prompted, enter the name of the audit user or group: *audit_user* or *audit_group*

- e. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

- f. Repeat these substeps for each additional user or group that has access to the audit share.

- 8. Optionally, verify your configuration: `validate-config`

The services are checked and displayed. You can safely ignore the following messages:

```
Can't find include file /etc/samba/includes/cifs-interfaces.inc
Can't find include file /etc/samba/includes/cifs-filesystem.inc
Can't find include file /etc/samba/includes/cifs-custom-config.inc
Can't find include file /etc/samba/includes/cifs-shares.inc
rlimit_max: increasing rlimit_max (1024) to minimum Windows limit
(16384)
```

- a. When prompted, press **Enter**.

The audit client configuration is displayed.

- b. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

- 9. Close the CIFS configuration utility: `exit`

- 10. Start the Samba service: `service smb start`

11. If the StorageGRID deployment is a single site, go to the next step.

or

Optionally, if the StorageGRID deployment includes Admin Nodes at other sites, enable these audit share as required:

- a. Remotely log in to a site's Admin Node:
 - i. Enter the following command: `ssh admin@grid_node_IP`
 - ii. Enter the password listed in the `Passwords.txt` file.
 - iii. Enter the following command to switch to root: `su -`
 - iv. Enter the password listed in the `Passwords.txt` file.
- b. Repeat the steps to configure the audit share for each additional Admin Node.
- c. Close the remote secure shell login to the remote Admin Node: `exit`

12. Log out of the command shell: `exit`

Configure audit clients for Active Directory

Perform this procedure for each Admin Node in a StorageGRID deployment from which you want to retrieve audit messages.

What you'll need

- You have the `Passwords.txt` file with the root/admin account password (available in the SAID package).
- You have the CIFS Active Directory username and password.
- You have the `Configuration.txt` file (available in the SAID package).



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

Steps

1. Log in to the primary Admin Node:
 - a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
 - b. Enter the password listed in the `Passwords.txt` file.
 - c. Enter the following command to switch to root: `su -`
 - d. Enter the password listed in the `Passwords.txt` file.

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

2. Confirm that all services have a state of Running or Verified: `storagegrid-status`

If all services are not Running or Verified, resolve issues before continuing.

3. Return to the command line, press **Ctrl+C**.
4. Start the CIFS configuration utility: `config_cifs.rb`

Shares	Authentication	Config
add-audit-share	set-authentication	validate-config
enable-disable-share	set-netbios-name	help
add-user-to-share	join-domain	exit
remove-user-from-share	add-password-server	
modify-group	remove-password-server	
	add-wins-server	
	remove-wins-server	

5. Set the authentication for Active Directory: `set-authentication`

In most deployments, you must set the authentication before adding the audit client. If authentication has already been set, an advisory message appears. If authentication has already been set, go to the next step.

- When prompted for Workgroup or Active Directory installation: `ad`
- When prompted, enter the name of the AD domain (short domain name).
- When prompted, enter the domain controller's IP address or DNS hostname.
- When prompted, enter the full domain realm name.

Use uppercase letters.

- When prompted to enable winbind support, type `y`.

Winbind is used to resolve user and group information from AD servers.

- When prompted, enter the NetBIOS name.
- When prompted, press **Enter**.

The CIFS configuration utility is displayed.

6. Join the domain:

- If not already started, start the CIFS configuration utility: `config_cifs.rb`
- Join the domain: `join-domain`
- You are prompted to test if the Admin Node is currently a valid member of the domain. If this Admin Node has not previously joined the domain, enter: `no`
- When prompted, provide the Administrator's username: `administrator_username`

where `administrator_username` is the CIFS Active Directory username, not the StorageGRID username.

- When prompted, provide the Administrator's password: `administrator_password`

where `administrator_password` is the CIFS Active Directory password, not the StorageGRID password.

password.

- f. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

7. Verify that you have correctly joined the domain:

- a. Join the domain: `join-domain`

- b. When prompted to test if the server is currently a valid member of the domain, enter: `y`

If you receive the message “Join is OK,” you have successfully joined the domain. If you do not get this response, try setting authentication and joining the domain again.

- c. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

8. Add an audit client: `add-audit-share`

- a. When prompted to add a user or group, enter: `user`

- b. When prompted to enter the audit user name, enter the audit user name.

- c. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

9. If more than one user or group is permitted to access the audit share, add additional users: `add-user-to-share`

A numbered list of enabled shares is displayed.

- a. Enter the number of the audit-export share.

- b. When prompted to add a user or group, enter: `group`

You are prompted for the audit group name.

- c. When prompted for the audit group name, enter the name of the audit user group.

- d. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

- e. Repeat this step for each additional user or group that has access to the audit share.

10. Optionally, verify your configuration: `validate-config`

The services are checked and displayed. You can safely ignore the following messages:

- Can't find include file `/etc/samba/includes/cifs-interfaces.inc`
- Can't find include file `/etc/samba/includes/cifs-filesystem.inc`
- Can't find include file `/etc/samba/includes/cifs-interfaces.inc`
- Can't find include file `/etc/samba/includes/cifs-custom-config.inc`

- Can't find include file `/etc/samba/includes/cifs-shares.inc`
- `rlimit_max`: increasing `rlimit_max` (1024) to minimum Windows limit (16384)



Do not combine the setting 'security=ads' with the 'password server' parameter. (by default Samba will discover the correct DC to contact automatically).

- When prompted, press **Enter** to display the audit client configuration.
- When prompted, press **Enter**.

The CIFS configuration utility is displayed.

- Close the CIFS configuration utility: `exit`
- If the StorageGRID deployment is a single site, go to the next step.

or

Optionally, if the StorageGRID deployment includes Admin Nodes at other sites, enable these audit shares as required:

- Remotely log in to a site's Admin Node:
 - Enter the following command: `ssh admin@grid_node_IP`
 - Enter the password listed in the `Passwords.txt` file.
 - Enter the following command to switch to root: `su -`
 - Enter the password listed in the `Passwords.txt` file.
- Repeat these steps to configure the audit shares for each Admin Node.
- Close the remote secure shell login to the Admin Node: `exit`

- Log out of the command shell: `exit`

Add a user or group to a CIFS audit share

You can add a user or group to a CIFS audit share that is integrated with AD authentication.

What you'll need

- You have the `Passwords.txt` file with the root/admin account password (available in the SAID package).
- You have the `Configuration.txt` file (available in the SAID package).

About this task

The following procedure is for an audit share integrated with AD authentication.



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

Steps

- Log in to the primary Admin Node:

- a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
- b. Enter the password listed in the `Passwords.txt` file.
- c. Enter the following command to switch to root: `su -`
- d. Enter the password listed in the `Passwords.txt` file.

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

2. Confirm that all services have a state of Running or Verified. Enter: `storagegrid-status`

If all services are not Running or Verified, resolve issues before continuing.

3. Return to the command line, press **Ctrl+C**.
4. Start the CIFS configuration utility: `config_cifs.rb`

Shares	Authentication	Config	

add-audit-share	set-authentication	validate-config	
enable-disable-share	set-netbios-name	help	
add-user-to-share	join-domain	exit	
remove-user-from-share	add-password-server		
modify-group	remove-password-server		
	add-wins-server		
	remove-wins-server		

5. Start adding a user or group: `add-user-to-share`

A numbered list of audit shares that have been configured is displayed.

6. When prompted, enter the number for the audit share (audit-export): `audit_share_number`

You are asked if you would like to give a user or a group access to this audit share.

7. When prompted, add a user or group: `user or group`
8. When prompted for the user or group name for this AD audit share, enter the name.

The user or group is added as read-only for the audit share both in the server's operating system and in the CIFS service. The Samba configuration is reloaded to enable the user or group to access the audit client share.

9. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

10. Repeat these steps for each user or group that has access to the audit share.
11. Optionally, verify your configuration: `validate-config`

The services are checked and displayed. You can safely ignore the following messages:

- Can't find include file /etc/samba/includes/cifs-interfaces.inc
- Can't find include file /etc/samba/includes/cifs-filesystem.inc
- Can't find include file /etc/samba/includes/cifs-custom-config.inc
- Can't find include file /etc/samba/includes/cifs-shares.inc
 - a. When prompted, press **Enter** to display the audit client configuration.
 - b. When prompted, press **Enter**.

12. Close the CIFS configuration utility: `exit`

13. Determine if you need to enable additional audit shares, as follows:

- If the StorageGRID deployment is a single site, go to the next step.
- If the StorageGRID deployment includes Admin Nodes at other sites, enable these audit shares as required:
 - a. Remotely log in to a site's Admin Node:
 - i. Enter the following command: `ssh admin@grid_node_IP`
 - ii. Enter the password listed in the `Passwords.txt` file.
 - iii. Enter the following command to switch to root: `su -`
 - iv. Enter the password listed in the `Passwords.txt` file.
 - b. Repeat these steps to configure the audit shares for each Admin Node.
 - c. Close the remote secure shell login to the remote Admin Node: `exit`

14. Log out of the command shell: `exit`

Remove a user or group from a CIFS audit share

You cannot remove the last user or group permitted to access the audit share.

What you'll need

- You have the `Passwords.txt` file with the root account passwords (available in the SAID package).
- You have the `Configuration.txt` file (available in the SAID package).

About this task

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

Steps

1. Log in to the primary Admin Node:

- a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
- b. Enter the password listed in the `Passwords.txt` file.
- c. Enter the following command to switch to root: `su -`
- d. Enter the password listed in the `Passwords.txt` file.

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

2. Start the CIFS configuration utility: `config_cifs.rb`

Shares	Authentication	Config	

add-audit-share	set-authentication	validate-config	
enable-disable-share	set-netbios-name	help	
add-user-to-share	join-domain	exit	
remove-user-from-share	add-password-server		
modify-group	remove-password-server		
	add-wins-server		
	remove-wins-server		

3. Start removing a user or group: `remove-user-from-share`

A numbered list of available audit shares for the Admin Node is displayed. The audit share is labeled audit-export.

4. Enter the number of the audit share: `audit_share_number`

5. When prompted to remove a user or a group: `user` or `group`

A numbered list of users or groups for the audit share is displayed.

6. Enter the number corresponding to the user or group you want to remove: `number`

The audit share is updated, and the user or group is no longer permitted access to the audit share. For example:

```
Enabled shares
 1. audit-export
Select the share to change: 1
Remove user or group? [User/group]: User
Valid users for this share
 1. audituser
 2. newaudituser
Select the user to remove: 1

Removed user "audituser" from share "audit-export".

Press return to continue.
```

7. Close the CIFS configuration utility: `exit`

8. If the StorageGRID deployment includes Admin Nodes at other sites, disable the audit share at each site as required.

9. Log out of each command shell when configuration is complete: `exit`

Change a CIFS audit share user or group name

You can change the name of a user or a group for a CIFS audit share by adding a new user or group and then deleting the old one.

About this task

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

Steps

1. Add a new user or group with the updated name to the audit share.
2. Delete the old user or group name.

Related information

- [Add a user or group to a CIFS audit share](#)
- [Remove a user or group from a CIFS audit share](#)

Verify CIFS audit integration

The audit share is read-only. Log files are intended to be read by computer applications and verification does not include opening a file. It is considered sufficient verification that the audit log files appear in a Windows Explorer window. Following connection verification, close all windows.

Configure the audit client for NFS

The audit share is automatically enabled as a read-only share.

What you'll need

- You have the `Passwords.txt` file with the root/admin password (available in the SAID package).
- You have the `Configuration.txt` file (available in the SAID 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 `Passwords.txt` file.

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

2. Confirm that all services have a state of Running or Verified. Enter: `storagegrid-status`

If any services are not listed as Running or Verified, resolve issues before continuing.

3. Return to the command line. Press **Ctrl+C**.
4. Start the NFS configuration utility. Enter: `config_nfs.rb`

```
-----  
| Shares                | Clients                | Config                |  
-----  
| add-audit-share       | add-ip-to-share        | validate-config       |  
| enable-disable-share  | remove-ip-from-share   | refresh-config        |  
|                       |                       | help                  |  
|                       |                       | exit                  |  
-----
```

5. Add the audit client: `add-audit-share`
 - a. When prompted, enter the audit client's IP address or IP address range for the audit share:
`client_IP_address`
 - b. When prompted, press **Enter**.
6. If more than one audit client is permitted to access the audit share, add the IP address of the additional user: `add-ip-to-share`
 - a. Enter the number of the audit share: `audit_share_number`
 - b. When prompted, enter the audit client's IP address or IP address range for the audit share:
`client_IP_address`
 - c. When prompted, press **Enter**.

The NFS configuration utility is displayed.
 - d. Repeat these substeps for each additional audit client that has access to the audit share.
7. Optionally, verify your configuration.
 - a. Enter the following: `validate-config`

The services are checked and displayed.
 - b. When prompted, press **Enter**.

The NFS configuration utility is displayed.
 - c. Close the NFS configuration utility: `exit`
8. Determine if you must enable audit shares at other sites.
 - If the StorageGRID deployment is a single site, go to the next step.
 - If the StorageGRID deployment includes Admin Nodes at other sites, enable these audit shares as required:

- a. Remotely log in to the site's Admin Node:
 - i. Enter the following command: `ssh admin@grid_node_IP`
 - ii. Enter the password listed in the `Passwords.txt` file.
 - iii. Enter the following command to switch to root: `su -`
 - iv. Enter the password listed in the `Passwords.txt` file.
- b. Repeat these steps to configure the audit shares for each additional Admin Node.
- c. Close the remote secure shell login to the remote Admin Node. Enter: `exit`

9. Log out of the command shell: `exit`

NFS audit clients are granted access to an audit share based on their IP address. Grant access to the audit share to a new NFS audit client by adding its IP address to the share, or remove an existing audit client by removing its IP address.

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.

What you'll need

- You have the `Passwords.txt` file with the root/admin account password (available in the SAID package).
- You have the `Configuration.txt` file (available in the SAID 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 `Passwords.txt` file.

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

2. Start the NFS configuration utility: `config_nfs.rb`

Shares	Clients	Config
add-audit-share	add-ip-to-share	validate-config
enable-disable-share	remove-ip-from-share	refresh-config
		help
		exit

3. Enter: `add-ip-to-share`

A list of NFS audit shares enabled on the Admin Node is displayed. The audit share is listed as:
`/var/local/audit/export`

4. Enter the number of the audit share: `audit_share_number`

5. When prompted, enter the audit client's IP address or IP address range for the audit share:
`client_IP_address`

The audit client is added to the audit share.

6. When prompted, press **Enter**.

The NFS configuration utility is displayed.

7. Repeat the steps for each audit client that should be added to the audit share.

8. Optionally, verify your configuration: `validate-config`

The services are checked and displayed.

a. When prompted, press **Enter**.

The NFS configuration utility is displayed.

9. Close the NFS configuration utility: `exit`

10. If the StorageGRID deployment is a single site, go to the next step.

Otherwise, if the StorageGRID deployment includes Admin Nodes at other sites, optionally enable these audit shares as required:

a. Remotely log in to a site's Admin Node:

- i. Enter the following command: `ssh admin@grid_node_IP`
- ii. Enter the password listed in the `Passwords.txt` file.
- iii. Enter the following command to switch to root: `su -`
- iv. Enter the password listed in the `Passwords.txt` file.

b. Repeat these steps to configure the audit shares for each Admin Node.

c. Close the remote secure shell login to the remote Admin Node: `exit`

11. Log out of the command shell: `exit`

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.

Steps

1. Verify connectivity (or variant for the client system) using the client-side IP address of the Admin Node hosting the AMS service. Enter: `ping IP_address`

Verify that the server responds, indicating connectivity.

2. Mount the audit read-only share using a command appropriate to the client operating system. A sample Linux command is (enter on one line):

```
mount -t nfs -o hard,intr Admin_Node_IP_address:/var/local/audit/export  
myAudit
```

Use the IP address of the Admin Node hosting the AMS service and the predefined share name for the audit system. The mount point can be any name selected by the client (for example, *myAudit* in the previous command).

3. Verify that the files are available from the audit share. Enter: `ls myAudit /*`

where *myAudit* is the mount point of the audit share. There should be at least one log file listed.

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.

What you'll need

- You have the `Passwords.txt` file with the root/admin account password (available in the SAID package).
- You have the `Configuration.txt` file (available in the SAID package).

About this task

You cannot remove the last IP address permitted to access the audit share.

Steps

1. Log in to the primary Admin Node:
 - a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
 - b. Enter the password listed in the `Passwords.txt` file.
 - c. Enter the following command to switch to root: `su -`
 - d. Enter the password listed in the `Passwords.txt` file.

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

2. Start the NFS configuration utility: `config_nfs.rb`

Shares	Clients	Config	

add-audit-share	add-ip-to-share	validate-config	
enable-disable-share	remove-ip-from-share	refresh-config	
		help	
		exit	

3. Remove the IP address from the audit share: `remove-ip-from-share`

A numbered list of audit shares configured on the server is displayed. The audit share is listed as:
`/var/local/audit/export`

4. Enter the number corresponding to the audit share: `audit_share_number`

A numbered list of IP addresses permitted to access the audit share is displayed.

5. Enter the number corresponding to the IP address you want to remove.

The audit share is updated, and access is no longer permitted from any audit client with this IP address.

6. When prompted, press **Enter**.

The NFS configuration utility is displayed.

7. Close the NFS configuration utility: `exit`

8. If your StorageGRID deployment is a multiple data center site deployment with additional Admin Nodes at the other sites, disable these audit shares as required:

a. Remotely log in to each site's Admin Node:

i. Enter the following command: `ssh admin@grid_node_IP`

ii. Enter the password listed in the `Passwords.txt` file.

iii. Enter the following command to switch to root: `su -`

iv. Enter the password listed in the `Passwords.txt` file.

b. Repeat these steps to configure the audit shares for each additional Admin Node.

c. Close the remote secure shell login to the remote Admin Node: `exit`

9. Log out of the command shell: `exit`

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

What an Archive Node is

Optionally, each StorageGRID data center site can be deployed with an Archive Node, which allows you to connect to a targeted external archival storage system, such as Tivoli Storage Manager (TSM).

The Archive Node provides an interface through which you can target an external archival storage system for the long term storage of object data. The Archive Node also monitors this connection and the transfer of object data between the StorageGRID system and the targeted external archival storage system.



The screenshot displays the StorageGRID Webconsole interface. On the left, the 'Grid Topology' pane shows a hierarchical view of the deployment, including Data Center 1, Data Center 2, and Data Center 3. Under Data Center 1, the 'DC1-ARC1-98-165' node is highlighted, showing its sub-components: SSM, ARC, Replication, Store, Retrieve, Target, Events, and Resources. The main pane shows the 'Overview' tab for the selected ARC node. The title is 'Overview: ARC (DC1-ARC1-98-165) - ARC', updated on 2015-09-30 at 10:29:18 PDT. The overview table lists various status metrics, all of which are 'Online' or 'No Errors'. Below the table, the 'Node Information' section provides details about the device type, version, build, node ID, and site ID.

Overview: ARC (DC1-ARC1-98-165) - ARC		Updated: 2015-09-30 10:29:18 PDT
ARC State:	Online	
ARC Status:	No Errors	
Tivoli Storage Manager State:	Online	
Tivoli Storage Manager Status:	No Errors	
Store State:	Online	
Store Status:	No Errors	
Retrieve State:	Online	
Retrieve Status:	No Errors	
Inbound Replication Status:	No Errors	
Outbound Replication Status:	No Errors	

Node Information	
Device Type:	Archive Node
Version:	10.2.0
Build:	20150928.2133.a27b3ab
Node ID:	19002524
Site ID:	10

After configuring connections to the external target, you can configure the Archive Node to optimize TSM performance, take an Archive Node offline when a TSM server is nearing capacity or unavailable, and configure replication and retrieve settings. You can also set Custom alarms for the Archive Node.

Object data that cannot be deleted, but is not regularly accessed, can at any time be moved off of a Storage Node's spinning disks and onto external archival storage such as the cloud or tape. This archiving of object data is accomplished through the configuration of a data center site's Archive Node and then the configuration of ILM rules where this Archive Node is selected as the "target" for content placement instructions. The Archive Node does not manage archived object data itself; this is achieved by the external archive device.



Object metadata is not archived, but remains on Storage Nodes.

What the ARC service is

The Archive (ARC) service on Archive Nodes provides the management interface you can use to configure connections to external archival storage, such as tape through TSM middleware.

It is the ARC service that interacts with an external archival storage system, sending object data for near-line storage and performing retrievals when a client application requests an archived object. When a client application requests an archived object, a Storage Node requests the object data from the ARC service. The ARC service makes a request to the external archival storage system, which retrieves the requested object data and sends it to the ARC service. The ARC service verifies the object data and forwards it to the Storage Node, which in turn returns the object to the requesting client application.

Requests for object data archived to tape through TSM middleware are managed for efficiency of retrievals. Requests can be ordered so that objects stored in sequential order on tape are requested in that same sequential order. Requests are then queued for submission to the storage device. Depending upon the archival device, multiple requests for objects on different volumes can be processed simultaneously.

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.



Moving objects from an Archive Node to an external archival storage system through the S3 API has been replaced by ILM Cloud Storage Pools, which offer more functionality. The **Cloud Tiering - Simple Storage Service (S3)** option is still supported, but you might prefer to implement Cloud Storage Pools instead.

If you are currently using an Archive Node with the **Cloud Tiering - Simple Storage Service (S3)** option, consider migrating your objects to a Cloud Storage Pool. See the instructions for [managing objects with ILM](#).

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.



Moving objects from an Archive Node to an external archival storage system through the S3 API has been replaced by ILM Cloud Storage Pools, which offer more functionality. The **Cloud Tiering - Simple Storage Service (S3)** option is still supported, but you might prefer to implement Cloud Storage Pools instead.

If you are currently using an Archive Node with the **Cloud Tiering - Simple Storage Service (S3)** option, consider migrating your objects to a Cloud Storage Pool. See [Manage objects with ILM](#).

What you'll need

- 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 cannot 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.

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 (98-127) - Target
Updated: 2015-09-24 15:48:22 PDT

Target Type: Cloud Tiering - Simple Storage Service (S3)

Cloud Tiering (S3) Account

Bucket Name:	<input type="text" value="name"/>
Region:	Virginia or Pacific Northwest (us-east-1)
Endpoint:	<input type="text" value="https://10.10.10.123:8082"/> <input type="checkbox"/> Use AWS
Endpoint Authentication:	<input type="checkbox"/>
Access Key:	<input type="text" value="ABCD123EFG45AB"/>
Secret Access Key:	<input type="password" value="•••••"/>
Storage Class:	Standard (Default)

Apply Changes 

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 unselect the check box to disable endpoint SSL certificate and hostname verification for the targeted external archival storage system. If another instance of a StorageGRID system is the target archival storage device and the system is configured with publicly signed certificates, you can keep the check box selected.
- **Storage Class:** Select **Standard (Default)** for regular storage. Select **Reduced Redundancy** only for objects that can be easily recreated. **Reduced Redundancy** provides lower cost storage with less reliability. If the targeted archival storage system is another instance of the StorageGRID system, **Storage Class** controls how many interim copies of the object are made at ingest on the target system, if dual commit is used when objects are ingested there.

6. Select **Apply Changes**.

The specified configuration settings are validated and applied to your StorageGRID system. Once configured, the target cannot 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.

What you'll need

- 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.

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 (98-127) - Target

Updated: 2015-09-24 15:48:22 PDT

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	<input type="checkbox"/>	Use AWS
Endpoint Authentication:	<input type="checkbox"/>		
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 cannot 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.

What you'll need


- You must be signed in to the Grid Manager using a [supported web browser](#).
- You must have specific access permissions.
- The Archive Node must be configured.

About this task

You can effectively take the Archive Node offline by changing the Cloud Tiering Service State to **Read-Write Disabled**.


Steps

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **Archive Node > ARC**.
3. Select **Configuration > Main**.




Overview Alarms Reports Configuration

Main Alarms

 Configuration: ARC (98-127) - ARC
Updated: 2015-09-24 17:18:29 PDT

ARC State Online

Cloud Tiering Service State Read-Write Enabled

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.

What you'll need

- 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 > Store**.
3. Select **Configuration > Main**.




Overview Alarms Reports Configuration

Main Alarms

 Configuration: ARC (98-127) - Store
Updated: 2015-09-29 17:54:42 PDT

Reset Store Failure Count ☐

Apply Changes 

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, consider migrating your objects to a Cloud Storage Pool instead. Cloud Storage Pools provide a scalable approach that takes advantage of all of the Storage Nodes in your StorageGRID system.

What you'll need

- You 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 cannot also be stored within StorageGRID. If the ILM rule you are currently using for Cloud Tiering is configured to store objects in multiple locations at the same time, consider whether you still want to perform this optional migration because you will lose that functionality. If you continue with this migration, you must create new rules instead of cloning the existing ones.

Steps

1. Create a Cloud Storage Pool.

Use a new S3 bucket for the Cloud Storage Pool to ensure it contains only the data managed by the Cloud Storage Pool.

2. Locate any ILM rules in the active ILM policy that cause objects to be stored in the Cloud Tiering bucket.
3. Clone each of these rules.
4. In the cloned rules, change the placement location to the new Cloud Storage Pool.
5. Save the cloned rules.
6. Create a new policy that uses the new rules.
7. Simulate and activate the new policy.

When the new policy is activated and ILM evaluation occurs, the objects are moved from the S3 bucket configured for Cloud Tiering to the S3 bucket configured for the Cloud Storage Pool. The usable space on

the grid is not affected. After the objects are moved to the Cloud Storage Pool, they are removed from the Cloud Tiering bucket.

Related information

[Manage objects with ILM](#)

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.

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 a number of settings.

What you'll need

- 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 (DC1-ARC1-98-165) - Target

Updated: 2015-09-28 09:56:36 PDT

Target Type:

Tivoli Storage Manager (TSM)

Tivoli Storage Manager State:

Online

Target (TSM) Account

Server IP or Hostname:

10.10.10.123

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

Apply Changes



4. From the **Target Type** drop-down list, select **Tivoli Storage Manager (TSM)**.
5. For the **Tivoli Storage Manager State**, select **Offline** to prevent retrievals from the TSM middleware server.

By default, the Tivoli Storage Manager State is set to Online, which means that the Archive Node is able to retrieve object data from the TSM middleware server.

6. Complete the following information:
 - **Server IP or Hostname:** Specify the IP address or fully qualified domain name of the TSM middleware server used by the ARC service. The default IP address is 127.0.0.1.
 - **Server Port:** Specify the port number on the TSM middleware server that the ARC service will connect to. The default is 1500.
 - **Node Name:** Specify the name of the Archive Node. You must enter the name (arc-user) that you registered on the TSM middleware server.
 - **User Name:** Specify the user name the ARC service uses to log in to the TSM server. Enter the default user name (arc-user) or the administrative user you specified for the Archive Node.
 - **Password:** Specify the password used by the ARC service to log in to the TSM server.
 - **Management Class:** Specify the default management class to use if a management class is not specified when the object is being saved to the StorageGRID system, or the specified management class is not defined on the TSM middleware server.
 - **Number of Sessions:** Specify the number of tape drives on the TSM middleware server that are dedicated to the Archive Node. The Archive Node concurrently creates a maximum of one session per mount point plus a small number of additional sessions (less than five).

You must change this value to be the same as the value set for MAXNUMMP (maximum number of mount points) when the Archive Node was registered or updated. (In the register command, the default value of MAXNUMMP used is 1, if no value is set.)

You must also change the value of MAXSESSIONS for the TSM server to a number that is at least as large as the Number of Sessions set for the ARC service. The default value of MAXSESSIONS on the TSM server is 25.

- **Maximum Retrieve Sessions:** Specify the maximum number of sessions that the ARC service can open to the TSM middleware server for retrieve operations. In most cases, the appropriate value is Number of Sessions minus Maximum Store Sessions. If you need to share one tape drive for storage and retrieval, specify a value equal to the Number of Sessions.
- **Maximum Store Sessions:** Specify the maximum number of concurrent sessions that the ARC service can open to the TSM middleware server for archive operations.

This value should be set to one except when the targeted archival storage system is full and only retrievals can be performed. Set this value to zero to use all sessions for retrievals.

7. 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.

What you'll need

- 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.

Steps

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **Archive Node > ARC > Target**.
3. Select **Configuration > Main**.
4. Change **Maximum Retrieve Sessions** to be the same as **Number of Sessions**.

Overview
Alarms
Reports
Configuration

Main
Alarms



Configuration: ARC (DC1-ARC1-98-165) - Target
Updated: 2015-09-28 09:56:36 PDT

Target Type:

Tivoli Storage Manager (TSM)

Tivoli Storage Manager State:

Online

Target (TSM) Account

Server IP or Hostname:

10.10.10.123

Server Port:

1500

Node Name:

ARC-USER

User Name:

arc-user

Password:

••••••

Management Class:

sg-mgmtclass

Number of Sessions:


2

Maximum Retrieve Sessions:

2

Maximum Store Sessions:

1

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.

What you'll need


- 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 > Store**.
3. Select **Configuration > Main**.

Overview
Alarms
Reports
Configuration

Main
Alarms



Configuration: ARC (DC1-ARC1-98-165) - Store
Updated: 2015-09-29 17:10:12 PDT

Store State


Online

Archive Store Disabled on Startup

☐

Reset Store Failure Count

☐

Apply Changes


4. Modify the following settings, as necessary:

- Store State: Set the component state to either:
 - Online: The Archive Node is available to process object data for storage to the archival storage system.
 - Offline: The Archive Node is not available to process object data for storage to the archival storage system.
- Archive Store Disabled on Startup: When selected, the Archive Store component remains in the Read-only state when restarted. Used to persistently disable storage to the targeted the archival storage system. Useful when the targeted 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.

What you'll need

- 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 cannot 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

 Configuration: ARC (DC1-ARC1-98-165) - Store
Updated: 2015-05-07 12:38:07 PDT

Store State Offline

Archive Store Disabled on Startup ☐

Reset Store Failure Count ☐

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.

What you'll need

- 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**.

The screenshot shows the 'Configuration: ARC (DC1-ARC1-98-165) - Retrieve' page. The 'Retrieve State' dropdown is set to 'Online'. There are checkboxes for 'Reset Request Failure Count' and 'Reset Verification Failure Count', both of which are currently unchecked. An 'Apply Changes' button with a right-pointing arrow is located at the bottom right of the configuration area.

4. Modify the following settings, as necessary:
 - **Retrieve State:** Set the component state to either:
 - Online: The grid node is available to retrieve object data from the archival media device.
 - Offline: The grid node is not available to retrieve object data.
 - **Reset Request Failures Count:** Select the check box to reset the counter for request failures. This can be used to clear the ARRF (Request Failures) alarm.
 - **Reset Verification Failure Count:** Select the check box to reset the counter for verification failures on retrieved object data. This can be used to clear the ARRV (Verification Failures) alarm.
5. 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.

What you'll need

- 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 > Replication**.
3. Select **Configuration > Main**.

Overview Alarms Reports **Configuration**

Main Alarms

Configuration: ARC (DC1-ARC1-98-165) - Replication
Updated: 2015-05-07 12:21:53 PDT

Reset Inbound Replication Failure Count ☐

Reset Outbound Replication Failure Count ☐

Inbound Replication

Disable Inbound Replication ☐

Outbound Replication

Disable Outbound Replication ☐

Apply Changes

4. Modify the following settings, as necessary:

- **Reset Inbound Replication Failure Count:** Select to reset the counter for inbound replication failures. This can be used to clear the RIRF (Inbound Replications — Failed) alarm.
- **Reset Outbound Replication Failure Count:** Select to reset the counter for outbound replication failures. This can be used to clear the RORF (Outbound Replications — Failed) alarm.
- **Disable Inbound Replication:** Select to disable inbound replication as part of a maintenance or testing procedure. Leave cleared during normal operation.

When inbound replication is disabled, object data can be retrieved from the ARC service for replication to other locations in the StorageGRID system, but objects cannot be replicated to this ARC service from other system locations. The ARC service is read-only.

- **Disable Outbound Replication:** Select the check box to disable outbound replication (including content requests for HTTP retrievals) as part of a maintenance or testing procedure. Leave unchecked during normal operation.

When outbound replication is disabled, object data can be copied to this ARC service to satisfy ILM rules, but object data cannot be retrieved from the ARC service to be copied to other locations in the StorageGRID system. The ARC service is write-only.

5. 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 [Monitor and troubleshoot](#).

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 that can be overwritten with newer data) might be treated differently than application data (fixed content that must be retained indefinitely).

The Archive Node can be integrated with a new or an existing TSM server; it does not require a dedicated TSM server. TSM servers can be shared with other clients, provided that the TSM server is sized appropriately for the maximum expected load. TSM must be installed on a server or virtual machine separate from the Archive Node.

It is possible to configure more than one Archive Node to write to the same TSM server; however, this configuration is only recommended if the Archive Nodes write different sets of data to the TSM server. Configuring more than one Archive Node to write to the same TSM server is not recommended when each Archive Node writes copies of the same object data to the archive. In the latter scenario, both copies are subject to a single point of failure (the TSM server) for what are supposed to be independent, redundant copies of object data.

Archive Nodes do not make use of the Hierarchical Storage Management (HSM) component of TSM.

Configuration best practices

When you are sizing and configuring your TSM server there are best practices you

should apply to optimize it to work with the Archive Node.

When sizing and configuring the TSM server, you should consider the following factors:

- Because the Archive Node does not aggregate objects before saving them to the TSM server, the TSM database must be sized to hold references to all objects that will be written to the Archive Node.
- Archive Node software cannot tolerate the latency involved in writing objects directly to tape or other removable media. Therefore, the TSM server must be configured with a disk storage pool for the initial storage of data saved by the Archive Node whenever removable media are used.
- You must configure TSM retention policies to use event-based retention. The Archive Node does not support creation-based TSM retention policies. Use the following recommended settings of `retmin=0` and `retver=0` in the retention policy (which indicates that retention begins when the Archive Node triggers a retention event, and is retained for 0 days after that). However, these values for `retmin` and `retver` are optional.

The disk pool must be configured to migrate data to the tape pool (that is, the tape pool must be the `NXTSTGPOOL` of the disk pool). The tape pool must not be configured as a copy pool of the disk pool with simultaneous write to both pools (that is, the tape pool cannot be a `COPYSTGPOOL` for the disk pool). To create offline copies of the tapes containing Archive Node data, configure the TSM server with a second tape pool that is a copy pool of the tape pool used for Archive Node data.

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 cannot be co-hosted with a TSM server.

Configure the TSM server

This section includes sample instructions for preparing a TSM server following TSM best practices.

The following instructions guide you through the process of:

- Defining a disk storage pool, and a tape storage pool (if required) on the TSM server

- Defining a domain policy that uses the TSM management class for the data saved from the Archive Node, and registering a node to use this domain policy

These instructions are provided for your guidance only; they are not intended to replace TSM documentation, or to provide complete and comprehensive instructions suitable for all configurations. Deployment specific instructions should be provided by a TSM administrator who is familiar both with your detailed requirements, and with the complete set of TSM Server documentation.

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 a number of steps on your TSM server before you can create a tape storage pool. (Create a tape library and at least one drive in the tape library. Define a path from the server to the library and from the server to the drives, and then define a device class for the drives.) The details of these steps can vary depending upon the hardware configuration and storage requirements of the site. For more information, see the TSM documentation.

The following set of instructions illustrates the process. You should be aware that the requirements for your site could be different depending on the requirements of your deployment. For configuration details and for instructions, see the TSM documentation.



You must log onto the server with administrative privileges and use the `dsmadm` tool to execute the following commands.

Steps

1. Create a tape library.

```
define library tapelibrary libtype=scsi
```

Where *tapelibrary* is an arbitrary name chosen for the tape library, and the value of *libtype* can vary depending upon the type of tape library.

2. Define a path from the server to the tape library.

```
define path servername tapelibrary srctype=server desttype=library device=lib-devicename
```

- ° *servername* is the name of the TSM server
- ° *tapelibrary* is the tape library name you defined
- ° *lib-devicename* is the device name for the tape library

3. Define a drive for the library.

```
define drive tapelibrary drivename
```

- *drivename* is the name you want to specify for the drive
- *tapelibrary* is the tape library name you defined

You might want to configure an additional drive or drives, depending upon your hardware configuration. (For example, if the TSM server is connected to a Fibre Channel switch that has two inputs from a tape library, you might want to define a drive for each input.)

4. Define a path from the server to the drive you defined.

```
define path servername drivename srctype=server desttype=drive
library=tapelibrary device=drive-dname
```

- *drive-dname* is the device name for the drive
- *tapelibrary* is the tape library name you defined

Repeat for each drive that you have defined for the tape library, using a separate *drivename* and *drive-dname* for each drive.

5. Define a device class for the drives.

```
define devclass DeviceClassName devtype=lto library=tapelibrary
format=tapetype
```

- *DeviceClassName* is the name of the device class
- *lto* is the type of drive connected to the server
- *tapelibrary* is the tape library name you defined
- *tapetype* is the tape type; for example, ultrium3

6. Add tape volumes to the inventory for the library.

```
checkin libvolume tapelibrary
```

tapelibrary is the tape library name you defined.

7. Create the primary tape storage pool.

```
define stgpool SGWSTapePool DeviceClassName description=description
collocate=filespace maxxscratch=XX
```

- *SGWSTapePool* is the name of the Archive Node's tape storage pool. You can select any name for the tape storage pool (as long as the name uses the syntax conventions expected by the TSM server).
- *DeviceClassName* is the name of the device class name for the tape library.
- *description* is a description of the storage pool that can be displayed on the TSM server using the query stgpool command. For example: "Tape storage pool for the Archive Node."
- *collocate=filespace* specifies that the TSM server should write objects from the same file space into a single tape.
- *xx* is one of the following:
 - The number of empty tapes in the tape library (in the case that the Archive Node is the only

application using the library).

- The number of tapes allocated for use by the StorageGRID system (in instances where the tape library is shared).

8. On a TSM server, create a disk storage pool. At the TSM server's administrative console, enter

```
define stgpool SGWSDiskPool disk description=description  
maxsize=maximum_file_size nextstgpool=SGWSTapePool highmig=percent_high  
lowmig=percent_low
```

- *SGWSDiskPool* is the name of the Archive Node's disk pool. You can select any name for the disk storage pool (as long as the name uses the syntax conventions expected by the TSM).
- *description* is a description of the storage pool that can be displayed on the TSM server using the query stgpool command. For example, "Disk storage pool for the Archive Node."
- *maximum_file_size* forces objects larger than this size to be written directly to tape, rather than being cached in the disk pool. It is recommended to set *maximum_file_size* to 10 GB.
- *nextstgpool=SGWSTapePool* refers the disk storage pool to the tape storage pool defined for the Archive Node.
- *percent_high* sets the value at which the disk pool begins to migrate its contents to the tape pool. It is recommended to set *percent_high* to 0 so that data migration begins immediately
- *percent_low* sets the value at which migration to the tape pool stops. It is recommended to set *percent_low* to 0 to clear out the disk pool.

9. On a TSM server, create a disk volume (or volumes) and assign it to the disk pool.

```
define volume SGWSDiskPool volume_name formatsize=size
```

- *SGWSDiskPool* is the disk pool name.
- *volume_name* is the full path to the location of the volume (for example, */var/local/arc/stage6.dsm*) on the TSM server where it writes the contents of the disk pool in preparation for transfer to tape.
- *size* is the size, in MB, of the disk volume.

For example, to create a single disk volume such that the contents of a disk pool fill a single tape, set the value of *size* to 200000 when the tape volume has a capacity of 200 GB.

However, it might be desirable to create multiple disk volumes of a smaller size, as the TSM server can write to each volume in the disk pool. For example, if the tape size is 250 GB, create 25 disk volumes with a size of 10 GB (10000) each.

The TSM server preallocates space in the directory for the disk volume. This can take some time to complete (more than three hours for a 200 GB disk volume).

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 are not supported by the Archive Node.

The value of MAXSESSIONS set for the TSM server controls the maximum number of sessions that can be opened to the TSM server by all client applications. The value of MAXSESSIONS specified on the TSM must be at least as large as the value specified for **ARC > Target > Configuration > Main > Number of Sessions** in the Grid Manager for the Archive Node. The Archive Node concurrently creates at most one session per mount point plus a small number (< 5) of additional sessions.

The TSM node assigned to the Archive Node uses a custom domain policy `tsm-domain`. The `tsm-domain` domain policy is a modified version of the “standard” domain policy, configured to write to tape and with the archive destination set to be the StorageGRID system’s storage pool (`SGWSDiskPool`).



You must log in to the TSM server with administrative privileges and use the `dsmadm` tool to create and activate the domain policy.

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 are not using an existing management class, enter one of the following:

```
define policyset tsm-domain standard
```

```
define mgmtclass tsm-domain standard default
```

default is the default management class for the deployment.

3. Create a copygroup to the appropriate storage pool. Enter (on one line):

```
define copygroup tsm-domain standard default type=archive
destination=SGWSDiskPool retinit=event retmin=0 retver=0
```

default is the default Management Class for the Archive Node. The values of `retinit`, `retmin`, and `retver` have been chosen to reflect the retention behavior currently used by the Archive Node



Do not set `retinit` to `retinit=create`. Setting `retinit=create` blocks the Archive Node from deleting content since retention events are used to remove content from the TSM server.

4. Assign the management class to be the default.

```
assign defmgmtclass tsm-domain standard default
```

5. Set the new policy set as active.

```
activate policyset tsm-domain standard
```

Ignore the “no backup copy group” warning that appears when you enter the activate command.

6. Register a node to use the new policy set on the TSM server. On the TSM server, enter (on one line):

```
register node arc-user arc-password passexp=0 domain=tsm-domain  
MAXNUMMP=number-of-sessions
```

`arc-user` and `arc-password` are same client node name and password as you define on the Archive Node, and the value of `MAXNUMMP` is set to the number of tape drives reserved for Archive Node store sessions.



By default, registering a node creates an administrative user ID with client owner authority, with the password defined for the node.

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.

The following section is a guide to understanding and planning a migration of large amounts of data into the StorageGRID system. It is not a general guide to data migration, and it does not include detailed steps for performing a migration. Follow the guidelines and instructions in this section to ensure that data is migrated efficiently into the StorageGRID system without interfering with day-to-day operations, and that the migrated data is handled appropriately by the StorageGRID system.

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 [Monitor and troubleshoot](#).

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.

Impact of migration on operations

A StorageGRID system is designed to provide efficient operation for object storage and retrieval, and to provide excellent protection against data loss through the seamless creation of redundant copies of object data and metadata.

However, data migration must be carefully managed according to the instructions in this chapter to avoid having an impact on day-to-day system operations, or, in extreme cases, placing data at risk of loss in case of a failure in the StorageGRID system.

Migration of large quantities of data places additional load on the system. When the StorageGRID system is heavily loaded, it responds more slowly to requests to store and retrieve objects. This can interfere with store and retrieve requests which are integral to day-to-day operations. Migration can also cause other operational issues. For example, when a Storage Node is nearing capacity, the heavy intermittent load due to batch ingest can cause the Storage Node to cycle between read-only and read-write, generating notifications.

If the heavy loading persists, queues can develop for various operations that the StorageGRID system must perform to ensure full redundancy of object data and metadata.

Data migration must be carefully managed according to the guidelines in this document to ensure safe and efficient operation of the StorageGRID system during migration. When migrating data, ingest objects in batches or continuously throttle ingest. Then, continuously monitor the StorageGRID system to ensure that various attribute values are not exceeded.

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, do not schedule data migration during periods of high activity. However, if it is not practical to completely avoid the high activity period, it is safe to proceed as long as you closely monitor the relevant attributes and take action if they exceed acceptable values.

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 evaluation	<ol style="list-style-type: none">1. Select SUPPORT > Tools > Grid topology.2. Select deployment > Overview > Main.3. In the ILM Activity section, monitor the number of objects shown for the following attributes:<ul style="list-style-type: none">◦ 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).4. If the number of objects shown for either of these attributes exceeds 100,000, throttle the ingest rate of objects to reduce the load on the StorageGRID system.
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: Overview

You manage the objects in a StorageGRID system by configuring information lifecycle management (ILM) rules and policies. The ILM rules and policies instruct StorageGRID how to create and distribute copies of object data and how to manage those copies over time.

About these instructions

Designing and implementing ILM rules and the ILM policy requires careful planning. You must understand your operational requirements, the topology of your StorageGRID system, your object protection needs, and the available storage types. Then, you must determine how you want different types of objects to be copied, distributed, and stored.

Use these instructions to:

- Learn about StorageGRID ILM, including how ILM operates throughout an object's life and what ILM policies and rules are.
- Learn how to configure storage pools, Erasure Coding profiles, and ILM rules.
- Learn how to create 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 are not deleted or overwritten for a specified amount of time.

Learn more

To learn more, review these videos:

- [Video: StorageGRID ILM Rules: Getting Started](#)



- [Video: StorageGRID ILM Policies](#)



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 the instructions for [monitoring and troubleshooting StorageGRID](#).
- **Object deletion:** Management of an object ends when all copies are removed from the StorageGRID system. Objects can be removed as a result of a delete request by a client, or as a result of deletion by ILM or deletion caused by the expiration of an S3 bucket lifecycle.



Objects in a bucket that has S3 Object Lock enabled cannot be deleted if they are under a legal hold or if a retain-until-date has been specified but not yet met.

The diagram summarizes how ILM operates throughout an object's lifecycle.




How objects are ingested

Data-protection options for ingest

When you create an ILM rule, you specify one of three options for protecting objects at ingest: Dual commit, Balanced, or Strict. Depending on your choice, StorageGRID makes interim copies and queues the objects for ILM evaluation later, or it uses synchronous placement and immediately makes copies to meet ILM requirements.

Flowchart of three 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.

When to use the Dual commit option

Use the Dual commit option in either of these cases:

- You are using multi-site ILM rules and client ingest latency is your primary consideration. When using Dual commit, you must ensure your grid can perform the additional work of creating and removing the dual-

commit copies if they do not satisfy ILM. Specifically:

- The load on the grid must be low enough to prevent an ILM backlog.
- The grid must have excess hardware resources (IOPS, CPU, memory, network bandwidth, and so on).
- You are using multi-site ILM rules and the WAN connection between the sites usually has high latency or limited bandwidth. In this scenario, using the Dual commit option can help prevent client timeouts. Before choosing the Dual commit option, you should test the client application with realistic workloads.

Strict

When you select the Strict option, StorageGRID uses synchronous placement on ingest and immediately makes all object copies specified in the rule's placement instructions. Ingest fails if StorageGRID cannot create all copies, for example, because a required storage location is temporarily unavailable. The client must retry the operation.

When to use the Strict option

Use the Strict option if you have an operational or regulatory requirement to immediately store objects only in the locations outlined in the ILM rule. For example, to satisfy a regulatory requirement, you might need to use the Strict option and a Location Constraint advanced filter to guarantee that objects are never stored at certain data center.

Example 5: ILM rules and policy for Strict ingest behavior

Balanced

When you select the Balanced option, StorageGRID also uses synchronous placement on ingest and immediately makes all copies specified in the rule's placement instructions. In contrast with the Strict option, if StorageGRID cannot immediately make all copies, it uses Dual commit instead.

When to use the Balanced option

Use the Balanced option to achieve the best combination of data protection, grid performance, and ingest success. Balanced is the default option in the ILM rule wizard.

Advantages, disadvantages, and limitations of the data-protection options

Understanding the advantages and disadvantages of each of the three options for protecting data at ingest (Balanced, Strict, or Dual commit) can help you decide which one to select for an ILM rule.

Advantages of the Balanced and Strict options

When compared to Dual commit, which creates interim copies during ingest, the two synchronous placement options can provide the following advantages:

- **Better data security:** Object data is immediately protected as specified in the ILM rule's placement instructions, which can be configured to protect against a wide variety of failure conditions, including the failure of more than one storage location. Dual commit can only protect against the loss of a single local copy.
- **More efficient grid operation:** Each object is processed only once, as it is ingested. Because the StorageGRID system does not need to track or delete interim copies, there is less processing load and less database space is consumed.

- **(Balanced) Recommended:** The Balanced option provides optimal ILM efficiency. Using the Balanced option is recommended unless Strict ingest behavior is required or the grid meets all of the criteria for using for Dual commit.
- **(Strict) Certainty about object locations:** The Strict option guarantees that objects are immediately stored according to the placement instructions in the ILM rule.

Disadvantages of the Balanced and Strict options

When compared to Dual commit, the Balanced and Strict options have some disadvantages:

- **Longer client ingests:** Client ingest latencies might be longer. When you use the Balanced and Strict options, an “ingest successful” message is not returned to the client until all erasure-coded fragments or replicated copies are created and stored. However, object data will most likely reach its final placement much faster.
- **(Strict) Higher rates of ingest failure:** With the Strict option, ingest fails whenever StorageGRID cannot immediately make all copies specified in the ILM rule. You might see high rates of ingest failure if a required storage location is temporarily offline or if network issues cause delays in copying objects between sites.
- **(Strict) S3 multipart upload placements might not be as expected in some circumstances:** With Strict, you expect objects either to be placed as described by the ILM rule or for ingest to fail. However, with an S3 multipart upload, ILM is evaluated for each part of the object as it ingested, and for the object as a whole when the multipart upload completes. In the following circumstances this might result in placements that are different than you expect:
 - **If ILM changes while an S3 multipart upload is in progress:** Because each part is placed according to the rule that is active when the part is ingested, some parts of the object might not meet current ILM requirements when the multipart upload completes. In these cases, ingest of the object does not fail. Instead, any part that is not placed correctly is queued for ILM re-evaluation, and is moved to the correct location later.
 - **When ILM rules filter on size:** When evaluating ILM for a part, StorageGRID filters on the size of the part, not the size of the object. This means that parts of an object can be stored in locations that do not meet ILM requirements for the object as a whole. For example, if a rule specifies that all objects 10 GB or larger are stored at DC1 while all smaller objects are stored at DC2, at ingest each 1 GB part of a 10-part multipart upload is stored at DC2. When ILM is evaluated for the object, all parts of the object are moved to DC1.
- **(Strict) Ingest does not fail when object tags or metadata are updated and newly required placements cannot be made:** With Strict, you expect objects either to be placed as described by the ILM rule or for ingest to fail. However, when you update metadata or tags for an object that is already stored in the grid, the object is not re-ingested. This means that any changes to object placement that are triggered by the update are not made immediately. Placement changes are made when ILM is re-evaluated by normal background ILM processes. If required placement changes cannot be made (for example, because a newly required location is unavailable), the updated object retains its current placement until the placement changes are possible.

Limitations on object placements with the Balanced or Strict options

The Balanced or Strict options cannot be used for ILM rules that have any of these placement instructions:

- Placement in a Cloud Storage Pool at day 0.
- Placement in an Archive Node at day 0.
- Placements in a Cloud Storage Pool or an Archive Node when the rule has a User Defined Creation Time as its Reference Time.

These restrictions exist because StorageGRID cannot synchronously make copies to a Cloud Storage Pool or an Archive Node, and a User Defined Creation Time could resolve to the present.

How ILM rules and consistency controls interact to affect data protection

Both your ILM rule and your choice of consistency control affect how objects are protected. These settings can interact.

For example, the ingest behavior selected for an ILM rule affects the initial placement of object copies, while the consistency control used when an object is stored affects the initial placement of object metadata. Because StorageGRID requires access to both an object's metadata and its data to fulfill client requests, selecting matching levels of protection for the consistency level and ingest behavior can provide better initial data protection and more predictable system responses.

Here is a brief summary of the consistency controls that are available in StorageGRID:

- **all**: All nodes receive object metadata immediately or the request will fail.
- **strong-global**: Object metadata is immediately distributed to all sites. Guarantees read-after-write consistency for all client requests across all sites.
- **strong-site**: Object metadata is immediately distributed to other nodes at the site. Guarantees read-after-write consistency for all client requests within a site.
- **read-after-new-write**: Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees.
- **available** (eventual consistency for HEAD operations): Behaves the same as the “read-after-new-write” consistency level, but only provides eventual consistency for HEAD operations.



Before selecting a consistency level, read the full description of consistency controls in the instructions for [S3](#) or [Swift](#) client applications. You should understand the benefits and limitations before changing the default value.

Example of how the consistency control and ILM rule can interact

Suppose you have a two-site grid with the following ILM rule and the following consistency level setting:

- **ILM rule**: Create two object copies, one at the local site and one at a remote site. The Strict ingest behavior is selected.
- **Consistency level**: “strong-global” (Object metadata is immediately distributed to all sites.)

When a client stores an object to the grid, StorageGRID makes both object copies and distributes metadata to both sites before returning success to the client.

The object is fully protected against loss at the time of the ingest successful message. For example, if the local site is lost shortly after ingest, copies of both the object data and the object metadata still exist at the remote site. The object is fully retrievable.

If you instead used the same ILM rule and the “strong-site” consistency level, the client might receive a success message after object data is replicated to the remote site but before object metadata is distributed there. In this case, the level of protection of object metadata does not match the level of protection for object data. If the local site is lost shortly after ingest, object metadata is lost. The object cannot be retrieved.

The inter-relationship between consistency levels and ILM rules can be complex. Contact NetApp if you require assistance.

Related information

- [Example 5: ILM rules and policy for Strict ingest behavior](#)


How objects are stored (replication or erasure coding)

What replication is

Replication is one of two methods used by StorageGRID to store object data. When objects match an ILM rule that uses replication, the system creates exact copies of object data and stores the copies on Storage Nodes or Archive Nodes.

When you configure an ILM rule to create replicated copies, you specify how many copies should be created, where those copies should be placed, and how long the copies should be stored at each location.

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



When StorageGRID matches objects to this rule, it creates two copies of the object, placing each copy on a different Storage Node in the storage pool. The two copies might be placed on any two of the three available Storage Nodes. In this case, the rule placed object copies on Storage Nodes 2 and 3. Because there are two copies, the object can be retrieved if any of the nodes in the storage pool fails.



StorageGRID can store only one replicated copy of an object on any given Storage Node. If your grid includes three Storage Nodes and you create a 4-copy ILM rule, only three copies will be made—one copy for each Storage Node. The **ILM placement unachievable** alert is triggered to indicate that the ILM rule could not be completely applied.

Related information

- [What a storage pool is](#)

- [Use multiple storage pools for cross-site replication](#)


Why you should not use single-copy replication

When creating an ILM rule to create replicated copies, you should always specify at least two copies for any time period in the placement instructions.




Do not use an ILM rule that creates only one replicated copy for any time period. If only one replicated copy of an object exists, that object is lost if a Storage Node fails or has a significant error. You also temporarily lose access to the object during maintenance procedures such as upgrades.

In the following example, the Make 1 Copy ILM rule specifies that one replicated copy of an object be placed in a storage pool that contains three Storage Nodes. When an object is ingested that matches this rule, StorageGRID places a single copy on only one Storage Node.



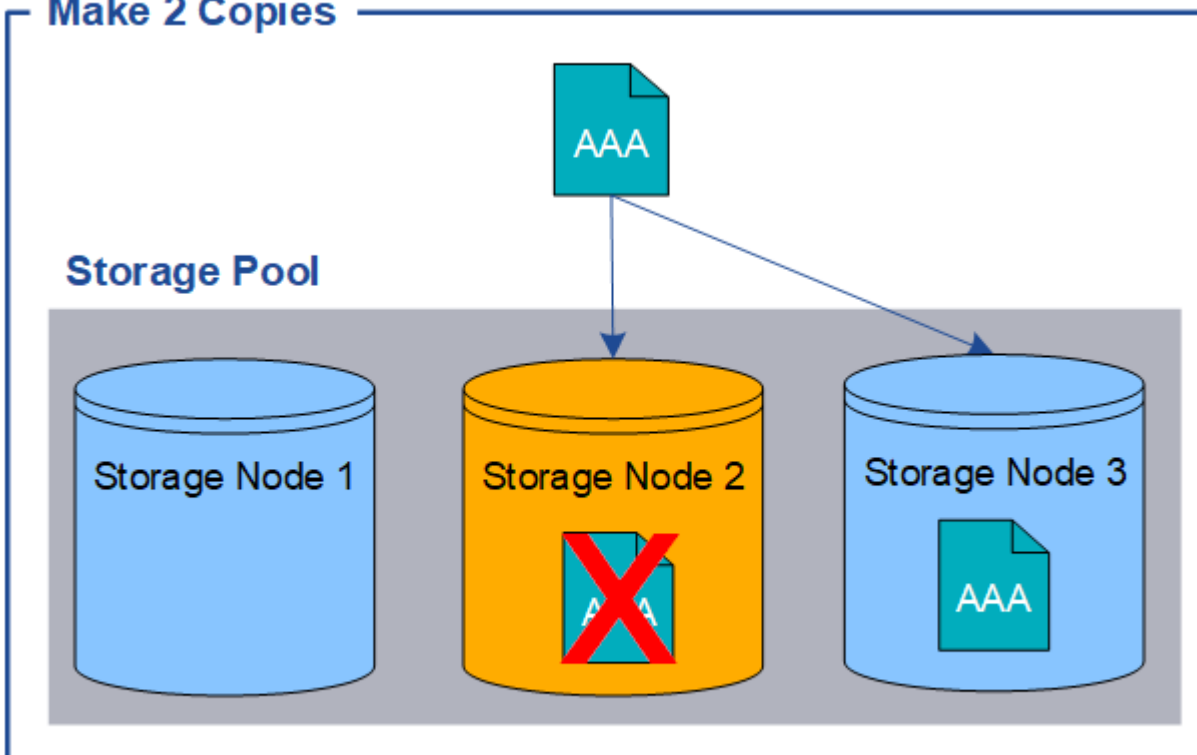
When an ILM rule creates only one replicated copy of an object, the object becomes inaccessible when the Storage Node is unavailable. In this example, you will temporarily lose access to object AAA whenever Storage Node 2 is offline, such as during an upgrade or other maintenance procedure. You will lose object AAA entirely if Storage Node 2 fails.

Make 1 Copy



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.


Make 2 Copies




What erasure coding is

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



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



The 4+2 erasure-coding scheme requires a minimum of nine Storage Nodes, with three Storage Nodes at each of three different sites. An object can be retrieved as long as any four of the six fragments (data or parity) remain available. Up to two fragments can be lost without loss of the object data. If an entire data center site is lost, the object can still be retrieved or repaired, as long as all of the other fragments remain accessible.



If more than two Storage Nodes are lost, the object is not retrievable.



Related information

- [What a storage pool is](#)
- [What erasure-coding schemes are](#)
- [Create an Erasure Coding profile](#)

What erasure-coding schemes are

When you configure the Erasure Coding profile for an ILM rule, you select an available erasure-coding scheme based on how many Storage Nodes and sites make up the storage pool you plan to use. Erasure-coding schemes control how many data fragments and how many parity fragments are created for each object.

The StorageGRID system uses the Reed-Solomon erasure-coding algorithm. The algorithm slices an object into k data fragments and computes m parity fragments. The $k + m = n$ fragments are spread across n Storage Nodes to provide data protection. An object can sustain up to m lost or corrupt fragments. k fragments are needed to retrieve or repair an object.

When configuring an Erasure Coding profile, use the following guidelines for storage pools:

- The storage pool must include three or more sites, or exactly one site.



You cannot configure an Erasure Coding profile if the storage pool includes two sites.

- [Erasure-coding schemes for storage pools containing three or more sites](#)
- [Erasure-coding schemes for one-site storage pools](#)
- Do not use the default storage pool, All Storage Nodes, or a storage pool that includes the default site, All Sites.
- The storage pool should include at least $k+m+1$ Storage Nodes.

The minimum number of Storage Nodes required is $k+m$. However, having at least one additional Storage Node can help prevent ingest failures or ILM backlogs if a required Storage Node is temporarily unavailable.

The storage overhead of an erasure-coding scheme is calculated by dividing the number of parity fragments (m) by the number of data fragments (k). You can use the storage overhead to calculate how much disk space each erasure-coded object requires:

$$\text{disk space} = \text{object size} + (\text{object size} * \text{storage overhead})$$

For example, if you store a 10 MB object using the 4+2 scheme (which has 50% storage overhead), the object consumes 15 MB of grid storage. If you store the same 10 MB object using the 6+2 scheme (which has 33% storage overhead), the object consumes approximately 13.3 MB.

Select the erasure-coding scheme with the lowest total value of $k+m$ that meets your needs. erasure-coding schemes with a lower number of fragments are overall more computationally efficient, as fewer fragments are created and distributed (or retrieved) per object, can show better performance due to the larger fragment size, and can require fewer nodes be added in an expansion when more storage is required. (See the instructions for expanding StorageGRID for information on planning a storage expansion.)

Erasure-coding schemes for storage pools containing three or more sites

The following table describes the erasure-coding schemes currently supported by StorageGRID for storage pools that include three or more sites. All of these schemes provide site loss protection. One site can be lost, and the object will still be accessible.

For erasure-coding schemes that provide site loss protection, the recommended number of Storage Nodes in the storage pool exceeds $k+m+1$ because each site requires a minimum of three Storage Nodes.

Erasure-coding scheme ($k+m$)	Minimum number of deployed sites	Recommended number of Storage Nodes at each site	Total recommended number of Storage Nodes	Site loss protection?	Storage overhead
4+2	3	3	9	Yes	50%
6+2	4	3	12	Yes	33%
8+2	5	3	15	Yes	25%
6+3	3	4	12	Yes	50%
9+3	4	4	16	Yes	33%
2+1	3	3	9	Yes	50%
4+1	5	3	15	Yes	25%
6+1	7	3	21	Yes	17%
7+5	3	5	15	Yes	71%



StorageGRID requires a minimum of three Storage Nodes per site. To use the 7+5 scheme, each site requires a minimum of four Storage Nodes. Using five Storage Nodes per site is recommended.

When selecting an erasure-coding scheme that provides site protection, balance the relative importance of the following factors:

- **Number of fragments:** Performance and expansion flexibility are generally better when the total number of fragments is lower.
- **Fault tolerance:** Fault tolerance is increased by having more parity segments (that is, when m has a higher value.)
- **Network traffic:** When recovering from failures, using a scheme with more fragments (that is, a higher total for $k+m$) creates more network traffic.
- **Storage overhead:** Schemes with higher overhead require more storage space per object.

For example, when deciding between a 4+2 scheme and 6+3 scheme (which both have 50% storage overhead), select the 6+3 scheme if additional fault tolerance is required. Select the 4+2 scheme if network resources are constrained. If all other factors are equal, select 4+2 because it has a lower total number of fragments.



If you are unsure of which scheme to use, select 4+2 or 6+3, or contact technical support.

Erasure-coding schemes for one-site storage pools

A one-site storage pool supports all of the erasure-coding schemes defined for three or more sites, provided that the site has enough Storage Nodes.

The minimum number of Storage Nodes required is $k+m$, but a storage pool with $k+m+1$ Storage Nodes is recommended. For example, the 2+1 erasure-coding scheme requires a storage pool with a minimum of three Storage Nodes, but four Storage Nodes is recommended.

Erasure-coding scheme ($k+m$)	Minimum number of Storage Nodes	Recommended number of Storage Nodes	Storage overhead
4+2	6	7	50%
6+2	8	9	33%
8+2	10	11	25%
6+3	9	10	50%
9+3	12	13	33%
2+1	3	4	50%
4+1	5	6	25%
6+1	7	8	17%
7+5	12	13	71%

Related information

[Expand your grid](#)

Advantages, disadvantages, and requirements for erasure coding

Before deciding whether to use replication or erasure coding to protect object data from loss, you should understand the advantages, disadvantages, and the requirements for erasure coding.

Advantages of erasure coding

When compared to replication, erasure coding offers improved reliability, availability, and storage efficiency.

- **Reliability:** Reliability is gauged in terms of fault tolerance—that is, the number of simultaneous failures that can be sustained without loss of data. With replication, multiple identical copies are stored on different nodes and across sites. With erasure coding, an object is encoded into data and parity fragments and distributed across many nodes and sites. This dispersal provides both site and node failure protection. When compared to replication, erasure coding provides improved reliability at comparable storage costs.
- **Availability:** Availability can be defined as the ability to retrieve objects if Storage Nodes fail or become

inaccessible. When compared to replication, erasure coding provides increased availability at comparable storage costs.

- **Storage efficiency:** For similar levels of availability and reliability, objects protected through erasure coding consume less disk space than the same objects would if protected through replication. For example, a 10 MB object that is replicated to two sites consumes 20 MB of disk space (two copies), while an object that is erasure coded across three sites with a 6+3 erasure-coding scheme only consumes 15 MB of disk space.



Disk space for erasure-coded objects is calculated as the object size plus the storage overhead. The storage overhead percentage is the number of parity fragments divided by the number of data fragments.

Disadvantages of erasure coding

When compared to replication, erasure coding has the following disadvantages:

- An increased number of Storage Nodes and sites is required. For example, if you use an erasure-coding scheme of 6+3, you must have at least three Storage Nodes at three different sites. In contrast, if you simply replicate object data, you require only one Storage Node for each copy.
- Increased cost and complexity of storage expansions. To expand a deployment that uses replication, you simply add storage capacity in every location where object copies are made. To expand a deployment that uses erasure coding, you must consider both the erasure-coding scheme in use and how full existing Storage Nodes are. For example, if you wait until existing nodes are 100% full, you must add at least $k+m$ Storage Nodes, but if you expand when existing nodes are 70% full, you can add two nodes per site and still maximize usable storage capacity. For more information, see [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. Do not 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 cannot be deleted by any method.
 - Before an object version's retain-until-date is reached, that version cannot be deleted by any method.
 - Objects in buckets with S3 Object Lock enabled are retained by ILM "forever." However, after its retain-until-date is reached, an object version can be deleted by a client request or the expiration of the bucket lifecycle. 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

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

Examples for object retention

To better understand the interactions between S3 Object Lock, bucket lifecycle settings, client delete requests, and ILM, consider the following examples.

Example 1: S3 bucket lifecycle keeps objects longer than ILM

ILM

Store two copies for 1 year (365 days)

Bucket lifecycle

Expire objects in 2 years (730 days)

Result

StorageGRID stores the object for 730 days. StorageGRID uses the bucket lifecycle settings to determine whether to delete or retain an object.



If the bucket lifecycle specifies that objects should be kept longer than specified by ILM, StorageGRID continues to use the ILM placement instructions when determining the number and type of copies to store. In this example, two copies of the object will continue to be stored in StorageGRID from days 366 to 730.

Example 2: S3 bucket lifecycle expires objects before ILM

ILM

Store two copies for 2 years (730 days)

Bucket lifecycle

Expire objects in 1 year (365 days)

Result

StorageGRID deletes both copies of the object after day 365.

Example 3: Client delete overrides bucket lifecycle and ILM

ILM

Store two copies on Storage Nodes “forever”

Bucket lifecycle

Expire objects in 2 years (730 days)

Client delete request

Issued on day 400

Result

StorageGRID deletes both copies of the object on day 400 in response to the client delete request.

Example 4: S3 Object Lock overrides client delete request

S3 Object Lock

Retain-until-date for an object version is 2026-03-31. A legal hold is not in effect.

Compliant ILM rule

Store two copies on Storage Nodes “forever.”

Client delete request

Issued on 2024-03-31.

Result

StorageGRID will not delete the object version because the retain-until-date is still 2 years away.

How objects are deleted

StorageGRID can delete objects either in direct response to a client request or automatically as a result of the expiration of an S3 bucket lifecycle or the requirements of the ILM policy. Understanding the different ways that objects can be deleted and how StorageGRID handles delete requests can help you manage objects more effectively.

StorageGRID can use one of two methods to delete objects:

- Synchronous deletion: When StorageGRID receives a client delete request, all object copies are removed immediately. The client is informed that deletion was successful after the copies have been removed.
- Objects are queued for deletion: When StorageGRID receives a delete request, the object is queued for deletion and the client is informed immediately that deletion was successful. Object copies are removed later by background ILM processing.

When deleting objects, StorageGRID uses the method that optimizes delete performance, minimizes potential delete backlogs, and frees space most quickly.

The table summarizes when StorageGRID uses each method.

Method of performing deletion	When used
Objects are queued for deletion	<p>When any of the following conditions are true:</p> <ul style="list-style-type: none">• Automatic object deletion has been triggered by one of the following events:<ul style="list-style-type: none">◦ 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.• An S3 or Swift client requests deletion and one or more of these conditions is true:<ul style="list-style-type: none">◦ Copies cannot be deleted within 30 seconds because, for example, an object location is temporarily unavailable.◦ Background deletion queues are idle. <p>Note: Objects in a bucket that has S3 Object Lock enabled cannot be deleted if they are under a legal hold or if a retain-until-date has been specified but not yet met.</p>

Method of performing deletion	When used
Objects are removed immediately (synchronous deletion)	<p>When an S3 or Swift client makes a delete request and all of the following conditions are met:</p> <ul style="list-style-type: none"> • 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 a number of objects to the delete queue. It then switches to performing synchronous deletion. Making sure that the background deletion queue has objects to process allows StorageGRID to process deletes more efficiently, especially for low concurrency clients, while helping to prevent client delete backlogs.

How long does it take 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 slow after a certain number of objects have been deleted. This change occurs when StorageGRID shifts from queuing objects for deletion to performing synchronous deletion. The apparent reduction in the deletion rate does not mean that object copies are being removed more slowly. On the contrary, it indicates that on average, space is now being freed more quickly.

If you are deleting large numbers of objects and your priority is to free space quickly, consider using a client request to delete objects rather than deleting them using ILM or other methods. In general, space is freed more quickly when deletion is performed by clients because StorageGRID can use synchronous deletion.

You should be aware that the amount of time required to free space after an object is deleted depends on a number of factors:

- Whether object copies are synchronously removed or are queued for removal later (for client delete requests).
- Other factors such as the number of objects in the grid or the availability of grid resources when object copies are queued for removal (for both client deletes and other methods).

How S3 versioned objects are deleted

When versioning is enabled for an S3 bucket, StorageGRID follows Amazon S3 behavior when responding to delete requests, whether those requests come from an S3 client, the expiration of an S3 bucket lifecycle, or the requirements of the ILM policy.

When objects are versioned, object delete requests do not delete the current version of the object and do not free space. Instead, an object delete request simply creates a delete marker as the current version of the object, which makes the previous version of the object “noncurrent.”

Even though the object has not been removed, StorageGRID behaves as though the current version of the object is no longer available. Requests to that object return 404 Not Found. However, because noncurrent object data has not been removed, requests that specify a noncurrent version of the object can succeed.

To free space when deleting versioned objects, you must do one of the following:

- **S3 client request:** Specify the object version number in the S3 DELETE Object request (`DELETE /object?versionId=ID`). Keep in mind that this request only removes object copies for the specified version (the other versions are still taking up space).
- **Bucket lifecycle:** Use the `NoncurrentVersionExpiration` action in the bucket lifecycle configuration. When the number of `NoncurrentDays` specified is met, StorageGRID permanently removes all copies of noncurrent object versions. These object versions cannot be recovered.
- **ILM:** Add two ILM rules to your ILM policy. Use **Noncurrent Time** as the Reference Time in the first rule to match the noncurrent versions of the object. Use **Ingest Time** in the second rule to match the current version. The **Noncurrent Time** rule must appear in the policy above the **Ingest Time** rule.

Related information

- [Use S3](#)
- [Example 4: ILM rules and policy for S3 versioned objects](#)

What an ILM policy is

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

How does an ILM policy evaluate objects?

The active ILM policy for your StorageGRID system controls the placement, duration, and data protection of all objects.

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

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

Example ILM policy

This example ILM policy uses three ILM rules.

Configure ILM Policy

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






Name Example ILM policy

Reason for change New policy

Rules

1. Select the rules you want to add to the policy.
2. Determine the order in which the rules will be evaluated by dragging and dropping the rows. The default rule will be automatically placed at the end of the policy and cannot be moved.

+ Select Rules

	Default	Rule Name	Tenant Account	Actions
+		Rule 1: 3 replicated copies for Tenant A 	Tenant A (58889986524346589742)	
+		Rule 2: Erasure coding for objects greater than 1 MB 	—	
	✓	Rule 3: 2 copies 2 data centers (default) 	—	


Cancel

Save

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

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

Rule 3 is the last and default rule in the policy, and it does not use filters. Rule 3 makes two replicated copies of all objects not matched by Rule 1 or Rule 2 (objects not belonging to Tenant A that are 1 MB or smaller).




What are proposed, active, and historical policies?

Every StorageGRID system must have one active ILM policy. A StorageGRID system might also have one proposed ILM policy and any number of historical policies.

When you first create an ILM policy, you create a proposed policy by selecting one or more ILM rules and arranging them in a specific order. After you have simulated the proposed policy to confirm its behavior, you activate it to create the active policy.

When you activate a new ILM policy, StorageGRID uses that policy to manage all objects, including existing objects and newly ingested objects. Existing objects might be moved to new locations when the ILM rules in the new policy are implemented.

Activating the proposed policy causes the previously active policy to become a historical policy. Historical ILM policies cannot be deleted.



Related information

[Create an ILM policy](#)

What an ILM rule is

To manage objects, you create a set of information lifecycle management (ILM) rules and organize them into an ILM policy. Every object ingested into the system is evaluated against the active policy. When a rule in the policy matches an object's metadata, the instructions in the rule determine what actions StorageGRID takes to copy and store that object.

ILM rules define:

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

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

Elements of an ILM rule

An ILM rule has three elements:

- **Filtering criteria:** A rule's basic and advanced filters define which objects the rule applies to. If an object matches all filters, StorageGRID applies the rule and creates the object copies specified in the rule's placement instructions.
- **Placement instructions:** A rule's placement instructions define the number, type, and location of object copies. Each rule can include a sequence of placement instructions to change the number, type, and location of object copies over time. When the time period for one placement expires, the instructions in the next placement are automatically applied by the next ILM evaluation.

- **Ingest behavior:** A rule's ingest behavior defines what happens when an S3 or Swift client saves an object to the grid. Ingest behavior controls whether object copies are immediately placed according to the instructions in the rule, or if interim copies are made and the placement instructions are applied later.

What ILM rule filtering is

When you create an ILM rule, you specify filters to identify which objects the rule applies to.

In the simplest case, a rule might not use any filters. Any rule that does not use filters applies to all objects, so it must be the last (default) rule in an ILM policy. The default rule provides storage instructions for objects that do not match the filters in another rule.

Basic filters allow you to apply different rules to large, distinct groups of objects. The basic filters on the Define Basics page of the Create ILM Rule wizard allow you to apply a rule to specific tenant accounts, specific S3 buckets or Swift containers, or both.

Create ILM Rule Step 1 of 3: Define Basics

Name

Description

Tenant Accounts (optional)

Select tenant accounts or enter tenant IDs

Bucket Name

matches all

Value

[Advanced filtering...](#) (0 defined)

Cancel

Next

These basic filters give you a simple way to apply different rules to large numbers of objects. For example, your company's financial records might need to be stored to meet regulatory requirements, while data from the marketing department might need to be stored to facilitate daily operations. After creating separate tenant accounts for each department or after segregating data from the different departments into separate S3 buckets, you can easily create one rule that applies to all financial records and a second rule that applies to all marketing data.

The **Advanced Filtering** page of the Create ILM Rule wizard gives you granular control. You can create filters to select objects based on the following object properties:

- Ingest time
- Last access time
- All or part of the object name (Key)
- S3 bucket region (Location Constraint)
- Object size
- User metadata
- S3 object tags

You can filter objects on very specific criteria. For example, objects stored by a hospital's imaging department might be used frequently when they are less than 30 days old and infrequently afterwards, while objects that contain patient visit information might need to be copied to the billing department at the health network's

headquarters. You can create filters that identify each type of object based on object name, size, S3 object tags, or any other relevant criteria, and then create separate rules to store each set of objects appropriately.

You can also combine basic and advanced filters as needed in a single rule. For example, the marketing department might want to store large image files differently than their vendor records, while the Human Resources department might need to store personnel records in a specific geography and policy information centrally. In this case you can create rules that filter by tenant account to segregate the records from each department, while using advanced filters in each rule to identify the specific type of objects that the rule applies to.


What ILM rule placement instructions are

Placement instructions determine where, when, and how object data is stored. An ILM rule can include one or more placement instructions. Each placement instruction applies to a single period of time.


When you create 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 the plus sign icon  to add more than one line for that time period.
- To place objects in different locations in different time periods, select the **Add** button to add the next time period. Then, specify one or more lines within the time period.

The example shows the Define Placements page of the Create ILM Rule wizard.

Placements  Sort by start day

From day store days Add Remove



Type Location Copies + x

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

Type Location Copies 1 + x

From day store Add Remove

Type Location Copies Temporary location 2 + x

	<p>The first placement instruction has two lines for the first year:</p> <ol style="list-style-type: none"> 1. The first line creates two replicated object copies at two data center sites. 2. The second line creates a 6+3 erasure-coded copy using three data center sites.
	<p>The second placement instruction creates two archived copies after one year and keeps those copies forever.</p>

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.

Example ILM rule

This example ILM rule applies to the objects belonging to Tenant A. It makes two replicated copies of those objects and stores each copy at a different site. The two copies are retained “forever,” which means that StorageGRID will not automatically delete them. Instead, StorageGRID will retain these objects until they are deleted by a client delete request or by the expiration of a bucket lifecycle.


This rule uses the Balanced option for ingest behavior: the two-site placement instruction is applied as soon as Tenant A saves an object to StorageGRID, unless it is not possible to immediately make both required copies. For example, if Site 2 is unreachable when Tenant A saves an object, StorageGRID will make two interim copies on Storage Nodes at Site 1. As soon as Site 2 becomes available, StorageGRID will make the required copy at that site.

Two copies at two sites for Tenant A

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

Matches all objects.

Retention Diagram:



Related information

- [Data-protection options for ingest](#)
- [What a storage pool is](#)
- [What a Cloud Storage Pool is](#)

Create storage grades, storage pools, EC profiles, and regions

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, instead of on all nodes at the site. For example, you might want certain objects to be stored on your fastest Storage Nodes, such as StorageGRID all-flash storage appliances.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

About this task

If you use more than one type of storage, you can optionally create a storage grade to identify each type. Creating storage grades allows you to select a specific type of Storage Node when configuring storage pools.

If storage grade is not a concern (for example, all Storage Nodes are identical), you can skip this procedure

and use the All Storage Nodes default storage grade when configuring storage pools.


When you add a new Storage Node in an expansion, that node is added to the All Storage Nodes default storage grade. As a result:

- If an ILM rule uses a storage pool with the All Storage Nodes grade, the new node can be used immediately after the expansion completes.
- If an ILM rule uses a storage pool with a custom storage grade, the new node will not be used until you manually assign the custom storage grade to the node, as described below.



When creating storage grades, do not create more storage grades than necessary. For example, do not create one storage grade for each Storage Node. Instead, assign each storage grade to two or more nodes. Storage grades assigned to only one node can cause ILM backlogs if that node becomes unavailable.

Steps

1. Select **ILM > Storage grades**.
2. Create a storage grade:
 - a. For each storage grade you need to define, select **Insert**  to add a row and enter a label for the storage grade.

The Default storage grade cannot be modified. It is reserved for new Storage Nodes added during a StorageGRID system expansion.



Storage Grades

Updated: 2017-05-26 11:22:39 MDT

Storage Grade Definitions

Storage Grade	Label	Actions
0	Default	
1	<input type="text" value="disk"/>	

Storage Grades

LDR	Storage Grade	Actions
Data Center 1/DC1-S1/LDR	Default	
Data Center 1/DC1-S2/LDR	Default	
Data Center 1/DC1-S3/LDR	Default	
Data Center 2/DC2-S1/LDR	Default	
Data Center 2/DC2-S2/LDR	Default	
Data Center 2/DC2-S3/LDR	Default	
Data Center 3/DC3-S1/LDR	Default	
Data Center 3/DC3-S2/LDR	Default	
Data Center 3/DC3-S3/LDR	Default	

Apply Changes

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



You cannot delete storage grades.










- c. Select **Apply Changes**.

These storage grades are now available for assignment to Storage Nodes.

3. Assign a storage grade to a Storage Node:

- a. For each Storage Node's LDR service, select **Edit** and select a storage grade from the list.

Storage Grades

LDR	Storage Grade	Actions
Data Center 1/DC1-S1/LDR	Default	
Data Center 1/DC1-S2/LDR	Default disk	
Data Center 1/DC1-S3/LDR	Default	
Data Center 2/DC2-S1/LDR	Default	
Data Center 2/DC2-S2/LDR	Default	
Data Center 2/DC2-S3/LDR	Default	
Data Center 3/DC3-S1/LDR	Default	
Data Center 3/DC3-S2/LDR	Default	
Data Center 3/DC3-S3/LDR	Default	

Apply Changes 



Assign a storage grade to a given Storage Node only once. A Storage Node recovered from failure maintains the previously assigned storage grade. Do not change this assignment after the ILM policy is activated. If the assignment is changed, data is stored based on the new storage grade.

b. Select **Apply Changes**.

Configure storage pools

What a storage pool is

A storage pool is a logical grouping of Storage Nodes or Archive Nodes. You configure storage pools to determine where the StorageGRID system stores object data and the type of storage used.

Storage pools have two attributes:

- **Storage grade:** For Storage Nodes, the relative performance of backing storage.
- **Site:** The data center where objects will be stored.

Storage pools are used in ILM rules to determine where object data is stored. When you configure ILM rules for replication, you select one or more storage pools that include either Storage Nodes or Archive Nodes. When you create Erasure Coding profiles, you select a storage pool that includes Storage Nodes.

Guidelines for creating storage pools

When configuring and using storage pools, follow these guidelines.

Guidelines for all storage pools

- StorageGRID includes a default storage pool, All Storage Nodes, that uses the default site, All Sites, and the default storage grade, All Storage Nodes. The All Storage Nodes storage pool is automatically updated

whenever you add new data center sites.



Using the All Storage Nodes storage pool or the All Sites site is not recommended because these items are automatically updated to include any new sites you add in an expansion, which might not be the behavior you want. Before using the All Storage Nodes storage pool or the default site, carefully review the guidelines for replicated and erasure-coded copies.

- Keep storage pool configurations as simple as possible. Do not create more storage pools than necessary.
- Create storage pools with as many nodes as possible. Each storage pool should contain two or more nodes. A storage pool with insufficient nodes can cause ILM backlogs if a node becomes unavailable.
- Avoid creating or using storage pools that overlap (contain one or more of the same nodes). If storage pools overlap, more than one copy of object data might be saved on the same node.

Guidelines for storage pools used for replicated copies

- Create a different storage pool for each site. Then, specify one or more site-specific storage pools in the placement instructions for each rule. Using a storage pool for each site ensures that replicated object copies are placed exactly where you expect (for example, one copy of every object at each site for site-loss protection).
- If you add a site in an expansion, create a new storage pool for the new site. Then, update ILM rules to control which objects are stored on the new site.
- In general, do not use the default storage pool, All Storage Nodes, or any storage pool that includes the default site, All Sites.

Guidelines for storage pools used for erasure-coded copies

- You cannot use Archive Nodes for erasure-coded data.
- The number of Storage Nodes and sites contained in the storage pool determine which erasure-coding schemes are available.
- If a storage pool includes only two sites, you cannot use that storage pool for erasure coding. No erasure-coding schemes are available for a storage pool that has two sites.
- In general, do not use the default storage pool, All Storage Nodes, or any storage pool that includes the default site, All Sites in any Erasure Coding profile.



If your grid includes only one site, you are prevented from using the All Storage Nodes storage pool or the All Sites default site in an Erasure Coding profile. This behavior prevents the Erasure Coding profile from becoming invalid if a second site is added.

- If you have high throughput requirements, creating a storage pool that includes multiple sites is not recommended if the network latency between sites is greater than 100 ms. As latency increases, the rate at which StorageGRID can create, place, and retrieve object fragments decreases sharply due to the decrease in TCP network throughput. The decrease in throughput affects the maximum achievable rates of object ingest and retrieval (when Strict or Balanced are selected as the Ingest Behavior) or could lead to ILM queue backlogs (when Dual Commit is selected as the Ingest Behavior).
- If possible, a storage pool should include more than the minimum number of Storage Nodes required for the erasure-coding scheme you select. For example, if you use a 6+3 erasure-coding scheme, you must have at least nine Storage Nodes. However, having at least one additional Storage Node per site is recommended.
- Distribute Storage Nodes across sites as evenly as possible. For example, to support a 6+3 erasure-coding

scheme, configure a storage pool that includes at least three Storage Nodes at three sites.

Guidelines for storage pools used for archived copies

- You cannot create a storage pool that includes both Storage Nodes and Archive Nodes. Archived copies require a storage pool that only includes Archive Nodes.
- When using a storage pool that includes Archive Nodes, you should also maintain at least one replicated or erasure-coded copy on a storage pool that includes Storage Nodes.
- If the global S3 Object Lock setting is enabled and you are creating a compliant ILM rule, you cannot use a storage pool that includes Archive Nodes. See the instructions for managing objects with S3 Object Lock.
- If an Archive Node's Target Type is Cloud Tiering - Simple Storage Service (S3), the Archive Node must be in its own storage pool. See [Administer StorageGRID](#).


Related information

- [What replication is](#)
- [What erasure coding is](#)
- [What erasure-coding schemes are](#)
- [Use multiple storage pools for cross-site replication](#)


Use multiple storage pools for cross-site replication

If your StorageGRID deployment includes more than one site, you can enable site-loss protection by creating a storage pool for each site and specifying both storage pools in the rule's placement instructions. For example, if you configure an ILM rule to make two replicated copies and specify storage pools at two sites, one copy of each object will be placed at each site. If you configure a rule to make two copies and specify three storage pools, the copies are distributed to balance disk usage among the storage pools, while ensuring that the two copies are stored at different sites.

The following example illustrates what can happen if an ILM rule places replicated object copies to a single storage pool containing Storage Nodes from two sites. Because the system uses any available nodes in the storage pool when it places the replicated copies, it might place all copies of some objects within only one of the sites. In this example, the system stored two copies of object AAA on Storage Nodes at Site 1, and two copies of object CCC on Storage Nodes at Site 2. Only object BBB is protected if one of the sites fails or becomes inaccessible.



In contrast, this example illustrates how objects are stored when you use multiple storage pools. In the example, the ILM rule specifies that two replicated copies of each object be created, and that the copies be distributed to two storage pools. Each storage pool contains all Storage Nodes at one site. Because a copy of each object is stored at each site, object data is protected from site failure or inaccessibility.



When using multiple storage pools, keep the following rules in mind:

- If you are creating n copies, you must add n or more storage pools. For example, if a rule is configured to make three copies, you must specify three or more storage pools.
- If the number of copies equals the number of storage pools, one copy of the object is stored in each storage pool.
- If the number of copies is less than the number of storage pools, the system distributes the copies to keep disk usage among the pools balanced and to ensure that two or more copies are not stored in the same storage pool.
- If the storage pools overlap (contain the same Storage Nodes), all copies of the object might be saved at

only one site. You must ensure that the selected storage pools do not contain the same Storage Nodes.

Use a storage pool as a temporary location (deprecated)

When you create an ILM rule with an object placement that includes a single storage pool, you are prompted to specify a second storage pool to use as a temporary location.

Temporary locations have been deprecated and will be removed in a future release. You should not select a storage pool as a temporary location for a new ILM rule.



If you select the Strict ingest behavior (Step 3 of the Create ILM Rule wizard), the temporary location is ignored.

Related information

[Data-protection options for ingest](#)

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.

What you'll need

- 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. Select the plus icon **+** to add an entry for each site.



Do not include the default All Sites site in a storage pool that will be used in an Erasure Coding profile. Instead, add a separate entry to the storage pool for each site that will store erasure coded data. See [this step](#) for an example.

- If you have more than one storage grade, do not create a storage pool that includes different storage grades at a single site. See the [Guidelines for creating storage pools](#).

Steps

1. Select **ILM > Storage pools**.

The Storage Pools page appears and lists all defined storage pools.

Storage Pools

Storage Pools

A storage pool is a logical group of Storage Nodes or Archive Nodes and is used in ILM rules to determine where object data is stored.

</

Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

+ Create	✎ Edit	✕ Remove	Clear Error
--------------------------	------------------------	--------------------------	-----------------------------

No Cloud Storage Pools found.

The list includes the system-default storage pool, All Storage Nodes, which uses the system-default site, All Sites, and the default storage grade, All Storage Nodes.



Because the All Storage Nodes storage pool is automatically updated whenever you add new data center sites, using this storage pool in ILM rules is not recommended.

- To create a new storage pool, select **Create**.

The Create Storage Pool dialog box appears.

Create Storage Pool

- For replication and single-site erasure coding, create a storage pool for each site.
- For erasure coding at three or more sites, click + to add each site to a single storage pool.
- Do not add more than one storage grade for a single site.

Name

Site

-- Choose One --

Storage Grade

All Storage Nodes

+

Viewing Storage Pool -

Site Name	Archive Nodes	Storage Nodes
-----------	---------------	---------------

Cancel

Save

- 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.

- 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, do not use the default 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 default All Storage Nodes storage grade includes all Storage Nodes at the selected site. The default Archive Nodes storage grade includes all Archive Nodes at the selected site. If you created additional storage grades for the Storage Nodes in your grid, they are listed in the drop-down.

6. If you want to use the storage pool in a multi-site Erasure Coding profile, select **+** to add an entry for each site to the storage pool.

Create Storage Pool

For replication and single-site erasure coding, create a storage pool for each site.

For erasure coding at three or more sites, select + to add each site to a single storage pool.

Do not select more than one storage grade for a single site.

Name

All 3 Sites for Erasure Coding

Site

Data Center 1

Storage Grade

All Storage Nodes

X

Site

Data Center 2

Storage Grade

All Storage Nodes

X

Site

Data Center 3

Storage Grade

All Storage Nodes

+X

Viewing Storage Pool - All 3 Sites for Erasure Coding

Site Name	Archive Nodes	Storage Nodes
Data Center 1	0	3
Data Center 2	0	3
Data Center 3	0	3

You are creating a multi-site storage pool, which should not be used for replication or single-site erasure coding.

Cancel

Save



You are prevented from creating duplicate entries or from creating a storage pool that includes both the **Archive Nodes** storage grade and any storage grade that contains Storage Nodes.

You are warned if you add more than one entry for a site but with different storage grades.

To remove an entry, select **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.

What you'll need

- 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 page appears. This page lists all defined storage pools.

Storage Pools

Storage Pools

A storage pool is a logical group of Storage Nodes or Archive Nodes and is used in ILM rules to determine where object data is stored.

<div><div>+ Create</div><div>Edit</div><div>✕ Remove</div><div>View Details</div></div>					
	Name ?	Used Space ?	Free Space ?	Total Capacity ?	ILM Usage ?
<input checked="" type="radio"/>	All Storage Nodes	1.88 MB	2.80 TB	2.80 TB	Used in 1 ILM rule
<input type="radio"/>	DC1	621.77 KB	932.42 GB	932.42 GB	Used in 2 ILM rules
<input type="radio"/>	DC2	675.82 KB	932.42 GB	932.42 GB	Used in 2 ILM rules
<input type="radio"/>	DC3	578.95 KB	932.42 GB	932.42 GB	Used in 1 ILM rule
<input type="radio"/>	All 3 Sites	1.88 MB	2.80 TB	2.80 TB	Used in 1 ILM rule and 1 EC profile
<input type="radio"/>	Archive	—	—	—	—

Displaying 6 storage pools.

Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

<div><div>+ Create</div><div>Edit</div><div>✕ Remove</div><div>Clear Error</div></div>			
No Cloud Storage Pools found.			

The table includes the following information for each storage pool that includes Storage Nodes:

- **Name:** The unique display name of the storage pool.
- **Used Space:** The amount of space that is currently being used to store objects in the storage pool.
- **Free Space:** The amount of space that remains available to store objects in the storage pool.
- **Total Capacity:** The size of the storage pool, which equals the total amount of usable space for object data for all nodes in the storage pool .
- **ILM Usage:** How the storage pool is currently being used. A storage pool might be unused or it might be used in one or more ILM rules, Erasure Coding profiles, or both.



You cannot remove a storage pool if it is being used.

2. To view details about a specific storage pool, select its radio button and select **View Details**.

The Storage Pool Details modal appears.

3. View the **Nodes Included** tab to learn about the Storage Nodes or Archive Nodes included in the storage pool.

Storage Pool Details - DC1

Nodes Included

ILM Usage

Number of Nodes: 3

Site - Storage Grade: DC1 - All Storage Nodes

Node Name	Site Name	Used (%) ?	↑↓
DC1-S3	DC1	0.000%	
DC1-S2	DC1	0.000%	
DC1-S1	DC1	0.000%	

Close

The table includes the following information for each node:

- Node Name
- Site Name
- Used (%): For Storage Nodes, the percentage of the total usable space for object data that has been used. This value does not include object metadata.



The same Used (%) value is also shown in the Storage Used - Object Data chart for each Storage Node (select **NODES** > **Storage Node** > **Storage**).

4. Select the **ILM Usage** tab to determine if the storage pool is currently being used in any ILM rules or Erasure Coding profiles.

In this example, the DC1 storage pool is used in three ILM rules: two rules that are in the active ILM policy and one rule that is not in the active policy.

Storage Pool Details - DC1

Nodes Included

ILM Usage

ILM Rules Using the Storage Pool

The following ILM rules in the active ILM policy (Example ILM policy) use this storage pool.

- 3 copies for Account01
- 2 copies for smaller objects

1 ILM rule that is not in the active ILM policy uses this storage pool.

If you want to remove this storage pool, you must delete or edit every rule where it is used. Go to the [ILM Rules page](#).

EC Profiles Using the Storage Pool

No Erasure Coding profiles use this storage pool.

Close



You cannot remove a storage pool if it is used in an ILM rule.

In this example, the All 3 Sites storage pool is used in an Erasure Coding profile. In turn, that Erasure Coding profile is used by one ILM rule in the active ILM policy.

Storage Pool Details - All 3 Sites

Nodes Included

ILM Usage

ILM Rules Using the Storage Pool

The following ILM rules in the active ILM policy (Example ILM policy) use this storage pool.

- EC larger objects

If you want to remove this storage pool, you must delete or edit every rule where it is used. Go to the [ILM Rules page](#).

EC Profiles Using the Storage Pool

The following Erasure Coding profiles use this storage pool.

Profile Name	Profile Status ?
6 plus 3	Used in 1 ILM Rule

Close



You cannot remove a storage pool if it is used in an Erasure Coding profile.

5. Optionally, go to the **ILM Rules page** to learn about and manage any rules that use the storage pool.

See the instructions for working with ILM rules.

6. When you are done viewing storage pool details, select **Close**.

Related information

[Work with ILM rules and ILM policies](#)

Edit storage pool

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

What you'll need

- 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 storage grade to a storage pool that is used in the active ILM policy, be aware that the Storage Nodes in the new storage grade will not be used automatically. To force StorageGRID to use a new storage grade, you must activate a new ILM policy after saving the edited storage pool.

Steps

1. Select **ILM > Storage pools**.

The Storage Pools page appears.

2. Select the radio button for the storage pool you want to edit.

You cannot edit the All Storage Nodes storage pool.

3. Select **Edit**.
4. As required, change the storage pool name.
5. As required, select other sites and storage grades.



You are prevented from changing the site or storage grade if the storage pool is used in an Erasure Coding profile and the change would cause the erasure-coding scheme to become invalid. For example, if a storage pool used in a Erasure Coding profile currently includes a storage grade with only one site, you are prevented from using a storage grade with two sites because the change would make the erasure-coding scheme invalid.

6. Select **Save**.

After you finish

If you added a new storage grade to a storage pool used in the active ILM policy, activate a new ILM policy to force StorageGRID to use the new storage grade. For example, clone your existing ILM policy and then activate the clone.

Remove a storage pool

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

What you'll need

- 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 page appears.

2. Look at the ILM Usage column in the table to determine whether you can remove the storage pool.

You cannot remove a storage pool if it is being used in an ILM rule or in an Erasure Coding profile. As required, select **View Details > ILM Usage** to determine where a storage pool is used.

3. If the storage pool you want to remove is not being used, select the radio button.
4. Select **Remove**.
5. Select **OK**.

Use Cloud Storage Pools

What a Cloud Storage Pool is

A Cloud Storage Pool lets you use ILM to move object data outside of your StorageGRID system. For example, you might want to move infrequently accessed objects to lower-cost cloud storage, such as Amazon S3 Glacier, S3 Glacier Deep Archive, or the Archive access tier in Microsoft Azure Blob storage. Or, you might want to maintain a cloud backup of StorageGRID objects to enhance disaster recovery.

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

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

	Storage pool	Cloud Storage Pool
How is it created?	Using the ILM > Storage pools option in Grid Manager. You must set up storage grades before you can create the storage pool.	Using the ILM > Storage pools option in Grid Manager. You must set up the external bucket or container before you can create the Cloud Storage Pool.
How many pools can you create?	Unlimited.	Up to 10.

	Storage pool	Cloud Storage Pool
Where are objects stored?	On one or more Storage Nodes or Archive Nodes within StorageGRID.	<p>In an Amazon S3 bucket or Azure Blob storage container that is external to the StorageGRID system.</p> <p>If the Cloud Storage Pool is an Amazon S3 bucket:</p> <ul style="list-style-type: none"> You can optionally configure a bucket lifecycle to transition objects to low-cost, long-term storage, such as Amazon S3 Glacier or S3 Glacier Deep Archive. The external storage system must support the Glacier storage class and the S3 POST Object restore API. You can create Cloud Storage Pools for use with AWS Commercial Cloud Services (C2S), which supports the AWS Secret Region. <p>If the Cloud Storage Pool is an Azure Blob storage container, StorageGRID transitions the object to the Archive tier.</p> <p>Note: In general, do not configure Azure Blob Storage lifecycle management for the container used for a Cloud Storage Pool. POST Object restore operations on objects in the Cloud Storage Pool can be affected by the configured lifecycle.</p>
What controls object placement?	An ILM rule in the active ILM policy.	An ILM rule in the active ILM policy.
What data protection method is used?	Replication or erasure coding.	Replication.
How many copies of each object are allowed?	Multiple.	<p>One copy in the Cloud Storage Pool and, optionally, one or more copies in StorageGRID.</p> <p>Note: You cannot store an object in more than one Cloud Storage Pool at any given time.</p>
What are the advantages?	Objects are quickly accessible at any time.	Low-cost storage.

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](#)
- [Azure: Lifecycle of a Cloud Storage Pool object](#)

S3: Lifecycle of a Cloud Storage Pool object


The figure shows the lifecycle stages of an object that is stored in an S3 Cloud Storage Pool.



In the figure and explanations, “Glacier” refers to both the Glacier storage class and the Glacier Deep Archive storage class, with one exception: the Glacier Deep Archive storage class does not support the Expedited restore tier. Only Bulk or Standard retrieval is supported.



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 GET Object request from StorageGRID, unless the object has been transitioned to Glacier storage.

3. Object transitioned to Glacier (non-retrievable state)

- Optionally, the object can be transitioned to Glacier storage. For example, the external S3 bucket might use lifecycle configuration to transition an object to Glacier storage immediately or after some number of days.



If you want to transition objects, you must create a lifecycle configuration for the external S3 bucket, and you must use a storage solution that implements the Glacier storage class and supports the S3 POST Object restore API.



Do not use Cloud Storage Pools for objects that have been ingested by Swift clients. Swift does not support POST Object restore requests, so StorageGRID will not be able to retrieve any Swift objects that have been transitioned to S3 Glacier storage. Issuing a Swift GET object request to retrieve these objects will fail (403 Forbidden).

- During the transition, the client application can use an S3 HEAD Object request to monitor the object's status.

4. Object restored from Glacier storage

If an object has been transitioned to Glacier storage, the client application can issue an S3 POST Object restore request to restore a retrievable copy to the S3 Cloud Storage Pool. The request specifies how many days the copy should be available in the Cloud Storage Pool and the data-access tier to use for the restore operation (Expedited, Standard, or Bulk). When the expiration date of the retrievable copy is reached, the copy is automatically returned to a non-retrievable state.




If one or more copies of the object also exist on Storage Nodes within StorageGRID, there is no need to restore the object from Glacier by issuing a POST Object restore request. Instead, the local copy can be retrieved directly, using a GET Object request.

5. Object retrieved

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

Azure: Lifecycle of a Cloud Storage Pool object

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



1. Object stored in StorageGRID

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

2. Object moved to Azure Cloud Storage Pool

When the object is matched by an ILM rule that uses an Azure Cloud Storage Pool as its placement location, StorageGRID moves the object to the external Azure Blob storage container specified by the Cloud Storage Pool.



Do not use Cloud Storage Pools for objects that have been ingested by Swift clients. Swift does not support POST Object restore requests, so StorageGRID will not be able to retrieve any Swift objects that have been transitioned to the Azure Blob storage Archive tier. Issuing a Swift GET object request to retrieve these objects will fail (403 Forbidden).

3. Object transitioned to Archive tier (non-retrievable state)

Immediately after moving the object to the Azure Cloud Storage Pool, StorageGRID automatically transitions the object to the Azure Blob storage Archive tier.

4. Object restored from Archive tier

If an object has been transitioned to the Archive tier, the client application can issue an S3 POST Object

restore request to restore a retrievable copy to the Azure Cloud Storage Pool.

When StorageGRID receives the POST Object Restore, it temporarily transitions the object to the Azure Blob storage Cool tier. As soon as the expiration date in the POST Object restore request is reached, StorageGRID transitions the object back to the Archive tier.



If one or more copies of the object also exist on Storage Nodes within StorageGRID, there is no need to restore the object from the Archive access tier by issuing a POST Object restore request. Instead, the local copy can be retrieved directly, using a GET Object request.

5. Object retrieved

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

Related information

[Use S3](#)

When to use Cloud Storage Pools

Cloud Storage Pools can provide significant benefits in several use cases.

Backing up StorageGRID data in an external location

You can use a Cloud Storage Pool to back up StorageGRID objects to an external location.

If the copies in StorageGRID are inaccessible, the object data in the Cloud Storage Pool can be used to serve client requests. However, you might need to issue S3 POST Object restore request to access the backup object copy in the Cloud Storage Pool.

The object data in a Cloud Storage Pool can also be used to recover data lost from StorageGRID because of a storage volume or Storage Node failure. If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID temporarily restores the object and creates a new copy on the recovered Storage Node.

To implement a backup solution:

1. Create a single Cloud Storage Pool.
2. Configure an ILM rule that simultaneously stores object copies on Storage Nodes (as replicated or erasure-coded copies) and a single object copy in the Cloud Storage Pool.
3. Add the rule to your ILM policy. Then, simulate and activate the policy.

Tiering data from StorageGRID to external location

You can use a Cloud Storage Pool to store objects outside of the StorageGRID system. For example, suppose you have a large number of objects that you need to retain, but you expect to access those objects rarely, if ever. You can use a Cloud Storage Pool to tier the objects to lower-cost storage and to free up space in StorageGRID.

To implement a tiering solution:

1. Create a single Cloud Storage Pool.
2. Configure an ILM rule that moves rarely used objects from Storage Nodes to the Cloud Storage Pool.

3. Add the rule to your ILM policy. Then, simulate and activate the policy.

Maintain multiple cloud endpoints

You can configure multiple Cloud Storage Pools if you want to tier or back up object data to more than one cloud. The filters in your ILM rules let you specify which objects are stored in each Cloud Storage Pool. For example, you might want to store objects from some tenants or buckets in Amazon S3 Glacier and objects from other tenant or buckets in Azure Blob storage. Or, you might want to move data between Amazon S3 Glacier and Azure Blob storage. When using multiple Cloud Storage Pools, keep in mind that an object can be stored in only one Cloud Storage Pool at a time.

To implement multiple cloud endpoints:

1. Create up to 10 Cloud Storage Pools.
2. Configure ILM rules to store the appropriate object data at the appropriate time in each Cloud Storage Pool. For example, store objects from bucket A in Cloud Storage Pool A, and store objects from bucket B in Cloud Storage Pool B. Or, store objects in Cloud Storage Pool A for some amount of time and then move them to Cloud Storage Pool B.
3. Add the rules to your ILM policy. Then, simulate and activate the policy.

Considerations for Cloud Storage Pools

If you plan to use a Cloud Storage Pool to move objects out of the StorageGRID system, you must review the considerations for configuring and using Cloud Storage Pools.

General considerations

- In general, cloud archival storage, such as Amazon S3 Glacier or Azure Blob storage, is an inexpensive place to store object data. However, the costs to retrieve data from cloud archival storage are relatively high. To achieve the lowest overall cost, you must consider when and how often you will access the objects in the Cloud Storage Pool. Using a Cloud Storage Pool is recommended only for content that you expect to access infrequently.
- Do not use Cloud Storage Pools for objects that have been ingested by Swift clients. Swift does not support POST Object restore requests, so StorageGRID will not be able to retrieve any Swift objects that have been transitioned to S3 Glacier storage or the Azure Blob storage Archive tier. Issuing a Swift GET object request to retrieve these objects will fail (403 Forbidden).
- Using Cloud Storage Pools with FabricPool is not supported because of the added latency to retrieve an object from the Cloud Storage Pool target.

Information required to create a Cloud Storage Pool

Before you can create a Cloud Storage Pool, you must create the external S3 bucket or the external Azure Blob storage container that you will use for the Cloud Storage Pool. Then, when you create the Cloud Storage Pool in StorageGRID, you must specify the following information:

- The provider type: Amazon S3 or Azure Blob storage.
- If you select Amazon S3, whether the Cloud Storage Pool is for use with the AWS Secret Region (**CAP (C2S Access Portal)**).
- The exact name of the bucket or container.
- The service endpoint needed to access the bucket or container.

- The authentication needed to access the bucket or container:
 - **S3**: Optionally, an access key ID and secret access key.
 - **C2S**: The complete URL for obtaining temporary credentials from the CAP server; a server CA certificate, a client certificate, a private key for the client certificate, and, if the private key is encrypted, the passphrase for decrypting it.
 - **Azure Blob storage**: An account name and account key. These credentials must have full permission for the container.
- Optionally, a custom CA certificate to verify TLS connections to the bucket or container.

Considerations for the ports used for Cloud Storage Pools

To ensure that the ILM rules can move objects to and from the specified Cloud Storage Pool, you must configure the network or networks that contain your system's Storage Nodes. You must ensure that the following ports can communicate with the Cloud Storage Pool.

By default, Cloud Storage Pools use the following ports:

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

You can specify a different port when you create or edit a Cloud Storage Pool.

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

Considerations for costs

Access to storage in the cloud using a Cloud Storage Pool requires network connectivity to the cloud. You must consider the cost of the network infrastructure you will use to access the cloud and provision it appropriately, based on the amount of data you expect to move between StorageGRID and the cloud using the Cloud Storage Pool.

When StorageGRID connects to the external Cloud Storage Pool endpoint, it issues various requests to monitor connectivity and to ensure it can perform the required operations. While some additional costs will be associated with these requests, the cost of monitoring a Cloud Storage Pool should only be a small fraction of the overall cost of storing objects in S3 or Azure.

More significant costs might be incurred if you need to move objects from an external Cloud Storage Pool endpoint back to StorageGRID. Objects might be moved back to StorageGRID in either of these cases:

- The only copy of the object is in a Cloud Storage Pool and you decide to store the object in StorageGRID instead. In this case, you simply reconfigure your ILM rules and policy. When ILM evaluation occurs, StorageGRID issues multiple requests to retrieve the object from the Cloud Storage Pool. StorageGRID then creates the specified number of replicated or erasure-coded copies locally. After the object is moved back to StorageGRID, the copy in the Cloud Storage Pool is deleted.
- Objects are lost because of Storage Node failure. If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID temporarily restores the object and creates a new copy on the recovered Storage Node.



When objects are moved back to StorageGRID from a Cloud Storage Pool, StorageGRID issues multiple requests to the Cloud Storage Pool endpoint for each object. Before moving large numbers of objects, contact technical support for help in estimating the time frame and associated costs.

S3: Permissions required for the Cloud Storage Pool bucket

The bucket policy for the external S3 bucket used for a Cloud Storage Pool must grant StorageGRID permission to move an object to the bucket, get an object's status, restore an object from Glacier storage when required, and more. Ideally, StorageGRID should have full-control access to the bucket (`s3: *`); however, if this is not possible, the bucket policy must grant the following S3 permissions to StorageGRID:

- `s3:AbortMultipartUpload`
- `s3:DeleteObject`
- `s3:GetObject`
- `s3:ListBucket`
- `s3:ListBucketMultipartUploads`
- `s3:ListMultipartUploadParts`
- `s3:PutObject`
- `s3:RestoreObject`

S3: Considerations for the external bucket's lifecycle

The movement of objects between StorageGRID and the external S3 bucket specified in the Cloud Storage Pool is controlled by ILM rules and the active ILM policy in StorageGRID. In contrast, the transition of objects from the external S3 bucket specified in the Cloud Storage Pool to Amazon S3 Glacier or S3 Glacier Deep Archive (or to a storage solution that implements the Glacier storage class) is controlled by that bucket's lifecycle configuration.

If you want to transition objects from the Cloud Storage Pool, you must create the appropriate lifecycle configuration on the external S3 bucket, and you must use a storage solution that implements the Glacier storage class and supports the S3 POST Object restore API.

For example, suppose you want all objects that are moved from StorageGRID to the Cloud Storage Pool to be transitioned to Amazon S3 Glacier storage immediately. You would create a lifecycle configuration on the external S3 bucket that specifies a single action (**Transition**) as follows:

```
<LifecycleConfiguration>
  <Rule>
    <ID>Transition Rule</ID>
    <Filter>
      <Prefix></Prefix>
    </Filter>
    <Status>Enabled</Status>
    <Transition>
      <Days>0</Days>
      <StorageClass>GLACIER</StorageClass>
    </Transition>
  </Rule>
</LifecycleConfiguration>
```

This rule would transition all bucket objects to Amazon S3 Glacier on the day they were created (that is, on the day they were moved from StorageGRID to the Cloud Storage Pool).



When configuring the external bucket's lifecycle, never use **Expiration** actions to define when objects expire. Expiration actions cause the external storage system to delete expired objects. If you later attempt to access an expired object from StorageGRID, the deleted object will not be found.

If you want to transition objects in the Cloud Storage Pool to S3 Glacier Deep Archive (instead of to Amazon S3 Glacier), specify `<StorageClass>DEEP_ARCHIVE</StorageClass>` in the bucket lifecycle. However, be aware that you cannot use the `Expedited` tier to restore objects from S3 Glacier Deep Archive.

Azure: Considerations for Access tier

When you configure an Azure storage account, you can set the default Access tier to Hot or Cool. When creating a storage account for use with a Cloud Storage Pool, you should use the Hot tier as the default tier. Even though StorageGRID immediately sets the tier to Archive when it moves objects to the Cloud Storage Pool, using a default setting of Hot ensures that you will not be charged an early deletion fee for objects removed from the Cool tier before the 30-day minimum.

Azure: Lifecycle management not supported

Do not use Azure Blob Storage lifecycle management for the container used with a Cloud Storage Pool. The lifecycle operations might interfere with Cloud Storage Pool operations.

Related information

- [Create a Cloud Storage Pool](#)
- [S3: Specify authentication details for a Cloud Storage Pool](#)
- [C2S S3: Specify authentication details for a Cloud Storage Pool](#)
- [Azure: Specify authentication details for a Cloud Storage Pool](#)

Comparing Cloud Storage Pools and CloudMirror replication

As you begin using Cloud Storage Pools, it might be helpful to understand the similarities

and differences between Cloud Storage Pools and the StorageGRID CloudMirror replication service.

	Cloud Storage Pool	CloudMirror replication service
What is the primary purpose?	A Cloud Storage Pool acts as an archive target. The object copy in the Cloud Storage Pool can be the only copy of the object, or it can be an additional copy. That is, instead of keeping two copies on-premise, you can keep only one copy within StorageGRID and send a copy to the Cloud Storage Pool.	The CloudMirror replication service enables a tenant to automatically replicate objects from a bucket in StorageGRID (source) to an external S3 bucket (destination). CloudMirror replication creates an independent copy of an object in an independent S3 infrastructure.
How is it set up?	Cloud Storage Pools are defined in the same way as storage pools, using the Grid Manager or the Grid Management API. A Cloud Storage Pool can be selected as the placement location in an ILM rule. While a storage pool consists of a group of Storage Nodes, a Cloud Storage Pool is defined using a remote S3 or Azure endpoint (IP address, credentials, and so on).	A tenant user configures CloudMirror replication by defining a CloudMirror endpoint (IP address, credentials, and so on) using the Tenant Manager or the S3 API. After the CloudMirror endpoint is set up, any bucket owned by that tenant account can be configured to point to the CloudMirror endpoint.
Who is responsible for setting it up?	Typically, a grid administrator	Typically, a tenant user
What is the destination?	<ul style="list-style-type: none"> Any compatible S3 infrastructure (including Amazon S3) Azure Blob Archive tier 	<ul style="list-style-type: none"> Any compatible S3 infrastructure (including Amazon S3)
What causes objects to be moved to the destination?	One or more ILM rules in the active ILM policy. The ILM rules define which objects StorageGRID moves to the Cloud Storage Pool and when the objects are moved.	The act of ingesting a new object into a source bucket that has been configured with a CloudMirror endpoint. Objects that existed in the source bucket before the bucket was configured with the CloudMirror endpoint are not replicated, unless they are modified.
How are objects retrieved?	Applications must make requests to StorageGRID to retrieve objects that have been moved to a Cloud Storage Pool. If the only copy of an object has been transitioned to archival storage, StorageGRID manages the process of restoring the object so it can be retrieved.	Because the mirrored copy in the destination bucket is an independent copy, applications can retrieve the object by making requests either to StorageGRID or to the S3 destination. For example, suppose you use CloudMirror replication to mirror objects to a partner organization. The partner can use its own applications to read or update objects directly from the S3 destination. Using StorageGRID is not required.

	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 in the Cloud Storage Pool.	The delete action is not replicated. A deleted object no longer exists in the StorageGRID bucket, but it continues to exist in the destination bucket. Similarly, objects in the destination bucket can be deleted without affecting the source.
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

When you create a Cloud Storage Pool, you specify the name and location of the external bucket or container that StorageGRID will use to store objects, the cloud provider type (Amazon S3 or Azure Blob Storage), and the information StorageGRID needs to access the external bucket or container.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.
- You have reviewed the guidelines for configuring Cloud Storage Pools.
- The external bucket or container referenced by the Cloud Storage Pool already exists.
- You have all of the authentication information needed to access the bucket or container.

About this task

A Cloud Storage Pool specifies a single external S3 bucket or Azure Blob storage container. StorageGRID validates the Cloud Storage Pool as soon as you save it, so you must ensure that the bucket or container specified in the Cloud Storage Pool exists and is reachable.

Steps

1. Select **ILM > Storage pools**.

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

Storage Pools

Storage Pools

A storage pool is a logical group of Storage Nodes or Archive Nodes and is used in ILM rules to determine where object data is stored.

</

Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

[+ Create](#) [Edit](#) [Remove](#) [Clear Error](#)

No Cloud Storage Pools found.

- In the Cloud Storage Pools section of the page, select **Create**.

The Create Cloud Storage Pool dialog box appears.

Create Cloud Storage Pool

Display Name

Provider Type

Bucket or Container

Cancel

Save

- Enter the following information:

Field	Description
Display Name	A name that briefly describes the Cloud Storage Pool and its purpose. Use a name that will be easy to identify when you configure ILM rules.
Provider Type	<p>Which cloud provider you will use for this Cloud Storage Pool:</p> <ul style="list-style-type: none">• Amazon S3: Select this option for an S3, C2S S3, or Google Cloud Platform (GCP) endpoint.• Azure Blob Storage <p>Note: When you select a Provider Type, the Service Endpoint, Authentication and Server Verification sections appear at the bottom on the page.</p>

Field	Description
Bucket or Container	The name of the external S3 bucket or Azure container that was created for the Cloud Storage Pool. The name you specify here must exactly match the bucket or container's name or Cloud Storage Pool creation will fail. You cannot change this value after the Cloud Storage Pool is saved.

4. Complete the Service Endpoint, Authentication and Server Verification sections of the page, based on the selected provider type.
 - [S3: Specify authentication details for a Cloud Storage Pool](#)
 - [C2S S3: Specify authentication details for a Cloud Storage Pool](#)
 - [Azure: Specify authentication details for a Cloud Storage Pool](#)


S3: Specifying authentication details for a Cloud Storage Pool


When you create a Cloud Storage Pool for S3, you must select the type of authentication that is required for the Cloud Storage Pool endpoint. You can specify Anonymous or enter an Access Key ID and Secret Access Key.


What you'll need

- You have entered the basic information for the Cloud Storage Pool and specified **Amazon S3** as the provider type.


Create Cloud Storage Pool


Display Name  S3 Cloud Storage Pool


Provider Type  Amazon S3 ▼


Bucket or Container  my-s3-bucket

Service Endpoint

Protocol  ☐ HTTP ☒ HTTPS

Hostname  example.com or 0.0.0.0


Port (optional)  443

URL Style  Auto-Detect ▼

Authentication

Authentication Type  ▼

Server Verification

Certificate Validation  Use operating system CA certificate ▼

Cancel

Save

- If you are using access key authentication, you know the Access Key ID and Secret Access Key for the external S3 bucket.

Steps

1. In the **Service Endpoint** section, provide the following information:

- a. Select which protocol to use when connecting to the Cloud Storage Pool.

The default protocol is HTTPS.

- b. Enter the server hostname or IP address of the Cloud Storage Pool.

For example:

`s3-aws-region.amazonaws.com`



Do not include the bucket name in this field. You include the bucket name in the **Bucket or Container** field.

- c. Optionally, specify the port that should be used when connecting to the Cloud Storage Pool.

Leave this field blank to use the default port: port 443 for HTTPS or port 80 for HTTP.

- d. Select the URL style for the Cloud Storage Pool bucket:

Option	Description
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, for example <code>https://bucket-name.s3.company.com/key-name</code> .
Path-Style	Use a path-style URL to access the bucket. Path-style URLs include the bucket name at the end, for example <code>https://s3.company.com/bucket-name/key-name</code> . Note: The path-style URL is being deprecated.
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.

2. In the **Authentication** section, select the type of authentication that is required for the Cloud Storage Pool endpoint.

Option	Description
Access Key	An Access Key ID and Secret Access Key are required to access the Cloud Storage Pool bucket.
Anonymous	Everyone has access to the Cloud Storage Pool bucket. An Access Key ID and Secret Access Key are not required.
CAP (C2S Access Portal)	Used for C2S S3 only. Go to C2S S3: Specifying authentication details for a Cloud Storage Pool .

3. If you selected Access Key, enter the following information:

Option	Description
Access Key ID	The Access Key ID for the account that owns the external bucket.
Secret Access Key	The associated Secret Access Key.

4. In the Server Verification section, select which method should be used to validate the certificate for TLS connections to the Cloud Storage Pool:

Option	Description
Use operating system CA certificate	Use the default Grid CA certificates installed on the operating system to secure connections.
Use custom CA certificate	Use a custom CA certificate. Select Select New , and upload the PEM-encoded CA certificate.
Do not verify certificate	The certificate used for the TLS connection is not verified.

5. Select **Save**.

When you save a Cloud Storage Pool, StorageGRID does the following:

- Validates that the bucket and the service endpoint exist and that they can be reached using the credentials that you specified.
- Writes a marker file to the bucket to identify the bucket as a Cloud Storage Pool. Never remove this file, which is named `x-ntap-sgws-cloud-pool-uuid`.

If Cloud Storage Pool validation fails, you receive an error message that explains why validation failed. For example, an error might be reported if there is a certificate error or if the bucket you specified does not already exist.

Error

422: Unprocessable Entity

Validation failed. Please check the values you entered for errors.

Cloud Pool test failed. Could not create or update Cloud Pool. Error from endpoint: NoSuchBucket: The specified bucket does not exist. status code: 404, request id: 4211567681, host id:

OK

See the instructions for [troubleshooting Cloud Storage Pools](#), resolve the issue, and then try saving the Cloud Storage Pool again.

C2S S3: Specify authentication details for a Cloud Storage Pool

To use the Commercial Cloud Services (C2S) S3 service as a Cloud Storage Pool, you must configure C2S Access Portal (CAP) as the authentication type, so that StorageGRID can request temporary credentials to access the S3 bucket in your C2S account.

What you'll need

- You have entered the basic information for an Amazon S3 Cloud Storage Pool, including the service endpoint.
- You know the complete URL that StorageGRID will use to obtain temporary credentials from the CAP server, including all the required and optional API parameters assigned to your C2S account.
- You have a server CA certificate issued by an appropriate Government Certificate Authority (CA). StorageGRID uses this certificate to verify the identity of the CAP server. The server CA certificate must use PEM encoding.
- You have a client certificate issued by an appropriate Government Certificate Authority (CA). StorageGRID uses this certificate to identify itself to the CAP server. The client certificate must use PEM encoding and must have been granted access to your C2S account.
- You have a PEM-encoded private key for the client certificate.
- If the private key for the client certificate is encrypted, you have the passphrase for decrypting it.

Steps


1. In the **Authentication** section, select **CAP (C2S Access Portal)** from the **Authentication Type** drop-down.

The CAP C2S authentication fields appear.

Create Cloud Storage Pool


Display Name  C2S Cloud Storage Pool

Provider Type  Amazon S3 ▼

Bucket or Container  my-c2s-bucket

Service Endpoint

Protocol  ☐ HTTP ☒ HTTPS

Hostname  s3-aws-region.amazonaws.com

Port (optional)  443

URL Style  Auto-Detect ▼

Authentication

Authentication Type  CAP (C2S Access Portal) ▼

Temporary Credentials URL  https://example.com/CAP/api/v1/cred


Server CA Certificate  [Select New](#)

Client Certificate  [Select New](#)

Client Private Key  [Select New](#)

Client Private Key
Passphrase (optional) 

Server Verification

Certificate Validation  Use operating system CA certificate ▼

Cancel

Save

2. Provide the following information:

- a. For **Temporary Credentials URL**, enter the complete URL that StorageGRID will use to obtain temporary credentials from the CAP server, including all the required and optional API parameters assigned to your C2S account.
- b. For **Server CA Certificate**, select **Select New**, and upload the PEM-encoded CA certificate that StorageGRID will use to verify the CAP server.
- c. For **Client Certificate**, select **Select New**, and upload the PEM-encoded certificate that StorageGRID will use to identify itself to the CAP server.
- d. For **Client Private Key**, select **Select New**, and upload the PEM-encoded private key for the client certificate.

If the private key is encrypted, the traditional format must be used. (PKCS #8 encrypted format is not supported.)

- e. If the client private key is encrypted, enter the passphrase for decrypting the client private key. Otherwise, leave the **Client Private Key Passphrase** field blank.

3. In the Server Verification section, provide the following information:

- a. For **Certificate Validation**, select **Use custom CA certificate**.
- b. Select **Select New**, and upload the PEM-encoded CA certificate.

4. Select **Save**.

When you save a Cloud Storage Pool, StorageGRID does the following:

- Validates that the bucket and the service endpoint exist and that they can be reached using the credentials that you specified.
- Writes a marker file to the bucket to identify the bucket as a Cloud Storage Pool. Never remove this file, which is named `x-ntap-sgws-cloud-pool-uuid`.

If Cloud Storage Pool validation fails, you receive an error message that explains why validation failed. For example, an error might be reported if there is a certificate error or if the bucket you specified does not already exist.

Error

422: Unprocessable Entity

Validation failed. Please check the values you entered for errors.

Cloud Pool test failed. Could not create or update Cloud Pool. Error from endpoint: NoSuchBucket:
The specified bucket does not exist. status code: 404, request id: 4211567681, host id:

OK

See the instructions for [troubleshooting Cloud Storage Pools](#), resolve the issue, and then try saving the Cloud Storage Pool again.

Azure: Specify authentication details for a Cloud Storage Pool

When you create a Cloud Storage Pool for Azure Blob storage, you must specify an account name and account key for the external container that StorageGRID will use to store objects.

What you'll need

- You have entered the basic information for the Cloud Storage Pool and specified **Azure Blob Storage** as the provider type. **Shared Key** appears in the **Authentication Type** field.

Create Cloud Storage Pool

Display Name ⓘ

Azure Cloud Storage Pool

Provider Type ⓘ

Azure Blob Storage ▼

Bucket or Container ⓘ

my-azure-container

Service Endpoint

URI ⓘ

https://myaccount.blob.core.windows.net

Authentication

Authentication Type ⓘ

Shared Key

Account Name ⓘ

Account Key ⓘ

Server Verification

Certificate Validation ⓘ

Use operating system CA certificate ▼

Cancel

Save

- You know the Uniform Resource Identifier (URI) used to access the Blob storage container used for the Cloud Storage Pool.
- You know the name of the storage account and the secret key. You can use the Azure portal to find these

values.

Steps

1. In the **Service Endpoint** section, enter the Uniform Resource Identifier (URI) used to access the Blob storage container used for the Cloud Storage Pool.

Specify the URI in one of the following formats:

- `https://host:port`
- `http://host:port`

If you do not specify a port, by default port 443 is used for HTTPS URIs and port 80 is used for HTTP URIs.

Example URI for Azure Blob storage container:

`https://myaccount.blob.core.windows.net`

2. In the **Authentication** section, provide the following information:
 - a. For **Account Name**, enter the name of the Blob storage account that owns the external service container.
 - b. For **Account Key**, enter the secret key for the Blob storage account.



For Azure endpoints, you must use Shared Key authentication.

3. In the **Server Verification** section, select which method should be used to validate the certificate for TLS connections to the Cloud Storage Pool:

Option	Description
Use operating system CA certificate	Use the Grid CA certificates installed on the operating system to secure connections.
Use custom CA certificate	Use a custom CA certificate. Select Select New , and upload the PEM-encoded certificate.
Do not verify certificate	The certificate used for the TLS connection is not verified.

4. Select **Save**.

When you save a Cloud Storage Pool, StorageGRID does the following:

- Validates that the container and the URI exist and that they can be reached using the credentials that you specified.
- Writes a marker file to the container to identify it as a Cloud Storage Pool. Never remove this file, which is named `x-ntap-sgws-cloud-pool-uuid`.

If Cloud Storage Pool validation fails, you receive an error message that explains why validation failed. For example, an error might be reported if there is a certificate error or if the container you specified does not already exist.

See the instructions for [troubleshooting Cloud Storage Pools](#), resolve the issue, and then try saving the Cloud

Storage Pool again.

Edit a Cloud Storage Pool

You can edit a Cloud Storage Pool to change its name, service endpoint, or other details; however, you cannot change the S3 bucket or Azure container for a Cloud Storage Pool.

What you'll need

- You 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**.

The Storage Pools page appears. The Cloud Storage Pools table lists the existing Cloud Storage Pools.

Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

+ Create	✎ Edit	✕ Remove	Clear Error			
	Pool Name	URI	Pool Type	Container	Used in ILM Rule	Last Error
<input checked="" type="radio"/>	azure-endpoint	https://storagegrid.blob.core.windows.net	azure	azure-3	✓	
<input type="radio"/>	s3-endpoint	https://s3.amazonaws.com	s3	s3-1	✓	
						Displaying 2 pools.

2. Select the radio button for the Cloud Storage Pool you want to edit.
3. Select **Edit**.
4. As required, change the display name, service endpoint, authentication credentials, or certificate validation method.



You cannot change the provider type or the S3 bucket or Azure container for a Cloud Storage Pool.

If you previously uploaded a server or client certificate, you can select **View Current** to review the certificate that is currently in use.

5. 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 that is not used in an ILM rule and that does not contain object data.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.
- You have confirmed that the S3 bucket or Azure container does not contain any objects. An error occurs if you attempt to remove a Cloud Storage Pool if it contains objects. See [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. Do not remove this file, which is named `x-ntap-sgws-cloud-pool-uuid`.

- You have already removed any ILM rules that might have used the pool.

Steps

1. Select **ILM > Storage pools**.

The Storage Pools page appears.

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

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

Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

+ Create

✎ Edit

✕ Remove

Clear Error

	Pool Name	URI	Pool Type	Container	Used in ILM Rule	Last Error
<input checked="" type="radio"/>	azure-endpoint	https://storagegrid.blob.core.windows.net	azure	azure-3	✓	
<input type="radio"/>	s3-endpoint	https://s3.amazonaws.com	s3	s3-1	✓	

Displaying 2 pools.

3. Select **Remove**.

A confirmation warning is displayed.

Warning

Remove Cloud Storage Pool

Are you sure you want to remove this Cloud Storage Pool: My Cloud Storage Pool?

Cancel

OK

4. Select **OK**.

The Cloud Storage Pool is removed.

Troubleshoot Cloud Storage Pools

If you encounter errors when creating, editing, or deleting a Cloud Storage Pool, use these troubleshooting steps to help resolve the issue.

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.

Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

+ Create

Edit

✕ Remove

Clear Error

	Pool Name	URI	Pool Type	Container	Used in ILM Rule	Last Error
<input checked="" type="radio"/>	S3	10.96.106.142:18082	s3	s3	✓	Endpoint failure: DC2-S1-106-147: Could not create or update Cloud Storage Pool. Error from endpoint: RequestError: send request failed caused by: Get https://10.96.106.142:18082/s3-targetbucket/x-ntap-sgws-cloud-pool-uuid: net/http: request canceled while waiting for connection (Client.Timeout exceeded while awaiting headers) 8 minutes ago
<input type="radio"/>	Azure	http://pboerkoe@10.96.100.254:10000/devstoreaccount1	azure	azure	✓	

Displaying 2 pools.

In addition, a **Cloud Storage Pool connectivity error** alert is triggered if the health check detects that one or more new Cloud Storage Pool errors have occurred within the past 5 minutes. If you receive an email notification for this alert, go to the Storage Pool page (select **ILM** > **Storage pools**), review the error messages in the Last Error column, and refer to the troubleshooting guidelines below.

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 radio button for the endpoint, and select **Clear Error**. A confirmation message indicates that StorageGRID has cleared the error for the Cloud Storage Pool.

Error successfully cleared. This error might reappear if the underlying problem is not resolved.



If the underlying problem has been resolved, the error message is no longer displayed. However, if the underlying problem has not been fixed (or if a different error is encountered), the error message will be shown in the Last Error column within a few minutes.

Error: This Cloud Storage Pool contains unexpected content

You might encounter this error when you try to create, edit, or delete a Cloud Storage Pool. This error occurs if the bucket or container includes the `x-ntap-sgws-cloud-pool-uuid` marker file, but that file does not

have the expected UUID.

Typically, you will only see this error if you are creating a new Cloud Storage Pool and another instance of StorageGRID is already using the same Cloud Storage Pool.

Try these steps to correct the issue:

- Check to make sure that no one in your organization is also using this Cloud Storage Pool.
- Delete the `x-ntap-sgws-cloud-pool-uuid` file and try configuring the Cloud Storage Pool again.

Error: Could not create or update Cloud Storage Pool. Error from endpoint

You might encounter this error when you try to create or edit a Cloud Storage Pool. This error indicates that some kind of connectivity or configuration issue is preventing StorageGRID from writing to the Cloud Storage Pool.

To correct the issue, review the error message from the endpoint.

- If the error message contains `Get url: EOF`, check that the service endpoint used for the Cloud Storage Pool does not use the HTTP protocol for a container or bucket that requires HTTPS.
- If the error message contains `Get url: net/http: request canceled while waiting for connection`, verify that the network configuration allows Storage Nodes to access the service endpoint used for the Cloud Storage Pool.
- For all other endpoint error messages, try one or more of the following:
 - Create an external container or bucket with the same name you entered for the Cloud Storage Pool, and try to save the new Cloud Storage Pool again.
 - Correct the container or bucket name you specified for the Cloud Storage Pool, and try to save the new Cloud Storage Pool again.

Error: Failed to parse CA certificate

You might encounter this error when you try to create or edit a Cloud Storage Pool. The error occurs if StorageGRID could not parse the certificate you entered when configuring the Cloud Storage Pool.

To correct the issue, check the CA certificate you provided for issues.

Error: A Cloud Storage Pool with this ID was not found

You might encounter this error when you try to edit or delete a Cloud Storage Pool. This error occurs if the endpoint returns a 404 response, which can mean either of the following:

- The credentials used for the Cloud Storage Pool do not have read permission for the bucket.
- The bucket used for the Cloud Storage Pool does not include the `x-ntap-sgws-cloud-pool-uuid` marker file.

Try one or more of these steps to correct the issue:

- Check that the user associated with the configured Access Key has the requisite permissions.
- Edit the Cloud Storage Pool with credentials that have the requisite permissions.
- If the permissions are correct, contact support.

Error: Could not check the content of the Cloud Storage Pool. Error from endpoint

You might encounter this error when you try to delete a Cloud Storage Pool. This error indicates that some kind of connectivity or configuration issue is preventing StorageGRID from reading the contents of the Cloud Storage Pool bucket.

To correct the issue, review the error message from the endpoint.

Error: Objects have already been placed in this bucket

You might encounter this error when you try to delete a Cloud Storage Pool. You cannot delete a Cloud Storage Pool if it contains data that was moved there by ILM, data that was in the bucket before you configured the Cloud Storage Pool, or data that was put in the bucket by some other source after the Cloud Storage Pool was created.

Try one or more of these steps to correct the issue:

- Follow the instructions for moving objects back to StorageGRID in “Lifecycle of a Cloud Storage Pool object.”
- If you are certain the remaining objects were not placed in the Cloud Storage Pool by ILM, manually delete the objects from the bucket.



Never manually delete objects from a Cloud Storage Pool that might have been placed there by ILM. If you later attempt to access a manually deleted object from StorageGRID, the deleted object will not be found.

Error: Proxy encountered an external error while trying to reach the Cloud Storage Pool

You might encounter this error if you have configured a non-transparent Storage proxy between Storage Nodes and the external S3 endpoint used for the Cloud Storage Pool. This error occurs if the external proxy server cannot reach the Cloud Storage Pool endpoint. For example, the DNS server might not be able to resolve the hostname or there might be an external networking issue.

Try one or more of these steps to correct the issue:

- Check the settings for the Cloud Storage Pool (**ILM > Storage pools**).
- Check the networking configuration of the Storage proxy server.

Related information

[Lifecycle of a Cloud Storage Pool object](#)

Configure Erasure Coding profiles

Create an Erasure Coding profile

To create an Erasure Coding profile, you associate a storage pool containing Storage Nodes with an erasure-coding scheme. This association determines the number of data and parity fragments created and where the system distributes these fragments.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).

- You have specific access permissions.
- You have created a storage pool that includes exactly one site or a storage pool that includes three or more sites. No erasure-coding schemes are available for a storage pool that has only two sites.

About this task

The storage pools used in Erasure Coding profiles must include exactly one site or three or more sites. If you want to provide site redundancy, the storage pool must have at least three sites.



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

Steps

1. Select **ILM > Erasure coding**.

The Erasure Coding Profiles page appears.

Erasure Coding Profiles

An Erasure Coding profile determines how many data and parity fragments are created and where those fragments are stored.

To create an Erasure Coding profile, select a [storage pool](#) and an erasure coding scheme. The storage pool must include Storage Nodes from exactly one site or from three or more sites. If you want to provide site redundancy, the storage pool must include nodes from at least three sites.

To deactivate an Erasure Coding profile that you no longer plan to use, first remove it from all ILM rules. Then, if the profile is still associated with object data, wait for those objects to be moved to new locations based on the new rules in the active ILM policy. Depending on the number of objects and the size of your StorageGRID system, it might take weeks or even months for the objects to be moved.

See [Managing objects with information lifecycle management](#) for important details.

<div> Create Rename Deactivate </div>									
Profile	Status	Storage Pool	Storage Nodes	Sites	Erasure Code	Storage Overhead (%)	Storage Node Redundancy	Site Redundancy	
No Erasure Coding profiles found.									

2. Select **Create**.

The Create EC Profile dialog box appears.

Create EC Profile

You cannot change the selected scheme and storage pool after saving the profile.

Profile Name

Storage Pool

Cancel

Save

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

Erasure Coding profile names must be unique. A validation error occurs if you use the name of an existing profile, even if that profile has been deactivated.



The Erasure Coding profile name is appended to the storage pool name in the placement instruction for an ILM rule.

From day store Add Remove

Type Location Copies + x

Erasure Coding profile name

Storage pool name

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



If your grid currently includes only one site, you are prevented from using the default storage pool, All Storage Nodes, or any storage pool that includes the default site, All Sites. This behavior prevents the Erasure Coding profile from becoming invalid if a second site is added.



If a storage pool includes exactly two sites, you cannot use that storage pool for erasure coding. No erasure-coding schemes are available for a storage pool that has two sites.

When you select a storage pool, the list of available erasure-coding schemes is shown, based on the number of Storage Nodes and sites in the pool.

Create EC Profile

You cannot change the selected scheme and storage pool after saving the profile.

Profile Name

Storage Pool

9 Storage Nodes across 3 site(s)

Scheme

	Erasure Code	Storage Overhead (%)	Storage Node Redundancy	Site Redundancy
<input checked="" type="radio"/>	6+3	50%	3	Yes
<input type="radio"/>	2+1	50%	1	Yes
<input type="radio"/>	4+2	50%	2	Yes

Cancel Save

The following information is listed for each available erasure-coding scheme:

- **Erasure Code:** The name of the erasure-coding scheme in the following format: data fragments + parity fragments.
- **Storage Overhead (%):** The additional storage required for parity fragments relative to the object's data size. Storage Overhead = Total number of parity fragments / Total number of data fragments.
- **Storage Node Redundancy:** The number of Storage Nodes that can be lost while still maintaining the ability to retrieve object data.
- **Site Redundancy:** Whether the selected erasure code allows the object data to be retrieved if a site is lost.

To support site redundancy, the selected storage pool must include multiple sites, each with enough Storage Nodes to allow any site to be lost. For example, to support site redundancy using a 6+3

erasure-coding scheme, the selected storage pool must include at least three sites with at least three Storage Nodes at each site.

Messages are displayed in these cases:

- The storage pool you selected does not provide site redundancy. The following message is expected when the selected storage pool includes only one site. You can use this Erasure Coding profile in ILM rules to protect against node failures.

Scheme

	Erasure Code ?	Storage Overhead (%) ?	Storage Node Redundancy ?	Site Redundancy ?
<input checked="" type="radio"/>	2+1	50%	1	No

The selected storage pool and erasure coding scheme cannot protect object data from loss if a site is lost.
To provide site redundancy, the storage pool must have at least three sites.

- The storage pool you selected does not satisfy the requirements for any erasure-coding scheme. For example, the following message is expected when the selected storage pool includes exactly two sites. If you want to use erasure coding to protect object data, you must select a storage pool with exactly one site or a storage pool with three or more sites.

Scheme

	Erasure Code ?	Storage Overhead (%) ?	Storage Node Redundancy ?	Site Redundancy ?
--	----------------	------------------------	---------------------------	-------------------

No erasure coding schemes are supported for the selected storage pool because it contains two sites. You must select a storage pool that contains exactly one site or a storage pool that contains at least three sites.

- Your grid includes only one site and you selected the default storage pool, All Storage Nodes, or any storage pool that includes the default site, All Sites.

Create EC Profile

You cannot change the selected scheme and storage pool after saving the profile.

Profile Name

Storage Pool

3 Storage Nodes across 1 site(s)

Scheme

	Erasure Code	Storage Overhead (%)	Storage Node Redundancy	Site Redundancy
--	--------------	----------------------	-------------------------	-----------------

No erasure coding schemes are available for the selected storage pool. The storage pool includes the **All Sites** site, so it cannot be used in an Erasure Coding profile for a one-site grid.


Cancel


Save

- The erasure-coding scheme and storage pool you selected overlap with another Erasure Coding profile.

Create EC Profile

You cannot change the selected scheme and storage pool after saving the profile.

Profile Name  2 plus 1 for three sites

Storage Pool  All 3 Sites 

9 Storage Nodes across 3 site(s)

Scheme

	Erasure Code 	Storage Overhead (%) 	Storage Node Redundancy 	Site Redundancy 
<input type="radio"/>	6+3	50%	3	Yes
<input checked="" type="radio"/>	2+1	50%	1	Yes
<input type="radio"/>	4+2	50%	2	Yes

The selected storage pool and erasure coding scheme overlap an existing Erasure Coding profile. Use caution if you apply this new profile to objects already protected by the other profile. When a new profile is applied to existing erasure-coded objects, entirely new erasure-coded fragments are created, which might cause resource issues.

Cancel Save

In this example, a warning message appears because another Erasure Coding profile is using the 2+1 scheme and the storage pool for the other profile also uses one of the sites in the All 3 Sites storage pool.

While you are not prevented from creating this new profile, you must be very careful when you start using it in the ILM policy. If this new profile is applied to existing erasure-coded objects already protected by the other profile, StorageGRID will create an entirely new set of object fragments. It will not reuse the existing 2+1 fragments. Resource issues might occur when you migrate from one Erasure Coding profile to the other, even though the erasure-coding schemes are the same.

5. If more than one erasure-coding scheme is listed, select the one you want to use.

When deciding which erasure-coding scheme to use, you should balance fault tolerance (achieved by having more parity segments) against the network traffic requirements for repairs (more fragments equals more network traffic). For example, when deciding between a 4+2 scheme and 6+3 scheme, select the 6+3 scheme if additional parity and fault tolerance are required. Select the 4+2 scheme if network resources are constrained to reduce network usage during node repairs.

6. Select **Save**.

Rename an Erasure Coding profile

You might want to rename an Erasure Coding profile to make it more obvious what the profile does.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

Steps

1. Select **ILM > Erasure coding**.

The Erasure Coding Profiles page appears. The **Rename** and **Deactivate** buttons are both disabled.

<div><div>+ Create</div><div><div>Rename</div><div>Deactivate</div></div></div>									
	Profile	Status	Storage Pool	Storage Nodes	Sites	Erasure Code	Storage Overhead (%)	Storage Node Redundancy	Site Redundancy
	DC1 2-1		DC1	3	1	2+1	50	1	No
	DC2 2-1		DC2	3	1	2+1	50	1	No
	DC3 2-1		DC3	3	1	2+1	50	1	No
	All sites 6-3	Deactivated	All 3 Sites	9	3	6+3	50	3	Yes

2. Select the profile you want to rename.

The **Rename** and **Deactivate** buttons become enabled.

3. Select **Rename**.

The Rename EC Profile dialog box appears.

Rename EC Profile

Profile Name

Cancel

Save

4. Enter a unique name for the Erasure Coding profile.

The Erasure Coding profile name is appended to the storage pool name in the placement instruction for an ILM rule.

From day store

Add

Remove

Type

Location

Copies

+

×

Storage pool name

Erasure Coding profile name



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.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).

- You have specific access permissions.
- You have confirmed that no erasure coded data repair operations or decommission procedures are in process. An error message is returned if you attempt to deactivate an Erasure Coding profile while either of these operations are in progress.

About this task

When you deactivate an Erasure Coding profile, the profile still appears on the Erasure Coding Profiles page, but its status is **Deactivated**.

+ Create

Rename

Deactivate

	Profile	Status	Storage Pool	Storage Nodes	Sites	Erasure Code	Storage Overhead (%)	Storage Node Redundancy	Site Redundancy
	DC1 2-1		DC1	3	1	2+1	50	1	No
	DC2 2-1		DC2	3	1	2+1	50	1	No
	DC3 2-1		DC3	3	1	2+1	50	1	No
	All sites 6-3	Deactivated	All 3 Sites	9	3	6+3	50	3	Yes

You can no longer use an Erasure Coding profile that has been deactivated. A deactivated profile is not shown when you create the placement instructions for an ILM rule. You cannot reactivate a deactivated profile.

StorageGRID prevents you from deactivating an Erasure Coding profile if either of the following is true:

- The Erasure Coding profile is currently used in an ILM rule.
- The Erasure Coding profile is no longer used in any ILM rules, but object data and parity fragments for the profile still exist.

Steps

1. Select **ILM > Erasure Coding**.

The Erasure Coding Profiles page appears. The **Rename** and **Deactivate** buttons are both disabled.

2. Review the **Status** column to confirm that the Erasure Coding profile you want to deactivate is not used in any ILM rules.

You cannot deactivate an Erasure Coding profile if it is used in any ILM rule. In the example, the **2_1 EC Profile** is used in at least one ILM rule.

+ Create

Rename

Deactivate

	Profile	Status	Storage Pool	Storage Nodes	Sites	Erasure Code	Storage Overhead (%)	Storage Node Redundancy	Site Redundancy
<input type="radio"/>	2_1 EC Profile	Used In ILM Rule	DC1	3	1	2+1	50	1	No
<input type="radio"/>	Site 1 EC Profile	Deactivated	DC1	3	1	2+1	50	1	No

3. If the profile is used in an ILM rule, follow these steps:

- a. Select **ILM > Rules**.
- b. For each rule listed, select the radio button and review the retention diagram to determine if the rule uses the Erasure Coding profile you want to deactivate.

In the example, the **Three site EC for larger objects** rule uses a storage pool called **All 3 Sites** and the **All sites 6-3** Erasure Coding profile. Erasure Coding profiles are represented by this icon: 

ILM Rules

Information lifecycle management (ILM) rules determine how and where object data is stored over time. Every object ingested into StorageGRID is evaluated against the ILM rules that make up the active ILM policy. Use this page to manage and view ILM rules. You cannot edit or remove an ILM rule that is used by an active or proposed ILM policy.

+ Create

Clone

Edit

Remove

Name	Used In Active Policy	Used In Proposed Policy
2 copy replication for smaller objects	✓	
Three site EC for larger objects	✓	
Make 2 Copies		

Three site EC for larger objects

Description:

6-3 erasure coding at 3 sites for objects larger than 200 KB

Ingest Behavior:

Balanced

Reference Time:

Ingest Time

Filtering Criteria:

Matches all of the following metadata:

System Metadata

Object Size (MB)

greater than

0.2

Retention Diagram:

Trigger

Day 0

Duration

Forever

- c. If the ILM rule uses the Erasure Coding profile you want to deactivate, determine if the rule is used in either the active ILM policy or a proposed policy.

In the example, the **Three site EC for larger objects** rule is used in the active ILM policy.

- d. Complete the additional steps in the table, based on where the Erasure Coding profile is used.

Where has the profile been used?	Additional steps to perform before deactivating the profile	Refer to these additional instructions
Never used in any ILM rule	No additional steps required. Continue with this procedure.	<i>None</i>
In an ILM rule that has never been used in any ILM policy	<ol style="list-style-type: none"> Edit or delete all affected ILM rules. If you edit the rule, remove all placements that use the Erasure Coding profile. Continue with this procedure. 	Work with ILM rules and ILM policies

Where has the profile been used?	Additional steps to perform before deactivating the profile	Refer to these additional instructions
In an ILM rule that is currently in the active ILM policy	<ol style="list-style-type: none"> 1. Clone the active policy. 2. Remove the ILM rule that uses the Erasure Coding profile. 3. Add one or more new ILM rules to ensure objects are protected. 4. Save, simulate, and activate the new policy. 5. Wait for the new policy to be applied and for existing objects to be moved to new locations based on the new rules you added. <p>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.</p> <p>While you can safely attempt to deactivate an Erasure Coding profile while it is still associated with data, the deactivation operation will fail. An error message will inform you if the profile is not yet ready to be deactivated.</p> <ol style="list-style-type: none"> 6. Edit or delete the rule you removed from the policy. If you edit the rule, remove all placements that use the Erasure Coding profile. 7. Continue with this procedure. 	<ul style="list-style-type: none"> • Create an ILM policy • Work with ILM rules and ILM policies
In an ILM rule that is currently in a proposed ILM policy	<ol style="list-style-type: none"> 1. Edit the proposed policy. 2. Remove the ILM rule that uses the Erasure Coding profile. 3. Add one or more new ILM rules to ensure all objects are protected. 4. Save the proposed policy. 5. Edit or delete the rule you removed from the policy. If you edit the rule, remove all placements that use the Erasure Coding profile. 6. Continue with this procedure. 	<ul style="list-style-type: none"> • Create an ILM policy • Work with ILM rules and ILM policies

Where has the profile been used?	Additional steps to perform before deactivating the profile	Refer to these additional instructions
In an ILM rule that is in a historical ILM policy	<ol style="list-style-type: none"> 1. Edit or delete the rule. If you edit the rule, remove all placements that use the Erasure Coding profile. (The rule will now appear as a historical rule in the historical policy.) 2. Continue with this procedure. 	Work with ILM rules and ILM policies

e. Refresh the Erasure Coding Profiles page to ensure that the profile is not used in an ILM rule.

4. If the profile is not used in an ILM rule, select the radio button and select **Deactivate**.

The Deactivate EC Profile dialog box appears.

Deactivate EC Profile

Are you sure you want to deactivate the profile 'All sites 6-3'?

StorageGRID will confirm that the profile is safe to remove (not used in any ILM rules and no longer associated with any object data). After this profile is deactivated, you can no longer use it.

Cancel
Deactivate

5. If you are sure you want to deactivate the profile, select **Deactivate**.

- If StorageGRID is able to deactivate the Erasure Coding profile, its status is **Deactivated**. You can no longer select this profile for any ILM rule.
- If StorageGRID is not able to deactivate the profile, an error message appears. For example, an error message appears if object data is still associated with this profile. You might need to wait several weeks before trying the deactivation process again.

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.

What you'll need

- 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 and must contain at least 2 and no more than 32 characters. Valid characters are numbers, letters, and hyphens.



EU is not considered to be an alias for eu-west-1. If you want to use the EU or eu-west-1 region, you must use the exact name.

- You cannot delete or modify a region if it is currently used within the active ILM policy or the proposed ILM policy.
- If the region used as the advanced filter in an ILM rule is invalid, it is still possible to add that rule to the proposed policy. However, an error occurs if you attempt to save or activate the proposed policy. (An invalid region can result if you use a region as an advanced filter in an ILM rule but you later delete that region, or if you use the Grid Management API to create a rule and specify a region that you have not defined.)
- If you delete a region after using it to create an S3 bucket, you will need to re-add the region if you ever want to use the Location Constraint advanced filter to find objects in that bucket.

Steps

1. Select **ILM > Regions**.

The Regions page appears, with the currently defined regions listed. **Region 1** shows the default region, us-east-1, which cannot be modified or removed.

Regions (optional and S3 only)

Define any regions you want to use for the Location Constraint advanced filter in ILM rules. Then, use these exact names when creating S3 buckets. (Region names are case sensitive.)

Region 1

us-east-1 (required)

Region 2

us-west-1

+

×

Save


2. To add a region:
 - a. Select the insert icon **+** to the right of the last entry.

- b. Enter the name of a region that you want to use when creating S3 buckets.

You must use this exact region name as the LocationConstraint request element when you create the corresponding S3 bucket.

3. To remove an unused region, select the delete icon ✕.

An error message appears if you attempt to remove a region that is currently used in the active policy or the proposed policy.



4. When you are done making changes, select **Save**.

You can now select these regions from the **Location Constraint** list on the Advanced Filtering page of the Create ILM rule wizard. See [Use advanced filters in ILM rules](#).

Create ILM rule

Access the Create 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 ILM Rule wizard.



If you are creating the default ILM rule for a policy, use this procedure instead: [Create a default ILM rule](#).

What you'll need

- 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.
- If you are creating replicated copies, you have configured any storage pools or Cloud Storage Pools you plan to use. See [Create storage pool](#) and [Create Cloud Storage Pool](#).
- If you are creating erasure-coded copies, you have configured an Erasure Coding profile. See [Create an Erasure Coding profile](#).
- You are familiar with the [data-protection options for ingest](#).

- 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: StorageGRID ILM Rules: Getting Started](#).



About this task

When creating ILM rules:

- Consider the StorageGRID system's topology and storage configurations.
- Consider what types of object copies you want to make (replicated or erasure coded) and the number of copies of each object that are required.
- Determine what types of object metadata are used in the applications that connect to the StorageGRID system. ILM rules filter objects based on their metadata.
- Consider where you want object copies to be placed over time.
- Decide which option to use for data protection option at ingest (Balanced, Strict, or Dual commit).

Steps

1. Select **ILM > Rules**.

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

CreateCloneEditRemove

Name	Used In Active Policy	Used In Proposed Policy
Make 2 Copies	✓	

Make 2 Copies

Ingest Behavior: Dual commit

Reference Time: Ingest Time

Filtering Criteria:

Matches all objects.

Retention Diagram:


Trigger

Day 0

All Storage Nodes

Duration

Forever



The ILM Rules page looks slightly different if the global S3 Object Lock setting has been enabled for the StorageGRID system. The summary table includes a **Compliant** column, and the details for the selected rule include a **Compliant** field.

2. Select **Create**.

Step 1 (Define Basics) of the Create ILM Rule wizard appears. You use the Define basics page to define which objects the rule applies to.

Step 1 of 3: Define basics

Step 1 (Define Basics) of the Create ILM Rule wizard allows you to define the rule’s basic and advanced filters.

About this task

When evaluating an object against an ILM rule, StorageGRID compares the object metadata to the rule’s filters. If the object metadata matches all filters, StorageGRID uses the rule to place the object. You can design a rule to apply to all objects, or you can specify basic filters, such as one or more tenant accounts or bucket names, or advanced filters, such as the object’s size or user metadata.

372

Name	<input type="text"/>
Description	<input type="text"/>
Tenant Accounts (optional)	<input type="text" value="Select tenant accounts or enter tenant IDs"/>
Bucket Name	<input type="text" value="matches all"/> <input type="button" value="Value"/>

[Advanced filtering...](#) (0 defined)

Cancel

Next

Steps

1. Enter a unique name for the rule in the **Name** field.

You must enter between 1 and 64 characters.

2. Optionally, enter a short description for the rule in the **Description** field.

You should describe the rule's purpose or function so you can recognize the rule later.

Name	<input type="text" value="Make 3 Copies"/>
Description	<input type="text" value="Save 1 copy at 3 sites for 1 year. Then, save EC copy forever"/>

3. Optionally, select one or more S3 or Swift tenant accounts to which this rule applies. If this rule applies to all tenants, leave this field blank.

If you do not have either the Root access permission or the Tenant accounts permission, you cannot select tenants from the list. Instead, enter the tenant ID or enter multiple IDs as a comma-delimited string.

4. Optionally, specify the S3 buckets or Swift containers to which this rule applies.

If **matches all** is selected (default), the rule applies to all S3 buckets or Swift containers.

5. Optionally, select **Advanced filtering** to specify additional filters.

If you do not configure advanced filtering, the rule applies to all objects that match the basic filters.

If this rule will create erasure-coded copies, add the **Object Size (MB)** advanced filter and set it to **greater than 1**. The size filter ensures that objects that are 1 MB or smaller will not be erasure coded.



Erasure coding is best suited for objects greater than 1 MB. Do not use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.

6. Select **Next**.

Step 2 (Define Placements) appears.

Related information

- [What an ILM rule is](#)
- [Use advanced filters in ILM rules](#)
- [Step 2 of 3: Define placements](#)

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 (microseconds)	<ul style="list-style-type: none">• equals• does not equal• less than• less than or equals• greater than• greater than or equals	<p>Time and date the object was ingested.</p> <p>Note: To avoid resource issues when activating a new ILM policy, you can use the Ingest Time advanced filter in any rule that might change the location of large numbers of existing objects. Set Ingest Time to be greater than or equal to the approximate time when the new policy will go into effect to ensure that existing objects are not moved unnecessarily.</p>
Key	<ul style="list-style-type: none">• equals• does not equal• contains• does not contain• starts with• does not start with• ends with• does not end with	<p>All or part of a unique S3 or Swift object key.</p> <p>For example, you might want to match objects that end with <code>.txt</code> or start with <code>test-object/</code>.</p>

Metadata type	Supported operators	Metadata value
Last Access Time (microseconds)	<ul style="list-style-type: none"> • equals • does not equal • less than • less than or equals • greater than • greater than or equals • exists • does not exist 	<p>Time and date the object was last retrieved (read or viewed).</p> <p>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.</p> <p>Use Last Access Time in ILM rules</p>
Location Constraint (S3 only)	<ul style="list-style-type: none"> • equals • does not equal 	<p>The region where an S3 bucket was created. Use ILM > Regions to define the regions that are shown.</p> <p>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.</p> <p>Configure regions (optional and S3 only)</p>
Object Size (MB)	<ul style="list-style-type: none"> • equals • not equals • less than • less than or equals • greater than • greater than or equals 	<p>The object's size in MB.</p> <p>Erasure coding is best suited for objects greater than 1 MB. Do not use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.</p> <p>Note: To filter on object sizes smaller than 1 MB, type in a decimal value. Your browser type and locale settings control whether you need to use a period or a comma as the decimal separator.</p>
User Metadata	<ul style="list-style-type: none"> • contains • ends with • equals • exists • does not contain • does not end with • does not equal • does not exist • does not start with • starts with 	<p>Key-value pair, where User Metadata Name is the key and User Metadata Value is the value.</p> <p>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 User Metadata Value.</p> <p>Note: User-metadata names are not case sensitive; user-metadata values are case sensitive.</p>

Metadata type	Supported operators	Metadata value
Object Tag (S3 only)	<ul style="list-style-type: none"> contains ends with equals exists does not contain does not end with does not equal does not exist does not start with starts with 	<p>Key-value pair, where Object Tag Name is the key and Object Tag Value is the value.</p> <p>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.</p> <p>Note: Object tag names and object tag values are case sensitive. You must enter these items exactly as they were defined for the object.</p>

Specifying multiple metadata types and values

When you define advanced filtering, you can specify multiple types of metadata and multiple metadata values. For example, if you want a rule to match objects between 10 MB and 100 MB in size, you would select the **Object Size** metadata type and specify two metadata values.

- The first metadata value specifies objects greater than or equal to 10 MB.
- The second metadata value specifies objects less than or equal to 100 MB.

Advanced Filtering

Use advanced filtering if you want a rule to apply only to specific objects. You can filter objects based on their system metadata, user metadata, or object tags (S3 only). When objects are evaluated, the rule is applied if the object's metadata matches the criteria in the advanced filter.

Objects between 10 and 100 MB

Matches all of the following metadata:

Object Size (MB)
greater than or equals
10
+
x

Object Size (MB)
less than or equals
100
+
x

+
x

Cancel

Remove Filters

Save

Using multiple entries allows you to have precise control over which objects are matched. In the following example, the rule applies to objects that have a Brand A or Brand B as the value of the camera_type user metadata. However, the rule only applies to those Brand B objects that are smaller than 10 MB.

Advanced Filtering

Use advanced filtering if you want a rule to apply only to specific objects. You can filter objects based on their system metadata, user metadata, or object tags (S3 only). When objects are evaluated, the rule is applied if the object's metadata matches the criteria in the advanced filter.

Multiple filters

Matches all of the following metadata:

User Metadata

▼

camera_type

equals

▼

Brand A

+

×

+

×

Or matches all of the following metadata:

User Metadata

▼

camera_type

equals

▼

Brand B

+

×

Object Size (MB)

▼

less than or equals

▼

10

+

×

+

×

Cancel

Remove Filters

Save

Step 2 of 3: Define placements

Step 2 (Define Placements) of the Create ILM Rule wizard allows you to define the placement instructions that determine how long objects are stored, the type of copies (replicated or erasure coded), the storage location, and the number of copies.

About this task

An ILM rule can include one or more placement instructions. Each placement instruction applies to a single period of time. When you use more than one instruction, the time periods must be contiguous, and at least one instruction must start on day 0. The instructions can continue either forever, or until you no longer require any object copies.

Each placement instruction can have multiple lines if you want to create different types of copies or use different locations during that time period.

This example ILM rule creates two replicated copies for the first year. Each copy is saved in a storage pool at a different site. After one year, a 2+1 erasure-coded copy is made and saved at only one site.

Configure placement instructions to specify how you want objects matched by this rule to be stored.

Example rule
Two copies for one year, then EC forever

Reference Time Ingest Time

Placements Sort by start day

From day 0 store for 365 days Add Remove

Type replicated Location DC1 DC2 Add Pool Copies 2 + -
Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

From day 365 store forever Add Remove

Type erasure coded Location DC1 (2 plus 1) Copies 1 + -

Retention Diagram Refresh

Trigger

Day 0

Year 1

DC1

DC2

DC1 (2 plus 1)

Duration

1 years

Forever

Cancel Back Next

Steps

1. For **Reference Time**, select the type of time to use when calculating the start time for a placement instruction.

Option	Description
Ingest Time	The time when the object was ingested.
Last Access Time	The time when the object was last retrieved (read or viewed). 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 .

Option	Description
Noncurrent Time	<p>The time an object version became noncurrent because a new version was ingested and replaced it as the current version.</p> <p>Note: Noncurrent Time applies only to S3 objects in versioning-enabled buckets.</p> <p>You can use this option to reduce the storage impact of versioned objects by filtering for noncurrent object versions. See Example 4: ILM rules and policy for S3 versioned objects.</p>
User Defined Creation Time	A time specified in user-defined metadata.



If you want to create a compliant rule, you must select **Ingest Time**.

- In the **Placements** section, select a starting time and a duration for the first time period.

For example, you might want to specify where to store objects for the first year (“day 0 for 365 days”). At least one instruction must start at day 0.

- If you want to create replicated copies:
 - From the **Type** drop-down list, select **replicated**.
 - In the **Location** field, select **Add Pool** for each storage pool you want to add.

If you specify only one storage pool, be aware that StorageGRID can store only one replicated copy of an object on any given Storage Node. If your grid includes three Storage Nodes and you select 4 as the number of copies, only three copies will be made—one copy for each Storage Node.



The **ILM placement unachievable** alert is triggered to indicate that the ILM rule could not be completely applied.

If you specify more than one storage pool, keep these rules in mind:

- The number of copies cannot be greater than the number of storage pools.
- If the number of copies equals the number of storage pools, one copy of the object is stored in each storage pool.
- If the number of copies is less than the number of storage pools, 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, do not specify the default All Storage Nodes storage pool and another storage pool.

Placements ⓘ ⇅ Sort by start day

From day store

Type Location Copies

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

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

A warning appears if you change the number of copies to 1. An ILM rule that creates only one replicated copy for any time period puts data at risk of permanent loss. See [Why you should not use single-copy replication](#).

Placements [?](#) Sort by start day

From day store Add Remove

Type Location Add Pool Copies Temporary location + x

An ILM rule that creates only one replicated copy for any time period puts data at risk of permanent loss. [View additional details](#).

To avoid these risks, do one or more of the following:

- Increase the number of copies for the time period.
- Select the plus sign icon **+** to create additional copies during the time period. Then, select a different storage pool or a Cloud Storage Pool.
- Select **erasure coded** for Type, instead of **replicated**. You can safely ignore this warning if this rule already creates multiple copies for all time periods.

- d. If you specified only one storage pool, ignore the **Temporary location** field.



Temporary locations are deprecated and will be removed in a future release. See [Use a storage pool as a temporary location \(deprecated\)](#).

4. If you want to create an erasure-coded copy:

- a. From the **Type** drop-down list, select **erasure coded**.

The number of copies changes to 1. A warning appears if the rule does not have an advanced filter to ignore objects that are 200 KB or smaller.

Erasure coding is best suited for objects greater than 1 MB. Do not use erasure coding for objects that are 200 KB or smaller. Select **Back** to return to Step 1. Then, use **Advanced filtering** to set the Object Size (MB) filter to any value greater than 0.2.



Erasure coding is best suited for objects greater than 1 MB. Do not use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.

- b. If the object size warning appeared, select **Back** to return to Step 1. Then, select **Advanced filtering** and set the Object Size (MB) filter to any value greater than 0.2.
- c. Select the storage location.

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

From day store Add Remove

Type Location Copies + ×

Erasure Coding profile name

Storage pool name

5. Optionally, add different time periods or create additional copies at different locations:

- Select the plus icon to create additional copies at a different location during the same time period.
- Select **Add** to add a different time period to the placement instructions.



Objects are automatically deleted at the end of the final time period unless the final time period ends with **forever**.

6. If you want to store objects in a Cloud Storage Pool:

- From the **Type** drop-down list, select **replicated**.
- In the **Location** field, select **Add Pool**. Then, select a Cloud Storage Pool.

From day store Add Remove

Type Location + ×

Copies

When using Cloud Storage Pools, keep these rules in mind:

- You cannot select more than one Cloud Storage Pool in a single placement instruction. Similarly, you cannot select a Cloud Storage Pool and a storage pool in the same placement instruction.

Type Location Copies


If you want to use a Cloud Storage Pool, you must remove any other storage pools or Cloud Storage Pools from this placement instruction.

- You can store only one copy of an object in any given Cloud Storage Pool. An error message appears if you set **Copies** to 2 or more.

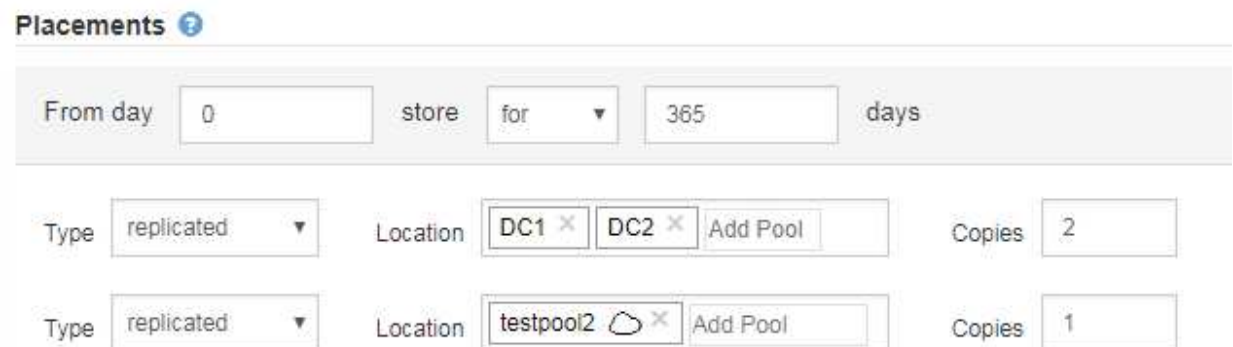
Type Location Copies

The number of copies cannot be more than one when a Cloud Storage Pool is selected.

- You cannot store more than one object copy in any Cloud Storage Pool at the same time. An error message appears if multiple placements that use a Cloud Storage Pool have overlapping dates or if multiple lines in the same placement use a Cloud Storage Pool.



- You can store an object in a Cloud Storage Pool at the same time that object is being stored as replicated or erasure coded copies in StorageGRID. However, as this example shows, you must include more than one line in the placement instruction for the time period, so you can specify the number and types of copies for each location.




7. Select **Refresh** to update the Retention Diagram and to confirm your placement instructions.

Each line in the diagram shows where and when object copies will be placed. The type of copy is represented by one of the following icons:

	Replicated copy
	Erasure-coded copy
	Cloud Storage Pool copy

In this example, two replicated copies will be saved to two storage pools (DC1 and DC2) for one year. Then, an erasure-coded copy will be saved for an additional 10 years, using a 6+3 erasure-coding scheme at three sites. After 11 years, the objects will be deleted from StorageGRID.



8. Select **Next**.

Step 3 (Define Ingest Behavior) appears.

Related information

- [What an ILM rule is](#)
- [Manage objects with S3 Object Lock](#)
- [Step 3 of 3: Define ingest behavior](#)

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	<ul style="list-style-type: none"> • No, for the source copy • Yes, for the destination copy 	<ul style="list-style-type: none"> • Yes, for the source copy • Yes, for the destination copy
Request to complete a multipart upload	Yes, for the assembled object	Yes, for the assembled object

Related information

- [Use S3](#)
- [Use a tenant account](#)

Step 3 of 3: Define ingest behavior

Step 3 (Define ingest behavior) of the Create ILM Rule wizard allows you to choose how the objects filtered by this rule are protected as they are ingested.

About this task

StorageGRID can make interim copies and queue the objects for ILM evaluation later, or it can make copies to meet the rule's placement instructions immediately.

Create ILM Rule
Step 3 of 3: Define ingest behavior

Select the data protection option to use when objects are ingested:

- ☐ Strict
Always uses this rule's placements on ingest. Ingest fails when this rule's placements are not possible.
- ☒ Balanced
Optimum ILM efficiency. Attempts this rule's placements on ingest. Creates interim copies when that is not possible.
- ☐ Dual commit
Creates interim copies on ingest and applies this rule's placements later.

Cancel
Back
Save

Steps

1. Select the data protection option to use when objects are ingested:

Option	Description
Strict	Always uses this rule's placements on ingest. Ingest fails when this rule's placements are not possible.

Option	Description
Balanced	Optimum ILM efficiency. Attempts this rule's placements on ingest. Creates interim copies when that is not possible.
Dual commit	Creates interim copies on ingest and applies this rule's placements later.

Balanced offers a combination of data security and efficiency that is suitable in most cases. Strict or Dual commit are generally used to meet specific requirements.

See [Data-protection options for ingest](#) and [Advantages, disadvantages, and limitations of the data-protection options](#) for more information.



An error message appears if you select the Strict or Balanced option and the rule uses one of these placements:

- A Cloud Storage Pool at day 0
- An Archive Node at day 0
- A Cloud Storage Pool or an Archive Node when the rule uses a User Defined Creation Time as a Reference Time

2. Select **Save**.

The ILM rule is saved. The rule does not become active until it is added to an ILM policy and that policy is activated.

Related information

- [Example 5: ILM rules and policy for Strict ingest behavior](#)
- [Create an ILM policy](#)

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 cannot use any filters. It must apply to all tenants, all buckets, and all object versions.

What you'll need

- 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 cannot use any filters or the Noncurrent reference time. The placement instructions for the default rule are applied to any objects that are not matched by another rule in the policy.

In this example policy, the first rule applies only to objects belonging to Tenant A. The default rule, which is last, applies to objects belonging to all other tenant accounts.

Configure ILM Policy

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

Name

Reason for change

Rules

1. Select the rules you want to add to the policy.
 2. Determine the order in which the rules will be evaluated by dragging and dropping the rows. The default rule will be automatically placed at the end of the policy and cannot be moved.

[+ Select Rules](#)

Default	Rule Name	Tenant Account	Actions
	EC for Tenant A	Tenant A (91643888913299990564)	
<input checked="" type="checkbox"/>	2 copies 2 sites	—	

[Cancel](#) [Save](#)

When you create the default rule, keep these requirements in mind:

- The default rule is automatically placed as the last rule in the policy.
- The default rule cannot use any basic or advanced filters.
- The default rule must apply to all object versions, so it cannot use the Noncurrent Time reference time.
- The default rule should create replicated copies.



Do not use a rule that creates erasure-coded copies as the default rule for a policy. Erasure-coding rules should use an advanced filter to prevent smaller objects from being erasure coded.

- In general, the default rule should retain objects forever.
- If you are using (or you plan to enable) the global S3 Object Lock setting, the default rule for the active or proposed policy must be compliant.

Steps

1. Select **ILM > Rules**.

The ILM Rules page appears.

2. Select **Create**.

Step 1 (Define Basics) of the Create ILM Rule wizard appears.

3. Enter a unique name for the rule in the **Name** field.
4. Optionally, enter a short description for the rule in the **Description** field.
5. Leave the **Tenant Accounts** field blank.

The default rule must apply to all tenant accounts.

6. Leave the **Bucket Name** field blank.

The default rule must apply to all S3 buckets and Swift containers.

7. Do not select **Advanced filtering**

The default rule cannot specify any filters.

8. Select **Next**.

Step 2 (Define Placements) appears.

9. For Reference Time, select any option except **Noncurrent Time**.

The default rule must apply all object versions.

10. Specify the placement instructions for the default rule.

- The default rule should retain objects forever. A warning appears when you activate a new policy if the default rule does not retain objects forever. You must confirm this is the behavior you expect.
- The default rule should create replicated copies.



Do not use a rule that creates erasure-coded copies as the default rule for a policy. Erasure-coding rules should include the **Object Size (MB) greater than 0.2** advanced filter to prevent smaller objects from being erasure coded.

- If you are using (or you plan to enable) the global S3 Object Lock setting, the default rule must be compliant:
 - It must create at least two replicated object copies or one erasure-coded copy.
 - These copies must exist on Storage Nodes for the entire duration of each line in the placement instructions.
 - Object copies cannot be saved in a Cloud Storage Pool.
 - Object copies cannot be saved on Archive Nodes.
 - At least one line of the placement instructions must start at day 0, using Ingest Time as the reference time.
 - At least one line of the placement instructions must be “forever.”

11. Select **Refresh** to update the Retention Diagram and to confirm your placement instructions.

12. Select **Next**.

Step 3 (Define Ingest Behavior) appears.

13. Select the data protection option to use when objects are ingested, and select **Save**.

Create ILM policy

Create ILM policy: Overview

When you create an ILM policy, you start by selecting and arranging the ILM rules. Then, you verify the behavior of your proposed policy by simulating it against previously ingested objects. When you are satisfied that the proposed policy is functioning as

intended, you can activate it to create the active policy.



An ILM policy that has been incorrectly configured can result in unrecoverable data loss. Before activating an ILM policy, carefully review the ILM policy and its ILM rules, and then simulate the ILM policy. Always confirm that the ILM policy will work as intended.

Considerations for creating an ILM policy

- Use the system's built-in policy, Baseline 2 Copies Policy, in test systems only. The Make 2 Copies rule in this policy uses the All Storage Nodes storage pool, which contains all sites. If your StorageGRID system has more than one site, two copies of an object might be placed on the same site.
- When designing a new policy, consider all of the different types of objects that might be ingested into your grid. Make sure the policy includes rules to match and place these objects as required.
- Keep the ILM policy as simple as possible. This avoids potentially dangerous situations where object data is not protected as intended when changes are made to the StorageGRID system over time.
- Make sure that the rules in the policy are in the correct order. When the policy is activated, new and existing objects are evaluated by the rules in the order listed, starting at the top. For example, if the first rule in a policy matches an object, that rule will not be evaluated by any other rule.
- The last rule in every ILM policy is the default ILM rule, which cannot use any filters. If an object has not been matched by another rule, the default rule controls where that object is placed and for how long it is retained.
- Before activating a new policy, review any changes that the policy is making to the placement of existing objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

Create a proposed ILM policy

You can create a proposed ILM policy from scratch, or you can clone the current active policy if you want to start with the same set of rules.



If the global S3 Object Lock setting has been enabled, use this procedure instead: [Create an ILM policy after S3 Object Lock is enabled](#).

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.
- You have created the ILM rules you want to add to the proposed policy. As required, you can save a proposed policy, create additional rules, and then edit the proposed policy to add the new rules.
- You have [created a default ILM rule](#) for the policy that does not contain any filters.
- Optionally, you have watched the video: [Video: StorageGRID ILM Policies](#)



About this task

Typical reasons for creating a proposed ILM policy include:

- You added a new site and need to use new ILM rules to place objects at that site.
- You are decommissioning a site and you need to remove all rules that refer to the site.
- You added a new tenant that has special data protection requirements.
- You started to use a Cloud Storage Pool.



Use the system's built-in policy, Baseline 2 Copies Policy, in test systems only. The Make 2 Copies rule in this policy uses the All Storage Nodes storage pool, which contains all sites. If your StorageGRID system has more than one site, two copies of an object might be placed on the same site.

Steps

1. Select **ILM > Policies**.

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

ILM Policies

Review the proposed, active, and historical policies. You can create, edit, or delete a proposed policy; clone the active policy; or view the details for any policy.

+ Create Proposed Policy
Clone
Edit
Remove

Policy Name	Policy State	Start Date	End Date
Baseline 2 Copies Policy	Active	2017-07-17 12:00:45 MDT	

Viewing Active Policy - Baseline 2 Copies Policy

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

Rules are evaluated in order, starting from the top.

Rule Name	Default	Tenant Account
Make 2 Copies	✓	Ignore

Simulate
Activate

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

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

The Configure ILM Policy dialog box appears.

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

Configure ILM Policy

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

Name

Reason for change

Rules

1. Select the rules you want to add to the policy.

2. Determine the order in which the rules will be evaluated by dragging and dropping the rows. The default rule will be automatically placed at the end of the policy and cannot be moved.

+ Select Rules

Default	Rule Name	Tenant Account	Actions
No rules selected.			

Cancel

Save

If you are cloning the active policy, the **Name** field shows the name of the active policy, appended by a version number ("v2" in the example). The rules used in the active policy are selected and shown in their current order.

Name	Baseline 2 Copies Policy (v2)
Reason for change	

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

You must enter at least 1 and no more than 64 characters. If you are cloning the active policy, you can use the current name with the appended version number or you can enter a new name.

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

You must enter at least 1 and no more than 128 characters.

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

The Select Rules for Policy dialog box appears, with all defined rules listed. If you are cloning a policy:

- The rules used by the policy you are cloning are selected.
- If the policy you are cloning used any rules with no filters that were not the default rule, you are prompted to remove all but one of those rules.
- If the default rule used a filter or the Noncurrent reference time, you are prompted to select a new default rule.
- If the default rule was not the last rule, a button allows you to move the rule to the end of the new policy.

Select Rules for Policy

Select Default Rule

This list shows the rules that do not use any filters. Select one rule to be the default rule for the policy. The default rule applies to any objects that do not match another rule in the policy and is always evaluated last. The default rule should retain objects forever.

Rule Name
<input checked="" type="radio"/> 2 copies 2 sites
<input type="radio"/> Make 2 Copies

Select Other Rules

The other rules in a policy are evaluated before the default rule and must use at least one filter. Each rule in this list uses at least one filter (tenant account, bucket name, advanced filter, or the noncurrent reference time).

Rule Name	Tenant Account
<input type="checkbox"/> EC for Tenant A	Tenant A (91643888913299990564)
<input type="checkbox"/> 2 copies 2 sites noncurrent time	—

Cancel

Apply

6. Select a rule name or the more details icon to view the settings for that rule.

This example shows the details of an ILM rule that makes two replicated copies at two sites.

Two-Site Replication for Other Tenants


Description: Two-Site Replication for Other Tenants

Ingest Behavior: Balanced

Reference Time: Ingest Time

Filtering Criteria: Matches all objects.

Retention Diagram:



Close

7. In the **Select Default Rule** section, select one default rule for the proposed policy.

The default rule applies to any objects that do not match another rule in the policy. The default rule cannot use any filters and is always evaluated last.



If no rule is listed in the Select Default Rule section, you must exit the ILM policy page and [create a default ILM rule](#).



Do not use the Make 2 Copies stock rule as the default rule for a policy. The Make 2 Copies rule uses a single storage pool, All Storage Nodes, which contains all sites. If your StorageGRID system has more than one site, two copies of an object might be placed on the same site.

8. In the **Select Other Rules** section, select any other rules you want to include in the policy.

The other rules are evaluated before the default rule and must use at least one filter (tenant account, bucket name, advanced filter, or the Noncurrent reference time).

9. When you are done selecting rules, select **Apply**.

The rules you selected are listed. The default rule is at the end, with the other rules above it.

Rules

1. Select the rules you want to add to the policy.
2. Determine the order in which the rules will be evaluated by dragging and dropping the rows. The default rule will be automatically placed at the end of the policy and cannot be moved.

+ Select Rules

	Default	Rule Name	Tenant Account	Actions
⬆⬇⬆⬆		3-site EC	Ignore	✕
⬆⬇⬆⬆		1-site EC	Ignore	✕
	✓	2 copies at 2 data centers	Ignore	✕

Cancel

Save

A warning appears if the default rule does not retain objects forever. When you activate this policy, you must confirm that you want StorageGRID to delete objects when the placement instructions for the default rule elapse (unless a bucket lifecycle keeps the objects for longer).



	Default	Rule Name	Tenant Account	Actions
⬆⬇⬆⬆		3-site EC	Ignore	✕
⬆⬇⬆⬆		1-site EC	Ignore	✕
	✓	2 copies at 2 data centers for 2 years	Ignore	✕

The default ILM rule in this policy does not retain objects forever. Confirm this is the behavior you expect. Otherwise, any objects that are not matched by another rule will be deleted after 720 days.

10. Drag and drop the rows for the non-default rules to determine the order in which these rules will be evaluated.

You cannot move the default rule.



You must confirm that the ILM rules are in the correct order. When the policy is activated, new and existing objects are evaluated by the rules in the order listed, starting at the top.

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

The ILM Policies page is updated:

- The policy you saved is shown as Proposed. Proposed policies do not have start and end dates.
- The **Simulate** and **Activate** buttons are enabled.

ILM Policies

Review the proposed, active, and historical policies. You can create, edit, or delete a proposed policy; clone the active policy; or view the details for any policy.

[+ Create Proposed Policy](#) [Clone](#) [Edit](#) [Remove](#)

Policy Name	Policy State	Start Date	End Date
• Data Protection for Three Sites	Proposed		
• Data Protection for Two Sites	Active	2020-09-18 16:01:24 MDT	
• Baseline 2 Copies Policy	Historical	2020-09-17 21:32:57 MDT	2020-09-18 16:01:24 MDT

Viewing Proposed Policy - Data Protection for Three Sites

Before activating a new ILM policy:

- Review and carefully simulate the policy. Errors in an ILM policy can cause irreparable data loss.
- Review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

See [Managing objects with information lifecycle management](#) for more information.

This policy contains a rule that makes an erasure-coded copy. Confirm that at least one rule uses the Object Size advanced filter to prevent objects that are 200 KB or smaller from being erasure coded. See [Managing objects with information lifecycle management](#) for more information.

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

Reason for change: Added a third site

Rules are evaluated in order, starting from the top.

Rule Name	Default	Tenant Account
One-Site Erasure Coding for Tenant A 🔗		Tenant A (20033011709864740158)
Three-Site Replication for Other Tenants 🔗	✓	Ignore

[Simulate](#) [Activate](#)

13. Go to [Simulate an ILM policy](#).

Related information

- [What an ILM policy is](#)
- [Manage objects with S3 Object Lock](#)

Create an ILM policy after S3 Object Lock is enabled

If the global S3 Object Lock setting is enabled, the steps for creating a policy are slightly different. You must ensure that the ILM policy is compliant with the requirements of buckets that have S3 Object Lock enabled.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.
- The global S3 Object Lock setting is already enabled for the StorageGRID system.



If the global S3 Object Lock setting has not been enabled, use the general instructions for [creating a proposed ILM policy](#).

- You have created the compliant and non-compliant ILM rules you want to add to the proposed policy. As required, you can save a proposed policy, create additional rules, and then edit the proposed policy to add the new rules. See [Example 7: Compliant ILM policy for S3 Object Lock](#).

- You have [created a default ILM rule](#) for the policy that is compliant.
- Optionally, you have watched the video: [Video: StorageGRID ILM Policies](#)



Steps

1. Select **ILM > Policies**.

The ILM Policies page appears. If the global S3 Object Lock setting is enabled, the ILM Policies page indicates which ILM rules are compliant.

ILM Policies

Review the proposed, active, and historical policies. You can create, edit, or delete a proposed policy; clone the active policy; or view the details for any policy.

+ Create Proposed Policy
Clone
Edit
Remove

Policy Name	Policy State	Start Date	End Date
Baseline 2 Copies Policy	Active	2021-02-04 01:04:29 MST	

Viewing Active Policy - Baseline 2 Copies Policy

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

Rules are evaluated in order, starting from the top. The policy's default rule must be compliant.

Rule Name	Default	Compliant	Tenant Account
Make 2 Copies	✓	✓	Ignore

Simulate
Activate

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

You must enter at least 1 and no more than 64 characters.

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

You must enter at least 1 and no more than 128 characters.

4. To add rules to the policy, select **Select Rules**.

The Select Rules for Policy dialog box appears, with all defined rules listed.

- The Select Default Rule section lists the rules that can be the default for a compliant policy. It includes compliant rules that do not use filters or the Noncurrent reference time.
- The Select Other Rules section lists the other compliant and non-compliant rules that can be selected for this policy.

Select Rules for Policy

Select Default Rule

This list shows the rules that are compliant and do not use any filters. Select one rule to be the default rule for the policy. The default rule applies to any objects that do not match another rule in the policy and is always evaluated last.

	Rule Name
<input type="radio"/>	Default Compliant Rule: Two Copies Two Data Centers
<input type="radio"/>	Make 2 Copies

Select Other Rules

The other rules in a policy are evaluated before the default rule and must use at least one filter. Each rule in this list uses at least one filter (tenant account, bucket name, advanced filter, or the noncurrent reference time).

	Rule Name	Compliant	Uses Filter	Is Selectable
<input type="checkbox"/>	Compliant Rule: EC for bank-records bucket - Bank of AB C	✓	✓	Yes
<input type="checkbox"/>	Non-Compliant Rule: Use Cloud Storage Pool			Yes

Cancel
Apply

5. Select a rule name or the more details icon to view the settings for that rule.
6. In the **Select Default Rule** section, select one default rule for the proposed policy.

The table in this section only lists the rules that are compliant and do not use any filters.



If no rule is listed in the Select Default Rule section, you must exit the ILM policy page and [create a default ILM rule](#) that is compliant.



Do not use the Make 2 Copies stock rule as the default rule for a policy. The Make 2 Copies rule uses a single storage pool, All Storage Nodes, which contains all sites. If you use this rule, multiple copies of an object might be placed on the same site.

7. In the **Select Other Rules** section, select any other rules you want to include in the policy.
 - a. If you need a different “default” rule for objects in non-compliant S3 buckets, optionally select one non-compliant rule that does not use a filter.

For example, you might want to use a Cloud Storage Pool or an Archive Node to store objects in buckets that do not have S3 Object Lock enabled.



You can only select one non-compliant rule that does not use a filter. As soon as you select one rule, the **Is Selectable** column shows **No** for any other non-compliant rules without filters.

- b. Select any other compliant or non-compliant rules you want to use in the policy.

The other rules must use at least one filter (tenant account, bucket name, or an advanced filter, such as object size).

- When you are done selecting the rules, select **Apply**.

The rules you selected are listed. The default rule is at the end, with the other rules above it. If you also selected a non-compliant “default” rule, that rule is added as the second-to-last rule in the policy.

In this example, the last rule, 2 Copies 2 Data Centers, is the default rule: it is compliant and has no filters. The second-to-last rule, Cloud Storage Pool, also has no filters but it is not compliant.

Configure ILM Policy

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

Name

Reason for change

Rules

- Select the rules you want to add to the policy.
- Determine the order in which the rules will be evaluated by dragging and dropping the rows. The default rule (and any non-compliant rule without a filter) will be automatically placed at the end of the policy and cannot be moved.

+ Select Rules				
Default	Rule Name	Compliant	Tenant Account	Actions
	Compliant Rule: EC for bank-records bucket - Bank of ABC	✓	Bank of ABC (90767802913525281639)	✕
	Non-Compliant Rule: Use Cloud Storage Pool		Ignore	✕
✓	Default Compliant Rule: Two Copies Two Data Centers	✓	Ignore	✕

Cancel

Save

- Drag and drop the rows for the non-default rules to determine the order in which these rules will be evaluated.

You cannot move the default rule or the non-compliant “default” rule.



You must confirm that the ILM rules are in the correct order. When the policy is activated, new and existing objects are evaluated by the rules in the order listed, starting at the top.

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

The ILM Policies page is updated:

- The policy you saved is shown as Proposed. Proposed policies do not have start and end dates.
- The **Simulate** and **Activate** buttons are enabled.

ILM Policies

Review the proposed, active, and historical policies. You can create, edit, or delete a proposed policy; clone the active policy; or view the details for any policy.

[+ Create Proposed Policy](#) [Clone](#) [Edit](#) [Remove](#)

Policy Name	Policy State	Start Date	End Date
Compliant ILM Policy for S3 Object Lock	Proposed		
Compliant ILM Policy	Active	2021-02-05 16:22:53 MST	
Non-Compliant ILM policy	Historical	2021-02-05 15:17:05 MST	2021-02-05 16:22:53 MST
Baseline 2 Copies Policy	Historical	2021-02-04 21:35:52 MST	2021-02-05 15:17:05 MST

Viewing Proposed Policy - Compliant ILM Policy for S3 Object Lock

Before activating a new ILM policy:

- Review and carefully simulate the policy. Errors in an ILM policy can cause irreparable data loss.
- Review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

See [Managing objects with information lifecycle management](#) for more information.

This policy contains a rule that makes an erasure-coded copy. Confirm that at least one rule uses the Object Size advanced filter to prevent objects that are 200 KB or smaller from being erasure coded. See [Managing objects with information lifecycle management](#) for more information.

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

Reason for change: [Example policy](#)

Rules are evaluated in order, starting from the top. The policy's default rule must be compliant.

Rule Name	Default	Compliant	Tenant Account
Compliant Rule: EC for bank-records bucket - Bank of ABC 🔗		✓	Bank of ABC (90767802913525281639)
Non-Compliant Rule: Use Cloud Storage Pool 🔗			Ignore
Default Compliant Rule: Two Copies Two Data Centers 🔗	✓	✓	Ignore

[Simulate](#) [Activate](#)

12. Go to [Simulate an ILM policy](#).

Simulate an ILM policy

You should simulate a proposed policy on test objects before activating the policy and applying it to your production data. The simulation window provides a standalone environment that is safe for testing policies before they are activated and applied to data in the production environment.

What you'll need


- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.
- You know the S3 bucket/object-key or the Swift container/object-name for each object you want to test, and you have already ingested those objects.

About this task

You must carefully select the objects you want the proposed policy to test. To simulate a policy thoroughly, you should test at least one object for each filter in each rule.

For example, if a policy includes one rule to match objects in bucket A and another rule to match objects in bucket B, you must select at least one object from bucket A and one object from bucket B to test the policy thoroughly. You must also select at least one object from another bucket to test the default rule.

When simulating a policy, the following considerations apply:

- After you make changes to a policy, save the proposed policy. Then, simulate the behavior of the saved proposed policy.
- When you simulate a policy, the ILM rules in the policy filter the test objects, so you can see which rule was applied to each object. However, no object copies are made and no objects are placed. Running a simulation does not modify your data, rules, or the policy in any way.
- The Simulation page retains the objects you tested until you close, navigate away from, or refresh the ILM Policies page.
- Simulation returns the name of the matched rule. To determine which storage pool or Erasure Coding profile is in effect, you can view the Retention Diagram by selecting the rule name or the more details icon .
- If S3 Versioning is enabled, the policy is only simulated against the current version of the object.

Steps

1. Select and arrange the rules, and save the proposed policy.

The policy in this example has three rules:

Rule Name	Filter	Type of Copies	Retention
X-men	<ul style="list-style-type: none"> • Tenant A • User metadata (series=x-men) 	2 copies at two data centers	2 years
PNGs	Key ends with .png	2 copies at two data centers	5 years
Two Copies Two Data Centers	<i>None</i>	2 copies at two data centers	Forever

Viewing Proposed Policy - Example ILM policy

Before activating a new ILM policy:


- Review and carefully simulate the policy. Errors in an ILM policy can cause irreparable data loss.
- Review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

See [Managing objects with information lifecycle management](#) for more information.

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

Reason for change: Example policy

Rules are evaluated in order, starting from the top.

Rule Name	Default	Tenant Account
X-men 		Tenant A (94793396288150002349)
PNGs 		Ignore
Two Copies at Two Data Centers 	✓	Ignore


Simulate
Activate

2. Using an S3 or Swift client or the [experimental S3 Console](#), which is available in Tenant Manager for each tenant, ingest the objects required to test each rule.
3. Select **Simulate**.

The Simulation ILM Policy dialog box appears.

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

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



Object

photos/test

Object 'photos/test' not found.

Simulate

5. Under **Simulation Results**, confirm that each object was matched by the correct rule.

In the example, the `Havok.png` and `Warpath.jpg` objects were correctly matched by the X-men rule. The `Fullsteam.png` object, which does not include `series=x-men` user metadata, was not matched by the X-men rule but was correctly matched by the PNGs rule. The default rule was not used because all three objects were matched by other rules.

Simulate ILM Policy - Demo




Simulates the active ILM policy or, if there is a proposed ILM policy, simulates the proposed ILM policy. Use this simulation to test the current configuration of ILM rules and determine whether ILM rules copy and place object data as intended.

Object

my-bucket/my-object-name or my-container/my-object-name

Simulate

Simulation Results ?

Object	Rule Matched	Previous Match	
photos/Havok.png	X-men 		✗
photos/Warpath.jpg	X-men 		✗
photos/Fullsteam.png	PNGs 		✗

Finish

Example 1: Verify rules when simulating a proposed ILM policy

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

In this example, the **Example ILM policy** is being simulated against the ingested objects in two buckets. The policy includes three rules, as follows:

- The first rule, **Two copies, two years for bucket-a**, applies only to objects in bucket-a.
- The second rule, **EC objects > 1 MB**, applies to all buckets but filters on objects greater than 1 MB.
- The third rule, **Two copies, two data centers**, is the default rule. It does not include any filters and does not use the Noncurrent reference time.

Before activating a new ILM policy:

- Review and carefully simulate the policy. Errors in an ILM policy can cause irreparable data loss.
- Review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

See the [instructions for managing objects with ILM](#) for more information.

This policy contains a rule that makes an erasure-coded copy. Confirm that at least one rule uses the Object Size advanced filter to prevent objects that are 200 KB or smaller from being erasure coded. Using EC is best suited for objects greater than 1 MB. See the [instructions for managing objects with ILM](#) for more information.

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

Reason for change:

Example policy

Rules are evaluated in order, starting from the top.

Rule Name	Default	Tenant Account
Two copies, two years for bucket-a 		—
EC objects > 1 MB 		—
Two copies, two data centers 	✓	—

Simulate

Activate

Steps

1. After adding the rules and saving the policy, select **Simulate**.

The Simulate ILM Policy dialog box appears.

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

The Simulation Results appear, showing which rule in the policy matched each object you tested.

Simulate ILM Policy - Example ILM policy

Simulates the active ILM policy or, if there is a proposed ILM policy, simulates the proposed ILM policy. Use this simulation to test the current configuration of ILM rules and determine whether ILM rules copy and place object data as intended.

Object

[Simulate](#)

Simulation Results

Object	Rule Matched	Previous Match	
bucket-a/bucket-a object.pdf	Two copies, two years for bucket-a 		✗
bucket-b/test object greater than 1 MB.pdf	EC objects > 1 MB 		✗
bucket-b/test object less than 1 MB.pdf	Two copies, two data centers 		✗

[Finish](#)

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

In this example:

- bucket-a/bucket-a object.pdf correctly matched the first rule, which filters on objects in bucket-a.
- bucket-b/test object greater than 1 MB.pdf is in bucket-b, so it did not match the first rule. Instead, it was correctly matched by the second rule, which filters on objects greater than 1 MB.
- bucket-b/test object less than 1 MB.pdf did not match the filters in the first two rules, so it will be placed by the default rule, which includes no filters.

Example 2: Reorder rules when simulating a proposed ILM policy

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

In this example, the **Demo** policy is being simulated. This policy, which is intended to find objects that have series=x-men user metadata, includes three rules, as follows:

- The first rule, **PNGs**, filters for key names that end in .png.
- The second rule, **X-men**, applies only to objects for Tenant A and filters for series=x-men user metadata.
- The last rule, **Two copies two data centers**, is the default rule, which matches any objects that do not match the first two rules.

Viewing Proposed Policy - Demo

Before activating a new ILM policy:

- Review and carefully simulate the policy. Errors in an ILM policy can cause irreparable data loss.
- Review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

See [Managing objects with information lifecycle management](#) for more information.

Review the rules in this policy. If this is a proposed policy, click **Simulate** to verify the policy and then click **Activate** to make the policy active.

Reason for change: new policy

Rules are evaluated in order, starting from the top.

Rule Name	Default	Tenant Account
PNGs 		Ignore
X-men 		Tenant A (24365814597594524591)
Two copies two data centers 	✓	Ignore

Simulate **Activate**

Steps

1. After adding the rules and saving the policy, 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.

Simulate ILM Policy - Demo

Simulates the active ILM policy or, if there is a proposed ILM policy, simulates the proposed ILM policy. Use this simulation to test the current configuration of ILM rules and determine whether ILM rules copy and place object data as intended.

Object **Simulate**

Simulation Results

Object	Rule Matched	Previous Match	
photos/Havok.png	PNGs 		✗

Finish

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

3. To resolve the issue, reorder the rules.
 - a. Select **Finish** to close the Simulate ILM Policy page.
 - b. Select **Edit** to edit the policy.
 - c. Drag the **X-men** rule to the top of the list.

Configure ILM Policy

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

Name Demo

Reason for change Reordering rules when simulating a proposed ILM policy

Rules

1. Select the rules you want to add to the policy.
2. Determine the order in which the rules will be evaluated by dragging and dropping the rows. The default rule will be automatically placed at the end of the policy and cannot be moved.

+ Select Rules

	Default	Rule Name	Tenant Account	Actions
+		X-men	Tenant A (48713995194927812566)	✕
+		PNGs	—	✕
	✓	Two copies, two data centers	—	✕

Cancel

Save

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.

Simulate ILM Policy - Demo

Simulates the active ILM policy or, if there is a proposed ILM policy, simulates the proposed ILM policy. Use this simulation to test the current configuration of ILM rules and determine whether ILM rules copy and place object data as intended.

Object my-bucket/my-object-name or my-container/my-object-name

Simulate

Simulation Results

Object	Rule Matched	Previous Match	
photos/Havok.png	X-men	PNGs	✕

Finish



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

Example 3: Correct a rule when simulating a proposed ILM policy

This example shows how to simulate a policy, correct a rule in the policy, and continue the simulation.

In this example, the **Demo** policy is being simulated. This policy is intended to find objects that have `series=x-men` user metadata. However, unexpected results occurred when simulating this policy against the `Beast.jpg` object. Instead of matching the X-men metadata rule, the object matched the default rule, Two

copies two data centers.

Simulate ILM Policy - Demo

Simulates the active ILM policy or, if there is a proposed ILM policy, simulates the proposed ILM policy. Use this simulation to test the current configuration of ILM rules and determine whether ILM rules copy and place object data as intended.

Object

my-bucket/my-object-name or my-container/my-object-name

Simulate

Simulation Results

Object	Rule Matched	Previous Match	
photos/Beast.jpg	Two copies two data centers		✖

Finish

When a test object is not matched by the expected rule in the policy, you must examine each rule in the policy and correct any errors.

Steps

- 1. For each rule in the policy, view the rule settings by selecting the rule name or the more details icon on any dialog box where the rule is displayed.
- 2. Review the rule’s tenant account, reference time, and filtering criteria.

In this example, the metadata for the X-men rule includes an error. The metadata value was entered as “x-men1” instead of “x-men.”

X-men

Ingest Behavior:Balanced

Tenant Account:06846027571548027538

Reference Time:Ingest Time

Filtering Criteria:

Matches all of the following metadata:

User Metadata

series

equals

x-men1

Retention Diagram:

Trigger

Day 0


All Storage Nodes

DurationForever

Close

- 3. To resolve the error, correct the rule, as follows:

- If the rule is part of the proposed policy, you can either clone the rule or remove the rule from the policy and then edit it.
- If the rule is part of the active policy, you must clone the rule. You cannot edit or remove a rule from the active policy.

Option	Description
Clone the rule	<ol style="list-style-type: none"> Select ILM > Rules. Select the incorrect rule, and select Clone. Change the incorrect information, and select Save. Select ILM > Policies. Select the proposed policy, and select Edit. Select Select Rules. Select the check box for the new rule, uncheck the check box for the original rule, and select Apply. Select Save.
Edit the rule	<ol style="list-style-type: none"> Select the proposed policy, and select Edit. Select the delete icon  to remove the incorrect rule, and select Save. Select ILM > Rules. Select the incorrect rule, and select Edit. Change the incorrect information, and select Save. Select ILM > Policies. Select the proposed policy, and select Edit. Select the corrected rule, select Apply, and select Save.

4. Perform the simulation again.



Because you navigated away from the ILM Policies page to edit the rule, the objects you previously entered for simulation are no longer displayed. You must re-enter the names of the objects.


In this example, the corrected X-men rule now matches the `Beast.jpg` object based on the `series=x-men` user metadata, as expected.

Simulate ILM Policy - Demo

Simulates the active ILM policy or, if there is a proposed ILM policy, simulates the proposed ILM policy. Use this simulation to test the current configuration of ILM rules and determine whether ILM rules copy and place object data as intended.

Object

Simulation Results

Object	Rule Matched	Previous Match	
photos/Beast.jpg	X-men 		

Activate the ILM policy

After you add ILM rules to a proposed ILM policy, simulate the policy, and confirm it behaves as you expect, you are ready to activate the proposed policy.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.
- You have saved and simulated the proposed ILM policy.



Errors in an ILM policy can cause unrecoverable data loss. Carefully review and simulate the policy before activating it to confirm that it will work as intended.



When you activate a new ILM policy, StorageGRID uses it to manage all objects, including existing objects and newly ingested objects. Before activating a new ILM policy, review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

About this task

When you activate an ILM policy, the system distributes the new policy to all nodes. However, the new active policy might not actually take effect until all grid nodes are available to receive the new policy. In some cases, the system waits to implement a new active policy to ensure that grid objects are not accidentally removed.

- If you make policy changes that increase data redundancy or durability, those changes are implemented immediately. For example, if you activate a new policy that includes a three-copies rule instead of a two-copies rule, that policy will be implemented right away because it increases data redundancy.
- If you make policy changes that could decrease data redundancy or durability, those changes will not be implemented until all grid nodes are available. For example, if you activate a new policy that uses a two-copies rule instead of a three-copies rule, the new policy will be marked as "Active," but it will not take effect until all nodes are online and available.

Steps

1. When you are ready to activate a proposed policy, select the policy on the ILM Policies page and select **Activate**.

A warning message is displayed, prompting you to confirm that you want to activate the proposed policy.

⚠ Warning

Activate the proposed policy

Errors in an ILM policy can cause irreparable data loss. Review and test the policy carefully before activating. Are you sure you want to activate the proposed policy?

Cancel

OK

A prompt appears in the warning message if the default rule for the policy does not retain objects forever. In this example, the retention diagram shows that the default rule will delete objects after 2 years. You must type **2** in the text box to acknowledge that any objects not matched by another rule in the policy will be removed from StorageGRID after 2 years.

⚠ Activate the proposed policy

Errors in an ILM policy can cause irreparable data loss. Review and test the policy carefully before activating.

The default rule in this policy does not retain objects forever. Confirm this is the behavior you want by referring to the retention diagram for the default rule:



Now, complete the following prompt:

Any objects that are not matched by another rule in this policy will be deleted after years.

Are you sure you want to activate the proposed policy?

Cancel

OK

2. Select **OK**.

Result

When a new ILM policy has been activated:

- The policy is shown with a Policy State of Active in the table on the ILM Policies page. The Start Date entry indicates the date and time the policy was activated.

ILM Policies

Review the proposed, active, and historical policies. You can create, edit, or delete a proposed policy; clone the active policy; or view the details for any policy.

+ Create Proposed Policy Clone Edit Remove			
Policy Name	Policy State	Start Date	End Date
<input checked="" type="radio"/> New Policy	Active	2017-07-20 18:49:53 MDT	
<input type="radio"/> Baseline 2 Copies Policy	Historical	2017-07-19 21:24:30 MDT	2017-07-20 18:49:53 MDT

- The previously active policy is shown with a Policy State of Historical. The Start Date and End Date entries

indicate when the policy became active and when it was no longer in effect.

Related information

Example 6: Changing an ILM policy

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.

What you'll need

- 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).

Object Metadata Lookup

Enter the identifier for any object stored in the grid to view its metadata.

Identifier	<input type="text" value="source/testobject"/>
Version ID (optional)	<input type="text" value="MEJGMkMyQzgtNEY5OC0xMUU3LTkzMEYtRDkyNTAwQkY5I"/>

Look Up

5. Select **Look Up**.

The object metadata lookup results appear. This page lists the following types of information:

- System metadata, including the object ID (UUID), the object name, the name of the container, the tenant account name or ID, the logical size of the object, the date and time the object was first created, and the date and time the object was last modified.
- Any custom user metadata key-value pairs associated with the object.
- For S3 objects, any object tag key-value pairs associated with the object.
- For replicated object copies, the current storage location of each copy.
- For erasure-coded object copies, the current storage location of each fragment.
- For object copies in a Cloud Storage Pool, the location of the object, including the name of the external bucket and the object's unique identifier.
- For segmented objects and multipart objects, a list of object segments including segment identifiers and data sizes. For objects with more than 100 segments, only the first 100 segments are shown.
- All object metadata in the unprocessed, internal storage format. This raw metadata includes internal system metadata that is not guaranteed to persist from release to release.

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

System Metadata

Object ID	A12E96FF-B13F-4905-9E9E-45373F6E7DA8
Name	testobject
Container	source
Account	t-1582139188
Size	5.24 MB
Creation Time	2020-02-19 12:15:59 PST
Modified Time	2020-02-19 12:15:59 PST

Replicated Copies

Node	Disk Path
99-97	/var/local/rangedb/2/p/06/0B/00nM8H\$ TFbnQQ CV2E
99-99	/var/local/rangedb/1/p/12/0A/00nM8H\$ TFboW28 CXG%

Raw Metadata

```
{
  "TYPE": "CNT",
  "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 cannot give you information about the correctness of the object data's placement or the completeness of the ILM policy. You must evaluate this yourself. For details, see [Review audit logs](#).

Related information

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

Work with ILM rules and ILM policies

Once you have created ILM rules and an ILM policy, you can continue to work with them, modifying their configuration as your storage requirements change.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

Delete an ILM rule

To keep the list of current ILM rules manageable, delete any ILM rules that you are not likely to use.

You cannot delete an ILM rule if it is currently used in the active policy or in the proposed policy. If you need to delete an ILM rule that is used a policy, you must perform these steps first:

1. Clone the active policy or edit the proposed policy.
2. Remove the ILM rule from the policy.
3. Save, simulate, and activate the new policy to make sure objects are protected as expected.


Steps

1. Select **ILM > Rules**.
2. Review the table entry for the rule you want to remove.

Confirm that the rule is not used in the active ILM policy or the proposed ILM policy.

3. If the rule you want to remove is not in use, select the radio button and select **Remove**.
4. Select **OK** to confirm that you want to delete the ILM rule.

The ILM rule is deleted.

If you delete a rule that is used in a historical policy, an  icon appears for the rule when you view the policy, which indicates that the rule has become a historical rule.



Viewing Historical Policy - Example ILM policy

Review the rules in this policy. If this is a proposed policy, click **Simulate**

Reason for change: new policy

Rules are evaluated in order, starting from the top

Rule Name

Erasure code larger objects

2 copies 2 sites  

This is a historical ILM rule. Historical rules are rules that were included in a policy and then edited or deleted after the policy became historical.

Edit an ILM rule

You might need to edit an ILM rule to change a filter or placement instruction.

You cannot edit a rule if it is being used in the proposed ILM policy or the active ILM policy. Instead, you can clone these rules and make any required changes to the cloned copy. You also cannot edit the stock ILM rule (Make 2 Copies) or ILM rules created before StorageGRID version 10.3.



Before adding an edited rule to the active ILM policy, be aware that a change to an object's placement instructions might cause an increased load on the system.


Steps

1. Select **ILM > Rules**.

The ILM Rules page appears. This page shows all available rules and indicates which rules are being used in the active policy or the proposed policy.

ILM Rules

Information lifecycle management (ILM) rules determine how and where object data is stored over time. Every object ingested into the StorageGRID Webscale is evaluated against the ILM rules that make up the active ILM policy. Use this page to manage and view ILM rules. You cannot edit or remove an ILM rule that is used by an active or proposed ILM policy.

<div> Create  Edit  Clone  Remove</div>			
Name	Used In Active Policy	Used In Proposed Policy	
<input type="radio"/> Make 2 Copies	✓	✓	
<input type="radio"/> PNGs		✓	
<input checked="" type="radio"/> JPGs			
<input type="radio"/> X-men		✓	

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

The Edit ILM Rule wizard opens.

Edit ILM Rule Step 1 of 3: Define Basics

Name: JPGs

Description:

Tenant Accounts (optional): Tenant-01 (16229710975421005503) Tenant-04 (83132053388229808098)

Bucket Name: contains az-01


Advanced filtering... (0 defined)

Cancel Next

- Complete the pages of the Edit ILM Rule wizard, following the steps for [creating an ILM rule](#) and [using advanced filters](#), as necessary.

When editing an ILM rule, you cannot change its name.

- Select **Save**.

If you edit a rule that is used in a historical policy, an  icon appears for the rule when you view the policy, which indicates that the rule has become a historical rule.



Viewing Historical Policy - Example ILM policy

Review the rules in this policy. If this is a proposed policy, click Simulat

Reason for change: new policy

Rules are evaluated in order, starting from the top

Rule Name

Erasure code larger objects

2 copies 2 sites



This is a historical ILM rule.
Historical rules are rules that were included a policy and then edited or deleted after the policy became historical.

Clone an ILM rule

You cannot edit a rule if it is being used in the proposed ILM policy or the active ILM policy. Instead, you can clone a rule and make any required changes to the cloned copy. Then, if required, you can remove the original rule from the proposed policy and replace it with the modified version. You cannot clone an ILM rule if it was created using StorageGRID version 10.2 or earlier.

Before adding a cloned rule to the active ILM policy, be aware that a change to an object's placement instructions might cause an increased load on the system.

Steps

- 1. Select **ILM > Rules**.

The ILM Rules page appears.

ILM Rules

Information lifecycle management (ILM) rules determine how and where object data is stored over time. Every object ingested into the StorageGRID Webscale is evaluated against the ILM rules that make up the active ILM policy. Use this page to manage and view ILM rules. You cannot edit or remove an ILM rule that is used by an active or proposed ILM policy.

+ Create

Edit

Clone

Remove

	Name	Used In Active Policy	Used In Proposed Policy
<input type="radio"/>	Make 2 Copies	✓	✓
<input type="radio"/>	PNGs		✓
<input checked="" type="radio"/>	JPGs		
<input type="radio"/>	X-men		✓

- 2. Select the ILM rule you want to clone, and select **Clone**.

The Create ILM Rule wizard opens.

- 3. Update the cloned rule by following the steps for editing an ILM rule and using advanced filters.

When cloning an ILM rule, you must enter a new name.

- 4. Select **Save**.


The new ILM rule is created.

View the ILM policy activity queue

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

Steps

- 1. Select **Dashboard**.



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

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

Use S3 Object Lock with ILM

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 are not deleted or overwritten for a specified amount of time.

What is S3 Object Lock?

The StorageGRID S3 Object Lock feature is an object-protection solution that is equivalent to S3 Object Lock in Amazon Simple Storage Service (Amazon S3).

As shown in the figure, when the global S3 Object Lock setting is enabled for a StorageGRID system, an S3 tenant account can create buckets with or without S3 Object Lock enabled. If a bucket has S3 Object Lock enabled, S3 client applications can optionally specify retention settings for any object version in that bucket. An object version must have retention settings specified to be protected by S3 Object Lock. In addition, each bucket that has S3 Object Lock enabled can optionally have a default retention mode and retention period, which apply if objects are added to the bucket without their own retention settings.

StorageGRID with S3 Object Lock setting enabled



The StorageGRID S3 Object Lock feature provides a single retention mode that is equivalent to the Amazon S3 compliance mode. By default, a protected object version cannot be overwritten or deleted by any user. The StorageGRID S3 Object Lock feature does not support a governance mode, and it does not allow users with special permissions to bypass retention settings or to delete protected objects.

If a bucket has S3 Object Lock enabled, the S3 client application can optionally specify either or both of the following object-level retention settings when creating or updating an object:

- **Retain-until-date:** If an object version's retain-until-date is in the future, the object can be retrieved, but it cannot be modified or deleted. As required, an object's retain-until-date can be increased, but this date cannot be decreased.
- **Legal hold:** Applying a legal hold to an object version immediately locks that object. For example, you might need to put a legal hold on an object that is related to an investigation or legal dispute. A legal hold has no expiration date, but remains in place until it is explicitly removed. Legal holds are independent of the retain-until-date.

For details about object retention settings, go to [Use S3 Object Lock](#).

For details about default bucket retention settings, go to [Use S3 Object Lock default bucket 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."

If you previously enabled the global Compliance setting, the global S3 Object Lock setting was enabled automatically. Tenant users are no longer be able to create new buckets with Compliance enabled; however, as required, tenant users can continue to use and manage any existing legacy Compliant buckets, which includes performing the following tasks:

- Ingesting new objects into an existing bucket that has legacy Compliance enabled.
- Increasing the retention period of an existing bucket that has legacy Compliance enabled.

- Changing the auto-delete setting for an existing bucket that has legacy Compliance enabled.
- Placing a legal hold on an existing bucket that has legacy Compliance enabled.
- Lifting a legal hold.

See [NetApp Knowledge Base: How to manage legacy Compliant buckets in StorageGRID 11.5](#) for instructions.

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 (new)	Compliance (legacy)
How is the feature enabled globally?	From the Grid Manager, select CONFIGURATION > System > S3 Object Lock .	No longer supported. Note: If you enabled the global Compliance setting using a previous version of StorageGRID, the S3 Object Lock setting is enabled in StorageGRID 11.6. You can continue to use StorageGRID to manage the settings of existing compliant buckets; however, you cannot create new compliant buckets.
How is the feature enabled for a bucket?	Users must enable S3 Object Lock when creating a new bucket using the Tenant Manager, the Tenant Management API, or the S3 REST API.	Users can no longer create new buckets with Compliance enabled; however, they can continue to add new objects to existing Compliant buckets.
Is bucket versioning supported?	Yes. Bucket versioning is required and is enabled automatically when S3 Object Lock is enabled for the bucket.	No. The legacy Compliance feature does not allow bucket versioning.
How is object retention set?	Users can set a retain-until-date for each object version.	Users must set a retention period for the entire bucket. The retention period applies to all objects in the bucket.
Can a bucket have default settings for retention and legal hold?	Yes. StorageGRID buckets that have S3 Object Lock enabled can have a default retention period that is applied to object versions that do not have their own retention settings specified during ingest.	Yes
Can the retention period be changed?	The retain-until-date for an object version can be increased but never decreased.	The bucket's retention period can be increased but never decreased.

	S3 Object Lock (new)	Compliance (legacy)
Where is legal hold controlled?	Users can place a legal hold or lift a legal hold for any object version in the bucket.	A legal hold is placed on the bucket and affects all objects in the bucket.
When can objects be deleted?	An object version can be deleted after the retain-until-date is reached, assuming the object is not under legal hold.	An object can be deleted after the retention period expires, assuming the bucket is not under legal hold. Objects can be deleted automatically or manually.
Is bucket lifecycle configuration supported?	Yes	No

Workflow for S3 Object Lock

As a grid administrator, you must coordinate closely with tenant users to ensure that the objects are protected in a manner that satisfies their retention requirements.

The workflow diagram shows the high-level steps for using S3 Object Lock. These steps are performed by the grid administrator and by tenant users.



Grid admin tasks

As the workflow diagram shows, a grid administrator must perform two high-level tasks before S3 tenant users can use S3 Object Lock:

1. Create at least one compliant ILM rule and make that rule the default rule in the active ILM policy.
2. Enable the global S3 Object Lock setting for the entire StorageGRID system.

Tenant user tasks

After the global S3 Object Lock setting has been enabled, tenants can perform these tasks:

1. Create buckets that have S3 Object Lock enabled.
2. Specify default retention settings for the bucket, which are applied to objects added to the bucket that do not specify their own retention settings.
3. Add objects to those buckets and specify object-level retention periods and legal hold settings.
4. As required, update a retention period or change the legal hold setting for an individual object.

Related information

- [Use a tenant account](#)
- [Use S3](#)
- [Use S3 Object Lock default bucket retention](#)

Requirements for S3 Object Lock

You must review the requirements for enabling the global S3 Object Lock setting, the requirements for creating compliant ILM rules and ILM policies, and the restrictions StorageGRID places on buckets and objects that use S3 Object Lock.

Requirements for using the global S3 Object Lock setting

- You must enable the global S3 Object Lock setting using the Grid Manager or the Grid Management API before any S3 tenant can create a bucket with S3 Object Lock enabled.
- Enabling the global S3 Object Lock setting allows all S3 tenant accounts to create buckets with S3 Object Lock enabled.
- After you enable the global S3 Object Lock setting, you cannot disable the setting.
- You cannot enable the global S3 Object Lock unless the default rule in the active ILM policy is *compliant* (that is, the default rule must comply with the requirements of buckets with S3 Object Lock enabled).
- When the global S3 Object Lock setting is enabled, you cannot create a new proposed ILM policy or activate an existing proposed ILM policy unless the default rule in the policy is compliant. After the global S3 Object Lock setting has been enabled, the ILM Rules and ILM Policies pages indicate which ILM rules are compliant.

In the following example, the ILM Rules page lists three rules that are compliant with buckets with S3 Object Lock enabled.

<div> <div>+ Create</div> <div>📄 Clone</div> <div>✎ Edit</div> <div>✕ Remove</div> </div>			
Name	Compliant	Used In Active Policy	Used In Proposed Policy
Make 2 Copies	✓	✓	
Compliant Rule: EC for objects in bank-records bucket	✓		
2 copies 10 years, Archive forever			
2 Copies 2 Data Centers	✓		

Compliant Rule: EC for objects in bank-records bucket

Description:

2+1 EC at one site

Ingest Behavior:

Balanced

Compliant:

Yes

Tenant Accounts:

Bank of ABC (94793396288150002349)

Bucket Name:

equals 'bank-records'

Reference Time:

Ingest Time

Requirements for compliant ILM rules

If you want to enable the global S3 Object Lock setting, you must ensure that the default rule in your active ILM policy is compliant. A compliant rule satisfies the requirements of both buckets with S3 Object Lock enabled and any existing buckets that have legacy Compliance enabled:

- It must create at least two replicated object copies or one erasure-coded copy.
- These copies must exist on Storage Nodes for the entire duration of each line in the placement instructions.
- Object copies cannot be saved in a Cloud Storage Pool.
- Object copies cannot be saved on Archive Nodes.
- At least one line of the placement instructions must start at day 0, using **Ingest Time** as the reference time.
- At least one line of the placement instructions must be “forever.”

For example, this rule satisfies the requirements of buckets with S3 Object Lock enabled. It stores two replicated object copies from Ingest Time (day 0) to “forever.” The objects will be stored on Storage Nodes at two data centers.

Compliant rule: 2 replicated copies at 2 sites

Description:

2 replicated copies on Storage Nodes from Day 0 to Forever

Ingest Behavior:

Balanced

Compliant:

Yes

Tenant Accounts:

Bank of ABC (94793396288150002349)

Reference Time:

Ingest Time

Filtering Criteria:

Matches all objects.

Retention Diagram:

Trigger

Day 0

dc1

dc2

Duration

Forever

Requirements for active and proposed ILM policies

When the global S3 Object Lock setting is enabled, active and proposed ILM policies can include both compliant and non-compliant rules.

- The default rule in the active or any proposed ILM policy must be compliant.
- Non-compliant rules only apply to objects in buckets that do not have S3 Object Lock enabled or that do not have the legacy Compliance feature enabled.
- Compliant rules can apply to objects in any bucket; S3 Object Lock or legacy Compliance does not need to be enabled for the bucket.

A compliant ILM policy might include these three rules:

1. A compliant rule that creates erasure-coded copies of the objects in a specific bucket with S3 Object Lock enabled. The EC copies are stored on Storage Nodes from day 0 to forever.
2. A non-compliant rule that creates two replicated object copies on Storage Nodes for a year and then moves one object copy to Archive Nodes and stores that copy forever. This rule only applies to buckets that do not have S3 Object Lock or legacy Compliance enabled because it stores only one object copy forever and it uses Archive Nodes.
3. A default, compliant rule that creates two replicated object copies on Storage Nodes from day 0 to forever. This rule applies to any object in any bucket that was not filtered out by the first two rules.

Requirements for buckets with S3 Object Lock enabled

- If the global S3 Object Lock setting is enabled for the StorageGRID system, you can use the Tenant Manager, the Tenant Management API, or the S3 REST API to create buckets with S3 Object Lock enabled.

This example from the Tenant Manager shows a bucket with S3 Object Lock enabled.

Buckets

Create buckets and manage bucket settings.

1 bucket Create bucket

Actions ▾

<input type="checkbox"/>	Name ▾	S3 Object Lock ? ▾	Region ▾	Object Count ? ▾	Space Used ? ▾	Date Created ▾
<input type="checkbox"/>	bank-records	✓	us-east-1	0	0 bytes	2021-01-06 16:53:19 MST

← Previous 1 Next →

- If you plan to use S3 Object Lock, you must enable S3 Object Lock when you create the bucket. You cannot enable S3 Object Lock for an existing bucket.
- Bucket versioning is required with S3 Object Lock. When S3 Object Lock is enabled for a bucket, StorageGRID automatically enables versioning for that bucket.
- After you create a bucket with S3 Object Lock enabled, you cannot disable S3 Object Lock or suspend versioning for that bucket.
- Optionally, you can configure default retention for a bucket. When an object version is uploaded, the default retention is applied to the object version. You can override the bucket default by specifying a retention mode and retain-until-date in the request to upload an object version.
- Bucket lifecycle configuration is supported for S3 Object Lifecycle buckets.

- CloudMirror replication is not supported for buckets with S3 Object Lock enabled.

Requirements for objects in buckets with S3 Object Lock enabled

- To protect an object version, the S3 client application must either configure bucket default retention, or specify retention settings in each upload request.
- You can increase the retain-until-date for an object version, but you can never decrease this value.
- If you are notified of a pending legal action or regulatory investigation, you can preserve relevant information by placing a legal hold on an object version. When an object version is under a legal hold, that object cannot be deleted from StorageGRID, even if it has reached its retain-until-date. As soon as the legal hold is lifted, the object version can be deleted if the retain-until-date has been reached.
- S3 Object Lock requires the use of versioned buckets. Retention settings apply to individual object versions. An object version can have both a retain-until-date and a legal hold setting, one but not the other, or neither. Specifying a retain-until-date or a legal hold setting for an object protects only the version specified in the request. You can create new versions of the object, while the previous version of the object remains locked.

Lifecycle of objects in buckets with S3 Object Lock enabled

Each object that is saved in a bucket with S3 Object Lock enabled goes through three stages:

1. Object ingest

- When adding an object version to a bucket with S3 Object Lock enabled, the S3 client application can use the default bucket retention settings or optionally specify retention settings for the object (retain-until-date, legal hold, or both). StorageGRID then generates metadata for that object, which includes a unique object identifier (UUID) and the ingest date and time.
- After an object version with retention settings is ingested, its data and S3 user-defined metadata cannot be modified.
- StorageGRID stores the object metadata independently of the object data. It maintains three copies of all object metadata at each site.

2. Object retention

- Multiple copies of the object are stored by StorageGRID. The exact number and type of copies and the storage locations are determined by the compliant rules in the active ILM policy.

3. Object deletion

- An object can be deleted when its retain-until-date is reached.
- An object that is under a legal hold cannot be deleted.

Related information

- [Use a tenant account](#)
- [Use S3](#)
- [Comparing S3 Object Lock to legacy Compliance](#)
- [Example 7: Compliant ILM policy for S3 Object Lock](#)
- [Review audit logs](#)
- [Use S3 Object Lock default bucket retention.](#)

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.

What you'll need

- 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 must understand the considerations.
- The default rule in the active ILM policy is compliant.
 - [Create a default ILM rule](#)
 - [Create an ILM policy](#)

About this task

A grid administrator must enable the global S3 Object Lock setting to allow tenant users to create new buckets that have S3 Object Lock enabled. After this setting is enabled, it cannot be disabled.



If you enabled the global Compliance setting using a previous version of StorageGRID, the S3 Object Lock setting is enabled in StorageGRID 11.6. You can continue to use StorageGRID to manage the settings of existing compliant buckets; however, you cannot create new compliant buckets. 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.

S3 Object Lock Settings

Enable S3 Object Lock for your entire StorageGRID system if S3 tenant accounts need to satisfy regulatory compliance requirements when saving object data. After this setting is enabled, it cannot be disabled.

S3 Object Lock

Before enabling S3 Object Lock, you must ensure that the default rule in the active ILM policy is compliant. A compliant rule satisfies the requirements of buckets with S3 Object Lock enabled.

- It must create at least two replicated object copies or one erasure-coded copy.
- These copies must exist on Storage Nodes for the entire duration of each line in the placement instructions.
- Object copies cannot be saved on Archive Nodes.
- At least one line of the placement instructions must start at day 0, using Ingest Time as the reference time.
- At least one line of the placement instructions must be "forever".

☐ Enable S3 Object Lock

Apply

If you had enabled the global Compliance setting using a previous version of StorageGRID, the page includes the following note:

The S3 Object Lock setting replaces the legacy Compliance setting. When this setting is enabled, tenant users can create buckets with S3 Object Lock enabled. Tenants who previously created buckets for the legacy Compliance feature can manage their existing buckets, but can no longer create new buckets with legacy Compliance enabled. See [Managing objects with information lifecycle management](#) for information.

2. Select **Enable S3 Object Lock**.
3. Select **Apply**.


A confirmation dialog box appears and reminds you that you cannot disable S3 Object Lock after it is enabled.



4. If you are sure you want to permanently enable S3 Object Lock for your entire system, select **OK**.

When you select **OK**:

- If the default rule in the active ILM policy is compliant, S3 Object Lock is now enabled for the entire grid and cannot be disabled.
- If the default rule is not compliant, an error appears, indicating that you must create and activate a new ILM policy that includes a compliant rule as its default rule. Select **OK**, and create a new proposed policy, simulate it, and activate it.



After you finish

After you enable the global S3 Object Lock setting, you might need to [create a default rule](#) that is compliant and [create an ILM policy](#) that is compliant. After the setting is enabled, the ILM policy can optionally include both a compliant default rule and a non-compliant default rule. For example, you might want to use a non-compliant rule that does not have filters for objects in buckets that do not have S3 Object Lock enabled.

Related information

- [Compare S3 Object Lock to legacy Compliance](#)


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 displays an error message. Specifically:

- Tenant Manager users see the following error message:



- Tenant Management API users and S3 API users receive a response code of 503 `Service Unavailable` with similar message text.

To resolve this error, follow these steps:

1. Attempt to make all Storage Nodes or sites available again as soon as possible.
2. If you are unable to make enough of the Storage Nodes at each site available, contact technical support, who can help you recover nodes and ensure that changes are consistently applied across the grid.
3. Once the underlying issue has been resolved, remind the tenant user to retry their configuration changes.

Related information

- [Use a tenant account](#)
- [Use S3](#)
- [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 the proposed policy to confirm it will work as intended to protect content from loss.

ILM rule 1 for example 1: Copy object data to two data centers

This example ILM rule copies object data to storage pools in two data centers.

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

Edit ILM Rule

Step 2 of 3: Define Placements

Configure placement instructions to specify how you want objects matched by this rule to be stored.

Two Copies Two Data Centers

Reference Time

Ingest Time

Placements

Sort by start day

From day

0

store

forever

Add

Remove

Type

replicated

Location

Storage Pool DC1Storage Pool DC2Add Pool

Copies

2

+

×

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

Retention Diagram

Refresh

Trigger

Day 0

Storage Pool DC1

Storage Pool DC2

Duration

Forever

Cancel

Back

Next

ILM rule 2 for example 1: Erasure Coding profile with bucket matching

This example ILM rule uses an Erasure Coding profile and an S3 bucket to determine where and how long the object is stored.

427

Rule definition	Example value
Erasure Coding Profile	<ul style="list-style-type: none"> • One storage pool across three data centers (All 3 sites) • Use 6+3 erasure-coding scheme
Rule Name	EC for S3 bucket finance-records
Reference Time	Ingest Time
Content Placement	For objects in the S3 bucket named finance-records, create one erasure-coded copy in the pool specified by the Erasure Coding profile. Keep this copy forever.

Create ILM Rule Step 2 of 3: Define Placements

Configure placement instructions to specify how you want objects matched by this rule to be stored.

EC for S3 bucket finance-records

Reference Time

Ingest Time

Placements

Sort by start day

From day

0

store

forever

Add

Remove

Type

erasure coded

Location

All 3 sites (6 plus 3)

Copies

1

+

×

Retention Diagram

Refresh

Trigger

Day 0

Duration

Forever

Cancel Back Next

ILM policy for example 1

The StorageGRID system allows you to design sophisticated and complex ILM policies; however, in practice, most ILM policies are simple.

A typical ILM policy for a multi-site topology might include ILM rules such as the following:

- At ingest, use 6+3 erasure coding to store all objects belonging to the S3 bucket named `finance-records` across three data centers.
- If an object does not match the first ILM rule, use the policy's default ILM rule, Two Copies Two Data Centers, to store a copy of that object in two data centers, DC1 and DC2.

Configure ILM Policy

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

Name	Object Storage Policy
Reason for change	new proposed policy

Rules

1. Select the rules you want to add to the policy.
2. Determine the order in which the rules will be evaluated by dragging and dropping the rows. The default rule will be automatically placed at the end of the policy and cannot be moved.

+ Select Rules

Default	Rule Name	Tenant Account	Actions
	EC for S3 bucket finance-records	Ignore	✕
✓	Two Copies Two Data Centers	Ignore	✕

Cancel

Save

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

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



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate the proposed policy to confirm it will work as intended to protect content from loss.

ILM rule 1 for example 2: Use EC for 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. Do not 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 Filtering for Object Size	Object Size (MB) greater than 1
Content Placement	Create a 2+1 erasure-coded copy using three sites

EC only objects > 1 MB

Matches all of the following metadata:

Object Size (MB) ▼	greater than ▼	1	+ ×
<div style="border: 1px solid #ccc; padding: 2px; display: inline-block;">+ ×</div>			

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 Filtering for Object Size	None
Content Placement	Create two replicated copies and save them at two data centers, DC1 and DC2

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.

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

Name

Reason for change

Rules

1. Select the rules you want to add to the policy.
2. Determine the order in which the rules will be evaluated by dragging and dropping the rows. The default rule will be automatically placed at the end of the policy and cannot be moved.

[+ Select Rules](#)

Default	Rule Name	Tenant Account	Actions
	EC only objects > 1 MB	—	
<input checked="" type="checkbox"/>	Two replicated copies	—	

[Cancel](#) [Save](#)

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

You can use the following example rules and policy to ensure that images 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 the proposed policy to confirm it will work as intended to protect content from loss.

ILM rule 1 for example 3: Use EC for image files 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. Do not use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.

Rule definition	Example value
Rule Name	EC image files > 1 MB
Reference Time	Ingest Time

Rule definition	Example value
Advanced Filtering for Object Size	Object Size (MB) greater than 1.0
Advanced Filtering for User Metadata	User Metadata type equals image
Content Placement	Create a 2+1 erasure-coded copy using three sites

EC image files > 1 MB

Matches all of the following metadata:

Object Size (MB)
greater than
1
+
x

User Metadata
type
equals
image
+
x

+
x

Because this rule is configured as the first rule in the policy, the erasure-coding placement instruction only applies to images 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 Filtering for User Metadata	User Metadata type equals image files
Content Placement	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).

Reason for change:		
new policy		
Rules are evaluated in order, starting from the top.		
Rule Name	Default	Tenant Account
EC image files > 1 MB 		—
2 copies for small images 		—
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.

As this example shows, you can control the amount of storage used by versioned objects by using different placement instructions for noncurrent object versions.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate the proposed policy to confirm it will work as intended to protect content from loss.



If you create ILM policies to manage noncurrent object versions, be aware that you must know the object version's UUID or CBID to simulate the policy. To find an object's UUID and CBID, use Object Metadata Lookup while the object is still current. See [Verify an ILM policy with object metadata lookup](#).

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 data centers for 10 years.

This rule applies to all objects, whether or not they are versioned.

Rule definition	Example value
Storage Pools	Three storage pools, each at different data centers, named DC1, DC2, and DC3.
Rule Name	Three Copies Ten Years
Reference Time	Ingest Time
Content Placement	On Day 0, keep three replicated copies for 10 years (3,652 days), one in DC1, one in DC2, and one in DC3. At the end of 10 years, delete all copies of the object.

Configure placement instructions to specify how you want objects matched by this rule to be stored.

Three Copies Ten Years

Save three copies for ten years

Reference Time

Ingest Time

Placements

Sort by start day

From day

0

store for

3652

days

Add

Remove

Type

replicated

Location

DC1

DC2

DC3

Add Pool

Copies

3

+

×

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

Retention Diagram

Refresh

Trigger

Day 0

Day 3652

DC1

DC2

DC3

Duration

3652 days

Forever

Cancel

Back

Next

ILM rule 2 for example 4: Save two copies of noncurrent versions for 2 years

This example ILM rule stores two copies of the noncurrent versions of an S3 versioned object for 2 years.

Because ILM rule 1 applies to all versions of the object, you must create another rule to filter out any noncurrent versions. This rule uses the **Noncurrent Time** option for Reference Time.

In this example, only two copies of the noncurrent versions are stored, and those copies will be stored for two years.

Rule definition	Example value
Storage Pools	Two storage pools, each at different data centers, named DC1 and DC2.
Rule Name	Noncurrent Versions: Two Copies Two Years
Reference Time	Noncurrent Time
Content Placement	On Day 0 relative to Noncurrent Time (that is, starting from the day the object version becomes the noncurrent version), keep two replicated copies of the noncurrent object versions for 2 years (730 days), one in DC1 and one in DC2. At the end of 2 years, delete the noncurrent versions.

Noncurrent Versions: Two Copies Two Years

Save two copies of noncurrent versions for two years

Reference Time

Noncurrent Time

Placements ?

Sort by start day

From day

0

store

for

730

days

Add

Remove

Type

replicated

Location

DC1

DC2

Add Pool

Copies

2

+

x

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

Retention Diagram ?

Refresh

Trigger

Day 0

Year 2

DC1

DC2

Duration

2 years

Forever

ILM policy for example 4: S3 versioned objects

If you want to manage older versions of an object differently than the current version, rules that use **Noncurrent Time** as the Reference Time must appear in the ILM policy before rules that apply to the current object version.

An ILM policy for S3 versioned objects might include ILM rules such as the following:

- Keep any older (noncurrent) versions of each object for 2 years, starting from the day the version became noncurrent.



The Noncurrent Time rules must appear in the policy before the rules that apply to the current object version. Otherwise, the noncurrent object versions will never be matched by the Noncurrent Time rule.

- At ingest, create three replicated copies and store one copy at each of three data centers. Keep copies of the current object version for 10 years.

Configure ILM Policy

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



Name ILM Policy for S3 Versioned Objects

Reason for change store 3 copies of current version for 10 years and 2 copies of noncurrent versions for 2 years

Rules

1. Select the rules you want to add to the policy.
2. Determine the order in which the rules will be evaluated by dragging and dropping the rows. The default rule will be automatically placed at the end of the policy and cannot be moved.

+ Select Rules

Default	Rule Name	Tenant Account	Actions
	Noncurrent Versions: Two Copies Two Years 	Ignore	✕
✓	Three Copies Ten Years 	Ignore	✕

The default ILM rule in this policy does not retain objects forever. Confirm this is the behavior you expect. Otherwise, any objects that are not matched by another rule will be deleted after 3652 days.

Cancel

Save

When you simulate the example policy, you would expect test objects to be evaluated as follows:

- Any noncurrent object versions would be matched by the first rule. If a noncurrent object version is older than 2 years, it is permanently deleted by ILM (all copies of the noncurrent version removed from the grid).



To simulate noncurrent object versions, you must use that version's UUID or CBID. While the object is still current, you can use Object Metadata Lookup to find its UUID and CBID.

- The current object version would be matched by the second rule. When the current object version has been stored for 10 years, the ILM process adds a delete marker as the current version of the object, and it makes the previous object version "noncurrent." The next time ILM evaluation occurs, this noncurrent version is matched by the first rule. As a result, the copy at DC3 is purged and the two copies at DC1 and DC2 are stored for 2 more years.

Example 5: ILM rules and policy for Strict ingest behavior

You can use a location filter and the Strict ingest behavior in a rule to prevent objects from being saved at a particular data center location.

In this example, a Paris-based tenant does not want to store some objects outside of the EU because of regulatory concerns. Other objects, including all objects from other tenant accounts, can be stored at either the Paris data center or the US data center.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate the proposed policy to confirm it will work as intended to protect content from loss.

Related information

- [Data-protection options for ingest](#)
- [Step 3 of 3: Define ingest behavior](#)

ILM rule 1 for example 5: Strict ingest to guarantee Paris data center

This example ILM rule uses the Strict ingest behavior to guarantee that objects saved by a Paris-based tenant to S3 buckets with the region set to eu-west-3 region (Paris) are never stored at the US data center.

This rule applies to objects that belong to the Paris tenant and that have the S3 bucket region set to eu-west-3 (Paris).

Rule definition	Example value
Tenant Account	Paris tenant
Advanced Filtering	Location Constraint equals eu-west-3
Storage Pools	DC1 (Paris)
Rule Name	Strict ingest to guarantee Paris data center
Reference Time	Ingest Time
Content Placement	On Day 0, keep two replicated copies forever in DC1 (Paris)
Ingest Behavior	Strict. Always use this rule's placements on ingest. Ingest fails if it is not possible to store two copies of the object at the Paris data center.

Strict ingest to guarantee Paris data center

Description:

Strict ingest to guarantee Paris data center

Ingest Behavior:

Strict

Tenant Account:

Paris tenant (25580610012441844135)

Reference Time:

Ingest Time

Filtering Criteria:

Matches all of the following metadata:

System Metadata

Location Constraint (S3 only)

equals

eu-west-3

Retention Diagram:

Trigger

Day 0

DC1 (Paris)

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 cannot be satisfied immediately, interim copies are stored at any available location.

This rule applies to objects that belong to any tenant and any region.

Rule definition	Example value
Tenant Account	Ignore
Advanced Filtering	<i>Not specified</i>
Storage Pools	DC1 (Paris) and DC2 (US)
Rule Name	2 Copies 2 Data Centers
Reference Time	Ingest Time
Content Placement	On Day 0, keep two replicated copies forever at two data centers
Ingest Behavior	Balanced. Objects that match this rule are placed according to the rule's placement instructions if possible. Otherwise, interim copies are made at any available location.

2 Copies 2 Data Centers

Description:

2 Copies 2 Data Centers

Ingest Behavior:

Balanced

Reference Time:

Ingest Time

Filtering Criteria:

Matches all objects.

Retention Diagram:

Trigger

Day 0

DC1 (Paris)

DC2 (US)

Duration

Forever

ILM policy for example 5: Combining ingest behaviors

The example ILM policy includes two rules that have different ingest behaviors.

An ILM policy that uses two different ingest behaviors might include ILM rules such as the following:

- Store objects that belong to the Paris tenant and that have the S3 bucket region set to eu-west-3 (Paris) only in the Paris data center. Fail ingest if the Paris data center is not available.
- Store all other objects (including those that belong to the Paris tenant but that have a different bucket region) in both the US data center and the Paris data center. Make interim copies in any available location if the placement instruction cannot be satisfied.

Configure ILM Policy

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

Name

Example policy for Strict ingest

Reason for change

Do not store certain objects for Paris tenant in US

Rules

1. Select the rules you want to add to the policy.

2. Determine the order in which the rules will be evaluated by dragging and dropping the rows. The default rule will be automatically placed at the end of the policy and cannot be moved.

+ Select Rules

Default	Rule Name	Tenant Account	Actions
	Strict ingest to guarantee Paris data center	Paris tenant (25580610012441844135)	✕
✓	2 Copies 2 Data Centers	Ignore	✕

Cancel

Save

When you simulate the example policy, you expect test objects to be evaluated as follows:

- Any objects that belong to the Paris tenant and that have the S3 bucket region set to eu-west-3 are matched by the first rule and are stored at the Paris data center. Because the first rule uses Strict ingest, these objects are never stored at the US data center. If the Storage Nodes at the Paris data center are not available, ingest fails.
- All other objects are matched by the second rule, including objects that belong to the Paris tenant and that do not have the S3 bucket region set to eu-west-3. One copy of each object is saved at each data center. However, because the second rule uses Balanced ingest, if one data center is unavailable, two interim copies are saved at any available location.

Example 6: Changing an ILM policy

You might need to create and activate a new ILM policy if your data protection needs change or you add new sites.

Before changing a policy, you must understand how changes in ILM placements can temporarily affect the overall performance of a StorageGRID system.

In this example, a new StorageGRID site has been added in an expansion and the active ILM policy needs to be revised to store data at the new site.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate the proposed policy to confirm it will work as intended to protect content from loss.

How does changing an ILM policy affect performance

When you activate a new ILM policy, the performance of your StorageGRID system might be temporarily affected, especially if the placement instructions in the new policy require many existing objects to be moved to new locations.



When you activate a new ILM policy, StorageGRID uses it to manage all objects, including existing objects and newly ingested objects. Before activating a new ILM policy, review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

The types of ILM policy changes that can temporarily affect StorageGRID performance include the following:

- Applying a different Erasure Coding profile to existing erasure-coded objects.



StorageGRID considers each Erasure Coding profile to be unique and does not reuse erasure-coding fragments when a new profile is used.

- Changing the type of copies required for existing objects; for example, converting a large percentage of replicated objects to erasure-coded objects.
- Moving copies of existing objects to a completely different location; for example, moving a large number of objects to or from a Cloud Storage Pool or to or from a remote site.

Related information

[Create an ILM policy](#)

Active ILM policy for example 6: Data protection at two sites

In this example, the active ILM policy was initially designed for a two-site StorageGRID system and uses two ILM rules.

ILM Policies

Review the proposed, active, and historical policies. You can create, edit, or delete a proposed policy; clone the active policy; or view the details for any policy.

[+ Create Proposed Policy](#) [Clone](#) [Edit](#) [Remove](#)

Policy Name	Policy State	Start Date	End Date
<input checked="" type="radio"/> Data Protection for Two Sites	Active	2020-06-10 16:42:09 MDT	
<input type="radio"/> Baseline 2 Copies Policy	Historical	2020-06-09 21:48:34 MDT	2020-06-10 16:42:09 MDT

Viewing Active Policy - Data Protection for Two Sites

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

Reason for change: Data Protection for Two Sites

Rules are evaluated in order, starting from the top.

Rule Name	Default	Tenant Account
One-Site Erasure Coding for Tenant A 🔗		Tenant A (49752734300032812036)
Two-Site Replication for Other Tenants 🔗	✓	Ignore

[Simulate](#) [Activate](#)

In this ILM policy, objects belonging to Tenant A are protected by 2+1 erasure coding at a single site, while objects belonging to all other tenants are protected across two sites using 2-copy replication.



The first rule in this example uses an advanced filter to ensure that erasure coding is not used for small objects. Any of Tenant A's objects that are smaller than 1 MB will be protected by the second rule, which uses replication.

Rule 1: One-site erasure coding for Tenant A

Rule definition	Example value
Rule Name	One-Site Erasure Coding for Tenant A
Tenant Account	Tenant A
Storage Pool	Data Center 1
Content Placement	2+1 erasure coding in Data Center 1 from day 0 to forever

Rule 2: Two-site replication for other tenants

Rule definition	Example value
Rule Name	Two-Site Replication for Other Tenants
Tenant Account	Ignore
Storage Pools	Data Center 1 and Data Center 2

Rule definition	Example value
Content Placement	Two replicated copies from day 0 to forever: one copy at Data Center 1 and one copy at Data Center 2.

Proposed ILM policy for example 6: Data protection at three sites

In this example, the ILM policy is being updated for a three-site StorageGRID system.

After performing an expansion to add the new site, the grid administrator created two new storage pools: a storage pool for Data Center 3 and a storage pool containing all three sites (not the same as the All Storage Nodes default storage pool). Then, the administrator created two new ILM rules and a new proposed ILM policy, which is designed to protect data at all three sites.

Viewing Proposed Policy - Data Protection for Three Sites

Before activating a new ILM policy:

- Review and carefully simulate the policy. Errors in an ILM policy can cause irreparable data loss.
- Review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

See [Managing objects with information lifecycle management](#) for more information.

This policy contains a rule that makes an erasure-coded copy. Confirm that at least one rule uses the Object Size advanced filter to prevent objects that are 200 KB or smaller from being erasure coded. See [Managing objects with information lifecycle management](#) for more information.

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

Reason for change: Data Protection for Three Sites

Rules are evaluated in order, starting from the top.

Rule Name	Default	Tenant Account
Three-Site Erasure Coding for Tenant A		Tenant A (49752734300032812036)
Three-Site Replication for Other Tenants	✓	Ignore

When this new ILM policy is activated, objects belonging to Tenant A will be protected by 2+1 erasure coding at three sites, while objects belonging to other tenants (and smaller objects belonging to Tenant A) will be protected across three sites using 3-copy replication.

Rule 1: Three-site erasure coding for Tenant A

Rule definition	Example value
Rule Name	Three-Site Erasure Coding for Tenant A
Tenant Account	Tenant A
Storage Pool	All 3 Data Centers (includes Data Center 1, Data Center 2, and Data Center 3)
Content Placement	2+1 erasure coding in All 3 Data Centers from day 0 to forever

Rule 2: Three-site replication for other tenants

Rule definition	Example value
Rule Name	Three-Site Replication for Other Tenants
Tenant Account	Ignore
Storage Pools	Data Center 1, Data Center 2, and Data Center 3
Content Placement	Three replicated copies from day 0 to forever: one copy at Data Center 1, one copy at Data Center 2, and one copy at Data Center 3.

Activating the proposed ILM policy for example 6

When you activate a new proposed ILM policy, existing objects might be moved to new locations or new object copies might be created for existing objects, based on the placement instructions in any new or updated rules.



Errors in an ILM policy can cause unrecoverable data loss. Carefully review and simulate the policy before activating it to confirm that it will work as intended.



When you activate a new ILM policy, StorageGRID uses it to manage all objects, including existing objects and newly ingested objects. Before activating a new ILM policy, review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

What happens when erasure-coding instructions change

In the currently active ILM policy for this example, objects belonging to Tenant A are protected using 2+1 erasure coding at Data Center 1. In the new proposed ILM policy, objects belonging to Tenant A will be protected using 2+1 erasure coding at Data Centers 1, 2, and 3.

When the new ILM policy is activated, the following ILM operations occur:

- New objects ingested by Tenant A are split into two data fragments and one parity fragment is added. Then, each of the three fragments is stored at a different data center.
- The existing objects belonging to Tenant A are re-evaluated during the ongoing ILM scanning process. Because the ILM placement instructions use a new Erasure Coding profile, entirely new erasure-coded fragments are created and distributed to the three data centers.



The existing 2+1 fragments at Data Center 1 are not reused. StorageGRID considers each Erasure Coding profile to be unique and does not reuse erasure-coding fragments when a new profile is used.

What happens when replication instructions change

In the currently active ILM policy for this example, objects belonging other tenants are protected using two replicated copies in storage pools at Data Centers 1 and 2. In the new proposed ILM policy, objects belonging to other tenants will be protected using three replicated copies in storage pools at Data Centers 1, 2, and 3.

When the new ILM policy is activated, the following ILM operations occur:

- When any tenant other than Tenant A ingests a new object, StorageGRID creates three copies and saves one copy at each data center.
- Existing objects belonging to these other tenants are re-evaluated during the ongoing ILM scanning process. Because the existing object copies at Data Center 1 and Data Center 2 continue to satisfy the replication requirements of the new ILM rule, StorageGRID only needs to create one new copy of the object for Data Center 3.

Performance impact of activating this policy

When the proposed ILM policy in this example is activated, the overall performance of this StorageGRID system will be temporarily affected. Higher than normal levels of grid resources will be required to create new erasure-coded fragments for Tenant A's existing objects and new replicated copies at Data Center 3 for other tenants' existing objects.

As a result of the ILM policy change, client read and write requests might temporarily experience higher than normal latencies. Latencies will return to normal levels after the placement instructions are fully implemented across the grid.

To avoid resource issues when activating an new ILM policy, you can use the Ingest Time advanced filter in any rule that might change the location of large numbers of existing objects. Set Ingest Time to be greater than or equal to the approximate time when the new policy will go into effect to ensure that existing objects are not moved unnecessarily.



Contact technical support if you need to slow or increase the rate at which objects are processed after an ILM policy change.

Example 7: Compliant ILM policy for S3 Object Lock

You can use the S3 bucket, ILM rules, and ILM policy in this example as a starting point when defining an ILM policy to meet the object protection and retention requirements for objects in buckets with S3 Object Lock enabled.



If you used the legacy Compliance feature in previous StorageGRID releases, you can also use this example to help manage any existing buckets that have the legacy Compliance feature enabled.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate the proposed policy to confirm it will work as intended to protect content from loss.

Related information

- [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)

Buckets

Create buckets and manage bucket settings.

1 bucket

Create bucket

Actions ▾

<input type="checkbox"/>	Name ▾	S3 Object Lock ? ▾	Region ▾	Object Count ? ▾	Space Used ? ▾	Date Created ▾
<input type="checkbox"/>	bank-records	✓	us-east-1	0	0 bytes	2021-01-06 16:53:19 MST

← Previous 1 Next →

Each object and object version that is added to the bank-records bucket will use the following values for `retain-until-date` and `legal hold` settings.

Setting for each object	Example value
<code>retain-until-date</code>	<p>"2030-12-30T23:59:59Z" (December 30, 2030)</p> <p>Each object version has its own <code>retain-until-date</code> setting. This setting can be increased, but not decreased.</p>
<code>legal hold</code>	<p>"OFF" (Not in effect)</p> <p>A legal hold can be placed or lifted on any object version at any time during the retention period. If an object is under a legal hold, the object cannot be deleted even if the <code>retain-until-date</code> has been reached.</p>

ILM rule 1 for S3 Object Lock example: Erasure Coding profile with bucket matching

This example ILM rule applies only to the S3 tenant account named Bank of ABC. It matches any object in the `bank-records` bucket and then uses erasure coding to store the object on Storage Nodes at three data center sites using a 6+3 Erasure Coding profile. This rule satisfies the requirements of buckets with S3 Object Lock enabled: an erasure-coded copy is kept on Storage Nodes from day 0 to forever, using Ingest Time as the reference time.

Rule definition	Example value
Rule Name	Compliant Rule: EC objects in bank-records bucket - Bank of ABC
Tenant Account	Bank of ABC
Bucket Name	bank-records
Advanced filtering	Object Size (MB) greater than 1 Note: This filter ensures that erasure coding is not used for objects 1 MB or smaller.

Create ILM Rule Step 1 of 3: Define Basics

Name

Description

Tenant Accounts (optional)

Bucket Name

[Advanced filtering...](#) (0 defined)

Cancel

Next

Rule definition	Example value
Reference Time	Ingest Time
Placements	From day 0 store forever
Erasure Coding Profile	<ul style="list-style-type: none"> • Create an erasure-coded copy on Storage Nodes at three data center sites • Uses 6+3 erasure-coding scheme

Edit ILM Rule Step 2 of 3: Define Placements

Configure placement instructions to specify how you want objects matched by this rule to be stored.

Compliant Rule: EC objects in bank-record bucket - Bank of ABC

Reference Time
Ingest Time

Placements
Sort by start day

From day 0 store forever
Add Remove

Type erasure coded Location Three Data Centers (6 plus 3) Copies 1
+ x

Retention Diagram
Refresh

Trigger
Day 0
Three Data Centers (6 plus 3)
Duration
Forever

Cancel Back Save

ILM rule 2 for S3 Object Lock example: Non-compliant rule

This example ILM rule initially stores two replicated object copies on Storage Nodes. After one year, it stores one copy on a Cloud Storage Pool forever. Because this rule uses a Cloud Storage Pool, it is not compliant and will not apply to the objects in buckets with S3 Object Lock enabled.

Rule definition	Example value
Rule Name	Non-Compliant Rule: Use Cloud Storage Pool
Tenant Accounts	Not specified
Bucket Name	Not specified, but will only apply to buckets that do not have S3 Object Lock (or the legacy Compliance feature) enabled.
Advanced filtering	Not specified

Create ILM Rule Step 1 of 3: Define Basics

Name
Non-Compliant Rule: Use Cloud Storage Pool

Description
DC1 and 2 for 1 year then move to CSP

Tenant Accounts (optional)
Select tenant accounts or enter tenant IDs

Bucket Name
matches all
Value

Advanced filtering... (0 defined)

Cancel Next

Rule definition	Example value
Reference Time	Ingest Time
Placements	<ul style="list-style-type: none"> • 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
Bucket Name	Not specified
Advanced filtering	Not specified

Create ILM Rule Step 1 of 3: Define Basics

Name

Description

Tenant Accounts (optional)

Bucket Name

[Advanced filtering...](#) (0 defined)

Cancel

Next

Rule definition	Example value
Reference Time	Ingest Time
Placements	From Day 0 to forever, keep two replicated copies—one on Storage Nodes in Data Center 1 and one on Storage Nodes in Data Center 2.

Compliant Rule: Two Copies Two Data Centers

Reference Time
Ingest Time

Placements
Sort by start day

From day 0 store forever
Add Remove

Type replicated Location Data Center 1 Data Center 2 Add Pool Copies 2
+

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

Retention Diagram
Refresh

Trigger
Day 0
Data Center 1
Data Center 2
Duration
Forever

Compliant ILM policy for S3 Object Lock example

To create an ILM policy that will effectively protect all objects in your system, including those in buckets with S3 Object Lock enabled, you must select ILM rules that satisfy the storage requirements for all objects. Then, you must simulate and activate the proposed policy.

Add rules to the policy

In this example, the ILM policy includes three ILM rules, in the following order:

1. A compliant rule that uses erasure coding to protect objects greater than 1 MB in a specific bucket with S3 Object Lock enabled. The objects are stored on Storage Nodes from day 0 to forever.
2. A non-compliant rule that creates two replicated object copies on Storage Nodes for a year and then moves one object copy to a Cloud Storage Pool forever. This rule does not apply to buckets with S3 Object Lock enabled because it uses a Cloud Storage Pool.
3. The default compliant rule that creates two replicated object copies on Storage Nodes from day 0 to forever.

Configure ILM Policy

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

Name Compliant ILM policy for S3 Object Lock example

Reason for change Example policy

Rules

1. Select the rules you want to add to the policy.
2. Determine the order in which the rules will be evaluated by dragging and dropping the rows. The default rule (and any non-compliant rule without a filter) will be automatically placed at the end of the policy and cannot be moved.

+ Select Rules

Default	Rule Name	Compliant	Tenant Account	Actions
	Compliant Rule: EC for bank-records bucket - Bank of ABC 	✓	Bank of ABC (90767802913525281639)	✕
	Non-Compliant Rule: Use Cloud Storage Pool 		Ignore	✕
✓	Default Compliant Rule: Two Copies Two Data Centers 	✓	Ignore	✕

Cancel

Save

Simulate the proposed policy

After you have added rules in your proposed policy, chosen a default compliant rule, and arranged the other rules, you should simulate the policy by testing objects from the bucket with S3 Object Lock enabled and from other buckets. For example, when you simulate the example policy, you would expect test objects to be evaluated as follows:

- The first rule will only match test objects that are 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.

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.

Related information

[Vulnerability Handling Policy](#)

Hardening guidelines for software upgrades

You must keep your StorageGRID system and related services up to date to defend against attacks.

Upgrades to StorageGRID software

Whenever possible, you should upgrade StorageGRID software to the most recent major release or to the previous major release. Keeping StorageGRID up to date helps reduce the amount of time that known vulnerabilities are active and reduces the overall attack surface area. In addition, the most recent releases of StorageGRID often contain security hardening features that are not included in earlier releases.

When a hotfix is required, NetApp prioritizes creating updates for the most recent releases. Some patches might not be compatible with earlier releases.

To download the most recent StorageGRID releases and hotfixes, go to the StorageGRID software download page. For step-by-step instructions for upgrading StorageGRID software, see the instructions for upgrading StorageGRID. For instructions on applying a hotfix, see the recovery and maintenance instructions.

Upgrades to external services

External services can have vulnerabilities that affect StorageGRID indirectly. You should ensure that the services that StorageGRID depends on are kept up to date. These services include LDAP, KMS (or KMIP server), DNS, and NTP.

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

Upgrades to hypervisors

If your StorageGRID nodes are running on VMware or another hypervisor, you must ensure that the hypervisor software and firmware are up to date.

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

Upgrade to Linux nodes

If your StorageGRID nodes are using Linux host platforms, you must ensure that security updates and kernel updates are applied to the host OS. Additionally, you must apply firmware updates to vulnerable hardware when these updates become available.

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

Related information

[NetApp Downloads: StorageGRID](#)

[Upgrade software](#)

[Recover and maintain](#)

[NetApp Interoperability Matrix Tool](#)

Hardening guidelines for StorageGRID networks

The StorageGRID system supports up to three network interfaces per grid node, allowing you to configure the networking for each individual grid node to match your security and access requirements.

Guidelines for 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 in the installation guide for your platform.
- If untrusted clients can access the Admin Network, block access to StorageGRID on the Admin Network with an external firewall.

Guidelines for Client Network

The Client Network is typically used for tenants and for communicating with external services, such as the CloudMirror replication service or another platform service. However, StorageGRID does not enforce this usage internally.

If you are using the Client Network, follow these guidelines:

- Block all internal traffic ports on the Client Network. See the list of internal ports in the installation guide for your platform.
- Accept inbound client traffic only on explicitly configured endpoints. See [Managing untrusted Client Networks](#).

Related information

[Networking guidelines](#)

[Grid primer](#)

[Administer StorageGRID](#)

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

[Install Ubuntu or Debian](#)

[Install VMware](#)

Hardening guidelines for StorageGRID nodes

StorageGRID nodes can be deployed on VMware virtual machines, within a 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.

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 uses an internal firewall that is managed automatically. While this internal firewall provides an additional layer of protection against some common threats, it does not remove the need for an external firewall.

For a list of all internal and external ports used by StorageGRID, see the installation guide for your platform.

Virtualization, containers, and shared hardware

For all StorageGRID nodes, avoid running StorageGRID on the same physical hardware as untrusted software. Do not assume that hypervisor protections will prevent malware from accessing StorageGRID-protected data if both StorageGRID and the malware exist on the same the physical hardware. For example,

the Meltdown and Spectre attacks exploit critical vulnerabilities in modern processors and allow programs to steal data in memory on the same computer.

Disable unused services

For all StorageGRID nodes, you should disable or block access to unused services. For example, if you are not planning to configure client access to the audit shares for CIFS or NFS, block or disable access to these services.

Protect nodes during installation

Do not allow untrusted users to access StorageGRID nodes over the network when the nodes are being installed. Nodes are not fully secure until they have joined the grid.

Guidelines for Admin Nodes

Admin Nodes provide management services such as system configuration, monitoring, and logging. When you sign in to the Grid Manager or the Tenant Manager, you are connecting to an Admin Node.

Follow these guidelines to secure the Admin Nodes in your StorageGRID system:

- Secure all Admin Nodes from untrusted clients, such as those on the open internet. Ensure that no untrusted client can access any Admin Node on the Grid Network, the Admin Network, or the Client Network.
- StorageGRID Groups control access to Grid Manager and Tenant Manager features. Grant each Group of users the minimum required permissions for their role, and use the read-only access mode to prevent users from changing configuration.
- When using StorageGRID load balancer endpoints, use Gateway Nodes instead of Admin Nodes for untrusted client traffic.
- If you have untrusted tenants, do not allow them to have direct access to the Tenant Manager or the Tenant Management API. Instead, have any untrusted tenants use a tenant portal or an external tenant management system, which interacts with the Tenant Management API.
- Optionally, use an Admin proxy for more control over AutoSupport communication from Admin Nodes to NetApp support. See the steps for creating an Admin proxy in the instructions for administering StorageGRID.
- Optionally, use the restricted 8443 and 9443 ports to separate Grid Manager and Tenant Manager communications. Block the shared port 443 and limit tenant requests to port 9443 for additional protection.
- Optionally, use separate Admin Nodes for grid administrators and tenant users.

For more information, see the instructions for administering StorageGRID.

Guidelines for Storage Nodes

Storage Nodes manage and store object data and metadata. Follow these guidelines to secure the Storage Nodes in your StorageGRID system.

- Do not 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.
- Do not enable outbound services for untrusted tenants. For example, when creating the account for an untrusted tenant, do not allow the tenant to use its own identity source and do not allow the use of platform services. See the steps for creating a tenant account in the instructions for administering StorageGRID.

- Use a third-party load balancer for untrusted client traffic. Third-party load balancing offers more control and additional layers of protection against attack.
- Optionally, use a Storage proxy for more control over Cloud Storage Pools and platform services communication from Storage Nodes to external services. See the steps for creating a Storage proxy in the instructions for administering StorageGRID.
- Optionally, connect to external services using the Client Network. Then, select **CONFIGURATION > Network > 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 instead of using the CLB service on Gateway Nodes. See the steps for managing load balancing in the instructions for administering StorageGRID.



The CLB service is deprecated.

- Use a third-party load balancer between the client and the Gateway Node or Storage Nodes for untrusted client traffic. Third-party load balancing offers more control and additional layers of protection against attack. If you do use a third-party load balancer, network traffic can still optionally be configured to go through an internal load balancer endpoint or be sent directly to Storage Nodes.
- If you are using load balancer endpoints, optionally have clients connect over the Client Network. Then, select **CONFIGURATION > Network > 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.
- If the storage controller in the appliance includes FDE or FIPS drives and the Drive Security feature is enabled, use SANtricity to configure Drive Security keys.
- For appliances without FDE or FIPS drives, enable node encryption using a Key Management Server (KMS).

See the installation and maintenance instructions for your StorageGRID hardware appliance.

Related information

- [Install Red Hat Enterprise Linux or CentOS](#)
- [Install Ubuntu or Debian](#)
- [Install VMware](#)
- [Administer StorageGRID](#)
- [Use a tenant account](#)
- [SG100 and SG1000 services appliances](#)
- [SG5600 storage appliances](#)
- [SG5700 storage appliances](#)
- [SG6000 storage appliances](#)

Hardening guidelines for server certificates

You should replace the default certificates created during installation with your own custom certificates.

For many organizations, the self-signed digital certificate for StorageGRID web access is not compliant with their information security policies. On production systems, you should install a CA-signed digital certificate for use in authenticating StorageGRID.

Specifically, you should use custom server certificates instead of these default certificates:

- **Management interface 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.



StorageGRID manages the certificates used for load balancer endpoints separately. To configure load balancer certificates, see the steps for configuring load balancer endpoints in the instructions for administering StorageGRID.

When using custom server certificates, follow these guidelines:

- Certificates should have a `subjectAltName` that matches DNS entries for StorageGRID. For details, see section 4.2.1.6, “Subject Alternative Name,” in [RFC 5280: PKIX Certificate and CRL Profile](#).
- When possible, avoid the use of wildcard certificates. An exception to this guideline is the certificate for an S3 virtual hosted style endpoint, which requires the use of a wildcard if bucket names are not known in advance.
- When you must use wildcards in certificates, you should take additional steps to reduce the risks. Use a wildcard pattern such as `*.s3.example.com`, and do not use the `s3.example.com` suffix for other applications. This pattern also works with path-style S3 access, such as `dc1-s1.s3.example.com/mybucket`.
- Set the certificate expiration times to be short (for example, 2 months), and use the Grid Management API to automate certificate rotation. This is especially important for wildcard certificates.

In addition, clients should use strict hostname checking when communicating with StorageGRID.

Other hardening guidelines

In addition to following the hardening guidelines for StorageGRID networks and nodes, you should follow the hardening guidelines for other areas of the StorageGRID system.

Logs and audit messages

Always protect StorageGRID logs and audit message output in a secure manner. StorageGRID logs and audit messages provide invaluable information from a support and system availability standpoint. In addition, the information and details contained in StorageGRID logs and audit message output are generally of a sensitive nature.

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 instructions for monitoring and troubleshooting for more information about StorageGRID logs. See the instructions for audit messages for more information about StorageGRID audit messages.

NetApp AutoSupport

The AutoSupport feature of StorageGRID allows you to proactively monitor the health of your system and automatically send messages and details to NetApp technical support, your organization's internal support team, or a support partner. By default, AutoSupport messages to NetApp technical support are enabled when StorageGRID is configured for the first time.

The AutoSupport feature can be disabled. However, NetApp recommends enabling it because AutoSupport helps speed problem identification and resolution should an issue arise on your StorageGRID system.

AutoSupport supports HTTPS, HTTP, and SMTP for transport protocols. Because of the sensitive nature of AutoSupport messages, NetApp strongly recommends using HTTPS as the default transport protocol for sending AutoSupport messages to NetApp support.

Optionally, you can configure an Admin proxy for more control over AutoSupport communication from Admin Nodes to NetApp technical support. See the steps for creating an Admin proxy in the instructions for administering StorageGRID.

Cross-Origin Resource Sharing (CORS)

You can configure Cross-Origin Resource Sharing (CORS) for an S3 bucket if you want that bucket and objects in that bucket to be accessible to web applications in other domains. In general, do not enable CORS unless it is required. If CORS is required, restrict it to trusted origins.

See the steps for configuring Cross-Origin Resource Sharing (CORS) in the instructions for using tenant accounts.

External security devices

A complete hardening solution must address security mechanisms outside of StorageGRID. Using additional infrastructure devices for filtering and limiting access to StorageGRID is an effective way to establish and maintain a stringent security posture. These external security devices include firewalls, intrusion prevention systems (IPSs), and other security devices.

A third-party load balancer is recommended for untrusted client traffic. Third-party load balancing offers more control and additional layers of protection against attack.

Related information

[Monitor and troubleshoot](#)

[Review audit logs](#)

[Use tenant account](#)

[Administer StorageGRID](#)

Configure FabricPool

Configure StorageGRID for FabricPool: Overview


If you use NetApp ONTAP software, you can use NetApp FabricPool to tier inactive, or cold, data to a NetApp StorageGRID object storage system.

About these instructions

Use these instructions to:

- Get an overview of configuring a StorageGRID object storage system for use with FabricPool.
- Learn how to obtain the information you provide to ONTAP when you attach StorageGRID as a FabricPool cloud tier.
- Learn the best practices for configuring the StorageGRID information lifecycle management (ILM) policy, a StorageGRID traffic classification policy, and other StorageGRID options for a FabricPool workload.

Configuration workflow



Before you start

- 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 these additional resources, which provide details about using and configuring FabricPool:

- [TR-4598: FabricPool Best Practices in ONTAP 9.9.1](#)
- [ONTAP 9 Documentation](#)

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. Data is stored either on the primary storage media or in the object datastore based on whether it is accessed frequently or not. Using FabricPool-enabled aggregates helps you reduce storage cost without compromising performance, efficiency, or protection.

No architectural changes are required, and you can continue managing your data and application environment from the central ONTAP storage system.

What is StorageGRID?

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 are not nested as files inside a directory inside other directories. Although object storage generally provides lower performance than file or block storage, it is significantly more scalable. StorageGRID buckets can hold petabytes of data and billions of objects.

Why use StorageGRID as a FabricPool cloud tier?

FabricPool can tier ONTAP data to a number of object store providers, including StorageGRID. Unlike public clouds that might set a maximum number of supported input/output operations per second (IOPS) at the bucket or container level, StorageGRID performance scales with the number of nodes in a system. Using StorageGRID as a FabricPool cloud tier allows you to keep your cold data in your own private cloud for highest performance and complete control over your data.

In addition, a FabricPool license is not required when you use StorageGRID as the cloud tier.

Can I use multiple ONTAP clusters with StorageGRID?

These instructions describe how to connect StorageGRID to a single ONTAP cluster. However, you might want to connect the same StorageGRID system to multiple ONTAP clusters.

The only requirement for tiering data from multiple ONTAP clusters to a single StorageGRID system is that you must use a different S3 bucket for each cluster. Based on your requirements, you can use the same high availability (HA) group, load balancer endpoint, and tenant account for all clusters, or you can configure each of these items for each cluster.

Attach StorageGRID as a 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 some configuration steps in StorageGRID and obtain certain values.

About this task

The following table lists the information you must provide to ONTAP when you attach StorageGRID as a cloud tier for FabricPool. The topics in this section explain how to use the StorageGRID Grid Manager and Tenant Manager to obtain the information you need.



The exact field names listed and the process you use to enter the required values in ONTAP depend on whether you are using the ONTAP CLI (storage aggregate object-store config create) or ONTAP System Manager (**Storage > Aggregates & Disks > Cloud Tier**).

For more information, refer to the following:

- [TR-4598: FabricPool Best Practices in ONTAP 9.9.1](#)
- [ONTAP 9 Documentation](#)

ONTAP field	Description
Object store name	Any unique and descriptive name. For example, StorageGRID_Cloud_Tier.
Provider type	StorageGRID (ONTAP System Manager) or SGWS (ONTAP CLI).
Port	<p>The port that FabricPool will use when it connects to StorageGRID. You determine which port number to use when you define the StorageGRID load balancer endpoint.</p> <p>Create a load balancer endpoint for FabricPool</p>
Server name	<p>The fully qualified domain name (FQDN) for the StorageGRID load balancer endpoint. For example, s3.storagegrid.company.com.</p> <p>Note the following:</p> <ul style="list-style-type: none">• The domain name that you specify here must match the domain name on the CA certificate you upload for the StorageGRID load balancer endpoint.• The DNS record for this domain name must map to each IP address you will use to connect to StorageGRID. <p>Configure the DNS server for StorageGRID IP addresses</p>
Container name	<p>The name of the StorageGRID bucket you will use with this ONTAP cluster. For example, fabricpool-bucket. You can create this bucket in the Tenant Manager or, starting with ONTAP 9.10 System Manager, you can create the bucket with the FabricPool setup wizard.</p> <p>Note the following:</p> <ul style="list-style-type: none">• The bucket name cannot be changed once the configuration is created.• The bucket cannot have versioning enabled.• You must use a different bucket for each ONTAP cluster that will tier data to StorageGRID. <p>Create an S3 bucket and obtain an access key</p>

ONTAP field	Description
Access key and secret password	<p>The access key and secret access key for the StorageGRID tenant account.</p> <p>You generate these values in the Tenant Manager.</p> <p>Create an S3 bucket and obtain an access key</p>
SSL	Must be enabled.
Object store certificate	<p>The CA certificate you uploaded when you created the StorageGRID load balancer endpoint.</p> <p>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.</p> <p>Create a load balancer endpoint for FabricPool</p>

After you finish

After you have obtained the required StorageGRID information, you can go to ONTAP to add StorageGRID as a cloud tier, add the cloud tier as an aggregate, and set volume tiering policies.

Best practices for load balancing

Before attaching StorageGRID as a FabricPool cloud tier, you must use the StorageGRID Grid Manager to configure at least one load balancer endpoint.

What is load balancing?

When data is tiered from FabricPool to a StorageGRID system, StorageGRID uses a load balancer to manage the ingest and retrieval workload. Load balancing maximizes speed and connection capacity by distributing the FabricPool workload across multiple Storage Nodes.

The StorageGRID Load Balancer service is installed on all Admin Nodes and all Gateway Nodes and provides Layer 7 load balancing. It performs Transport Layer Security (TLS) termination of client requests, inspects the requests, and establishes new secure connections to the Storage Nodes.

The Load Balancer service on each node operates independently when forwarding client traffic to the Storage Nodes. Through a weighting process, the Load Balancer service routes more requests to Storage Nodes with higher CPU availability.

Although the StorageGRID Load Balancer service is the recommended load balancing mechanism, you might want to integrate a third-party load balancer instead. For information, contact your NetApp account representative or refer to [TR-4626: StorageGRID third-party and global load balancers](#).



The separate Connection Load Balancer (CLB) service on Gateway Nodes is deprecated and no longer recommended for use with FabricPool.

Best practices for StorageGRID load balancing

As a general best practice, each site in your StorageGRID system should include two or more nodes with the Load Balancer service. For example, a site might include 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 SG100 or SG1000 services appliances, bare metal nodes, or virtual machine (VM) based nodes.

You must configure a StorageGRID load balancer endpoint to define the port that Gateway Nodes and Admin Nodes will use for incoming and outgoing FabricPool requests.

Best practices for the load balancer endpoint certificate

When creating a load balancer endpoint for use with FabricPool, you should use HTTPS as the protocol. Communicating with StorageGRID without TLS encryption is supported but not recommended.

You can then either upload a certificate that is signed by either a publicly trusted or a private certificate authority (CA), or you can generate a self-signed certificate. The certificate allows ONTAP to authenticate with StorageGRID.

As a best practice, you should use a CA server certificate to secure the connection. Certificates signed by a CA can be rotated nondisruptively.

When requesting a CA certificate for use with the load balancer endpoint, ensure that the domain name on the certificate matches the server name you enter in ONTAP for that load balancer endpoint. If possible, use a wildcard (*) to allow for virtual-host-style URLs. For example:

```
*.s3.storagegrid.company.com
```

When you add StorageGRID as a FabricPool cloud tier, you must install the same certificate to the ONTAP cluster, as well as the root and any subordinate certificate authority (CA) certificates.



StorageGRID uses server certificates for a number of purposes. If you are connecting to the Load Balancer service, you can optionally use the S3 and Swift API certificate.

To learn more about the server certificate for a load balancing endpoint:

- [Configure load balancer endpoints](#)
- [Hardening guidelines for server certificates](#)

Best practices for high availability groups

Before attaching StorageGRID as a FabricPool cloud tier, you should use the StorageGRID Grid Manager to configure a high availability (HA) group.

What is a high availability (HA) group?

To ensure that the Load Balancer service is always available to manage FabricPool data, you can group the network interfaces of multiple Admin and Gateway Nodes into a single entity, known as a high availability (HA) group. If the active node in the HA group fails, another node in the group can continue to manage the workload.

Each HA group provides highly available access to the shared services on the associated nodes. For example, an HA group 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 create an HA group, you perform these general steps:

1. Select network interfaces for one or more Admin Nodes or Gateway Nodes. You can select the Grid Network interface (eth0), Client Network interface (eth2), or a VLAN interface.



If you plan to use a VLAN interface to segregate FabricPool traffic, a network administrator must first configure a trunk interface and the corresponding VLAN. Each VLAN is identified by a numeric ID or tag. For example, your network might use VLAN 100 for FabricPool traffic.

2. Assign one or more virtual IP (VIP) addresses to the group. Clients applications, such as FabricPool, can use any of these VIP addresses to connect to StorageGRID.
3. Specify one interface to be the Primary interface and determine the priority order for any Backup interfaces. The Primary interface is the active interface unless a failure occurs.

If the HA group includes more than one interface and the Primary interface fails, the 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. This failover process generally takes only a few seconds and is fast enough that client applications should experience little impact and can rely on normal retry behaviors to continue operation.

When the 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.

Best practices for high availability (HA) groups

The best practices for creating a StorageGRID HA group for FabricPool depend on the workload, as follows:

- If you plan to use FabricPool with primary workload data, you must create 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](#).

Configure the DNS server for StorageGRID IP addresses

After configuring high availability groups and load balancer endpoints, you must ensure that the domain name system (DNS) for the ONTAP system includes a record to associate the StorageGRID server name (fully qualified domain name) to the IP address that FabricPool will use to make connections.

The IP address you enter in the DNS record depends on whether you are using an HA group of load-balancing nodes:

- If you have configured an HA group, FabricPool will connect to the virtual IP addresses of that HA group.

- If you are not using an HA group, FabricPool can connect to the StorageGRID Load Balancer service using the IP address of any Gateway Node or Admin Node.

You must also ensure that the DNS record references all required endpoint domain names, including any wildcard names.

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 consists of one or more network interfaces on Admin Nodes, Gateway Nodes, or both.

What you'll need

- 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](#).

About this task

Each HA group uses virtual IP addresses (VIPs) to provide highly available access to the shared services on the associated nodes.

For details about this task, see [Manage high availability groups](#).

Steps

1. Select **CONFIGURATION > Network > High availability groups**.
2. Select **Create**.
3. Enter a unique name and optionally a description.
4. 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.

5. Determine the Primary interface and any backup interfaces for this HA group.

Drag and drop 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 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. 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.

7. Optionally, if the ONTAP IP addresses used to access StorageGRID are not 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.

8. Enter one or more virtual IP addresses for the HA group. You can add up to 10 IP addresses. All VIPs must be within the VIP subnet.

You must provide at least one IPv4 address. Optionally, you can specify additional IPv4 and IPv6 addresses.

9. Select **Create HA group** and then select **Finish**.

Create a load balancer endpoint for FabricPool

When configuring StorageGRID for use with FabricPool, you must configure a load balancer endpoint and upload the load balancer endpoint certificate, which is used to secure the connection between ONTAP and StorageGRID.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the Root access permission.
- You have the following files:
 - Server Certificate: The custom server certificate file.
 - Server Certificate Private Key: The custom server certificate private key file.
 - CA Bundle: A single 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.

About this task

For details about this task, see [Configure load balancer endpoints](#).

Steps

1. Select **CONFIGURATION > Network > Load balancer endpoints**.
2. Select **Create**.

×

Create a load balancer endpoint

1 Enter endpoint details

2 Select binding mode

3 Attach certificate

Endpoint details

Name ?

Port ?

Enter an unused port or accept the suggested port.

10443

Client type ?

Select the type of client application that will use this endpoint.

☒ S3
 ☐ Swift

Network protocol ?

Select the network protocol clients will use with this endpoint. If you select HTTPS, attach the security certificate before saving the endpoint.

☐ HTTPS (recommended)
 ☒ HTTP

Cancel

Continue

3. Enter endpoint details.

Field	Description
Name	A descriptive name for the endpoint
Port	<p>The StorageGRID port you want to use for load balancing. This field defaults to 10433, but you can enter any unused external port. If you enter 80 or 443, the endpoint is configured only on Gateway Nodes, since these ports are reserved on Admin Nodes.</p> <p>Note: Ports used by other grid services are not permitted. See the Network port reference.</p> <p>You must provide this same port number to ONTAP when you attach StorageGRID as a FabricPool cloud tier.</p>
Client type	Select S3 .
Network protocol	<p>Select HTTPS.</p> <p>Note: Using HTTP is supported but not recommended.</p>

4. Select **Continue**.
5. Specify the binding mode.

Use the **Global** setting (recommended) or restrict the accessibility of this endpoint to one of the following:

- Specific network interfaces of specific nodes.
- Specific high availability (HA) virtual IP addresses (VIPs). Use this selection only if you require much higher levels of isolation of workloads.

6. Select **Continue**.
7. Select **Upload Certificate** (recommended) and then browse to your server certificate, certificate private key, and optional CA bundle.
8. 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.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

About this task

Tenant accounts allow client applications to store and retrieve objects on StorageGRID. Each tenant account has its own account ID, authorized groups and users, buckets, and objects.

You can use the same tenant account for multiple ONTAP clusters. Or, you can create a dedicated tenant account for each ONTAP cluster as required.



These instructions assume that you have configured single sign-on (SSO) for the Grid Manager. If SSO is not enabled, use [these instructions to create a tenant account](#) instead.

Steps

1. Select **TENANTS**.
2. Select **Create**.
3. Enter a display name and a description.
4. Select **S3**.
5. Leave the **Storage quota** field blank.
6. Select **Allow platform services** to enable the use of platform services.

If platform services are enabled, a tenant can use features, such as CloudMirror replication, that access external services.

7. Do not select **Use own identity source**.
8. Do not select **Allow S3 Select**.

9. Select an existing federated group from the Grid Manager to have the initial Root Access permission for the tenant.
10. Select **Create tenant**.

Create an S3 bucket and obtain an access key

Before using StorageGRID with a FabricPool workload, you must create an S3 bucket for your FabricPool data. You also need to obtain an access key and secret access key for the tenant account you will use for FabricPool.

What you'll need

- You have created a tenant account for FabricPool use.

About this task

These instructions describe how to use the StorageGRID Tenant Manager to create a bucket and obtain access keys. You can also perform these tasks using the Tenant Management API or the StorageGRID S3 REST API. Or, if you are using ONTAP 9.10, you can create the bucket using the FabricPool setup wizard instead.

To learn more:

- [Use a tenant account](#)
- [Use S3](#)

Steps

1. Sign in to the Tenant Manager.

You can do either of the following:

- From the Tenant Accounts page in the Grid Manager, select the **Sign in** link for the tenant, and enter your credentials.
- Enter the URL for the tenant account in a web browser, and enter your credentials.

2. Create an S3 bucket for FabricPool data.

You must create a unique bucket for each ONTAP cluster you plan to use.

- a. Select **STORAGE (S3) > Buckets**.
- b. Select **Create bucket**.
- c. Enter the name of the StorageGRID bucket you will use with FabricPool. For example, `fabricpool-bucket`.



You cannot change the bucket name after creating the bucket.

Bucket names must comply with these rules:

- Must be unique across each StorageGRID system (not just unique within the tenant account).
- Must be DNS compliant.
- Must contain at least 3 and no more than 63 characters.

- Can be a series of one or more labels, with adjacent labels separated by a period. Each label must start and end with a lowercase letter or a number and can only use lowercase letters, numbers, and hyphens.
- Must not look like a text-formatted IP address.
- Should not use periods in virtual hosted style requests. Periods will cause problems with server wildcard certificate verification.

d. Select the region for this bucket.

By default, all buckets are created in the `us-east-1` region.

Create bucket
×

Enter bucket details

Enter the bucket's name and select the bucket's region.

Bucket name ?

fabricpool-bucket

Region ?

us-east-1 ▼

Cancel
Create bucket

e. Select **Create bucket**.



For FabricPool buckets, the recommended bucket consistency level is **Read-after-new-write**, which is the default setting for a new bucket. Do not edit FabricPool buckets to use **Available** or any other consistency level.

3. Create an access key and a secret access key.

- Select **STORAGE (S3) > My access keys**.
- Select **Create key**.
- Select **Create access key**.
- Copy the access key ID and the secret access key to a safe location, or select **Download .csv** to save a spreadsheet file containing the access key ID and secret access key.

You will enter these values in ONTAP when you configure StorageGRID as a FabricPool cloud tier.



If you create a new access key and secret access key in the future, remember to update the corresponding values in ONTAP immediately to ensure that ONTAP can store and retrieve data in StorageGRID without interruption.

Use StorageGRID information lifecycle management with FabricPool data

If you are using FabricPool to tier data to StorageGRID, you must understand the requirements for creating StorageGRID information lifecycle management (ILM) rules and an ILM policy to manage FabricPool data. You must ensure the ILM rules that apply to FabricPool data are not disruptive.



FabricPool has no knowledge of StorageGRID ILM rules or policies. Data loss can occur if the StorageGRID ILM policy is misconfigured. See [Manage objects with ILM](#) for detailed ILM instructions.

Review these guidelines to ensure that your ILM rules and ILM policy are suitable for FabricPool data and your business requirements. If you are already using StorageGRID ILM, you might need to update your active ILM policy to meet these guidelines.

- You can use any combination of replication and erasure-coding rules to protect cloud tier data.

The recommended best practice is to use 2+1 erasure coding within a site for cost-efficient data protection. Erasure coding uses more CPU, but 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.

- Do not use an ILM rule that will expire or delete FabricPool cloud tier data. Set the retention period in each ILM rule to "forever" to ensure that FabricPool objects are not deleted by StorageGRID ILM.
- Do not create rules that will move FabricPool cloud tier data out of the bucket to another location. You cannot use ILM rules to archive FabricPool data to tape using an Archive Node or use a Cloud Storage Pool to move FabricPool data to another object store.



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

- Starting with ONTAP 9.8, you can optionally create object tags to help classify and sort tiered data for easier management. For example, you can set tags only on FabricPool volumes attached to StorageGRID. Then, when you create ILM rules in StorageGRID, you can use the Object Tag advanced filter to select and place this data.

Example ILM policy for FabricPool data

Use this simple example policy as a starting point for your own ILM rules and policy.

This example assumes you are designing the ILM rules and an ILM policy for a StorageGRID system that has four Storage Nodes at a single data center in Denver, Colorado. The FabricPool data in this example uses a bucket named `fabricpool-bucket`.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate the proposed policy to confirm it will work as intended to protect content from loss. To learn more, see [Manage objects with ILM](#).

Steps

1. Create a storage pool named **DEN**. Select the Denver site.
2. Create an Erasure Coding profile named **2 plus 1**. Select the 2+1 erasure-coding scheme and the **DEN** storage pool.
3. Create an ILM rule that applies only to the data in `fabricpool-bucket`. This example rule creates erasure-coded copies.

Rule definition	Example value
Rule Name	2 plus 1 erasure coding for FabricPool data
Bucket Name	<code>fabricpool-bucket</code> You could also filter on the FabricPool tenant account.
Advanced Filtering	Object Size (MB) greater than 0.2 MB. 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
Placement	From day 0 store forever
Type	Erasure coded
Location	DEN (2 plus 1)
Ingest Behavior	Balanced

4. Create an ILM rule that will create two replicated copies of any objects not matched by the first rule. Do not select a basic filter (tenant account or bucket name) or any advanced filters.

Rule definition	Example value
Rule Name	Two replicated copies
Bucket Name	<i>none</i>
Advanced Filtering	<i>none</i>
Reference Time	Ingest Time

Rule definition	Example value
Placement	From day 0 store forever
Type	Replicated
Location	DEN
Copies	2
Ingest Behavior	Balanced

5. Create a proposed ILM policy and select the two rules. Because the replication rule does not use any filters, it can be the default (last) rule for the policy.
6. Ingest test objects into the grid.
7. Simulate the policy with the test objects to verify the behavior.
8. Activate the policy.

When this policy is activated, StorageGRID places object data as follows:

- The data tiered from FabricPool in `fabricpool-bucket` will be erasure coded using the 2+1 erasure-coding scheme. Two data fragments and one parity fragment will be placed on three different Storage Nodes.
- All objects in all other buckets will be replicated. Two copies will be created and placed on two different Storage Nodes.
- The erasure-coded and replicated copies will be maintained in StorageGRID until they are deleted by the S3 client. StorageGRID ILM will never delete these items.

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.

What you'll need

- 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 the majority of bandwidth. You can create a traffic classification policy to limit all other workloads.



In general, FabricPool read operations are more important to prioritize than write operations.

For example, if other S3 clients use this StorageGRID system, you should create a traffic classification policy. You can limit network traffic for the other buckets, tenants, IP subnets, or load balancer endpoints.

- As a general rule, you should not impose quality of service limits on any FabricPool workload; you should only limit the other workloads.
- The limits placed on other workloads 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.

To learn more: [Manage traffic classification policies](#)

Steps

1. Select **CONFIGURATION > Network > Traffic classification**.
2. Enter a name and a description.
3. In the Matching Rules section, create at least one rule.
 - a. Select **Create**.
 - b. Select **Endpoint**, and select the load balancer endpoint you created for FabricPool.

You can also select the FabricPool tenant account or bucket.

- c. If you want this traffic policy to limit traffic for the other endpoints, select **Inverse Match**.
4. Optionally, create one or more limits.



Even if no limits are set for a traffic classification policy, metrics are collected so you can understand traffic trends.

- a. Select **Create**.
- b. Select the type of traffic you want to limit and the limit to apply.

This example FabricPool traffic classification policy shows the types of network traffic you can limit and the types of values you can select. The limits for an actual policy would be based on your specific requirements.

Policy

Name

Description (optional)

Matching Rules

Traffic that matches any rule is included in the policy.

+ Create
Edit
Remove

Type	Inverse Match	Match Value
<input checked="" type="radio"/> Endpoint	<input checked="" type="checkbox"/>	FabricPool (https 10443)

Displaying 1 matching rule.

Limits (Optional)

+ Create
Edit
Remove

Type	Value	Units
<input type="radio"/> Concurrent Read Requests	50	Concurrent Requests
<input type="radio"/> Concurrent Read Requests	15	Concurrent Requests
<input type="radio"/> Read Request Rate	100	Requests/Second
<input type="radio"/> Write Request Rate	25	Requests/Second
<input type="radio"/> Per-Request Bandwidth In	2000000	Bytes/Second
<input checked="" type="radio"/> Per-Request Bandwidth Out	10000000	Bytes/Second

- After creating the traffic classification policy, select the policy and then select **Metrics** to determine if the policy is limiting traffic as expected.

Traffic Classification Policies

Traffic classification policies can be used to identify network traffic for metrics reporting and optional traffic limiting.

+ Create Edit Remove Metrics		
Name	Description	ID
<input checked="" type="radio"/> FabricPool	Limit traffic other than FabricPool	587f53b2-7cf2-44b9-af5c-694ebbd4a2c5
Displaying 1 traffic classification policy.		

Other best practices for StorageGRID and FabricPool

When configuring a StorageGRID system for use with FabricPool, you should avoid setting global options that might affect how your data is saved.

Object encryption

When configuring StorageGRID, you can optionally enable the global **Stored Object Encryption** setting if data encryption is required for other StorageGRID clients (**CONFIGURATION > System > Grid options**). The data that is tiered from FabricPool to StorageGRID is already encrypted, so enabling the StorageGRID setting is not required. Client-side encryption keys are owned by ONTAP.

Object compression

When configuring StorageGRID, do not enable the global **Compress Stored Objects** setting (**CONFIGURATION > System > Grid options**). The data that is tiered from FabricPool to StorageGRID is already compressed. Enabling **Compress Stored Objects** will not further reduce an object's size.

Consistency level

For FabricPool buckets, the recommended bucket consistency level is **Read-after-new-write**, which is the default setting for a new bucket. Do not edit FabricPool buckets to use **Available** or any other consistency level.

FabricPool tiering

If the StorageGRID node uses storage assigned from a NetApp 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.

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.