



Use a tenant account

StorageGRID

NetApp

October 03, 2025

Table of Contents

Use a tenant account	1
Use a tenant account: Overview	1
What is a tenant account?	1
How to create a tenant account	1
Use the Tenant Manager	2
How to sign in and sign out	2
Sign in to Tenant Manager	2
Sign out of Tenant Manager	5
Understand Tenant Manager Dashboard	6
Tenant account summary	7
Storage and quota usage	7
Quota usage alerts	9
Endpoint errors	9
Tenant Management API	9
Understand Tenant Management API	9
Tenant Management API versioning	12
Protect against Cross-Site Request Forgery (CSRF)	13
Manage system access	14
Use identity federation	14
Manage groups	19
Manage local users	32
Manage S3 tenant accounts	34
Manage S3 access keys	34
Manage S3 buckets	45
Manage S3 platform services	63
What are platform services?	63
Considerations for using platform services	68
Configure platform services endpoints	70
Configure CloudMirror replication	87
Configure event notifications	91
Use search integration service	95

Use a tenant account

Use a tenant account: Overview

A tenant account allows you to use either the Simple Storage Service (S3) REST API or the Swift REST API to store and retrieve objects in a StorageGRID system.

What is a tenant account?

Each tenant account has its own federated or local groups, users, S3 buckets or Swift containers, and objects.

Optionally, tenant accounts can be used to segregate stored objects by different entities. For example, multiple tenant accounts can be used for either of these use cases:

- **Enterprise use case:** If the StorageGRID system is being used within an enterprise, the grid's object storage might be segregated by the different departments in the organization. For example, there might be tenant accounts for the Marketing department, the Customer Support department, the Human Resources department, and so on.



If you use the S3 client protocol, you can also use S3 buckets and bucket policies to segregate objects between the departments in an enterprise. You do not need to create separate tenant accounts. See the [instructions for implementing S3 client applications](#).

- **Service provider use case:** If the StorageGRID system is being used by a service provider, the grid's object storage might be segregated by the different entities that lease the storage. For example, there might be tenant accounts for Company A, Company B, Company C, and so on.

How to create a tenant account

Tenant accounts are created by a [StorageGRID grid administrator](#) using the Grid Manager. When creating a tenant account, the grid administrator specifies the following information:

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

In addition, grid administrators can enable the S3 Object Lock setting for the StorageGRID system if S3 tenant accounts need to comply with regulatory requirements. When S3 Object Lock is enabled, all S3 tenant accounts can create and manage compliant buckets.

Configure S3 tenants

After an [S3 tenant account is created](#), you can access the Tenant Manager to perform tasks such as the following:

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



While you can create and manage S3 buckets with the Tenant Manager, you must have [S3 access keys and use the S3 REST API to ingest and manage objects](#).

Configure Swift tenants

After a [Swift tenant account is created](#), you can access the Tenant Manager to perform tasks such as the following:

- Setting up identity federation (unless the identity source is shared with the grid), 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.

Use the Tenant Manager

The Tenant Manager allows you to manage all aspects of a StorageGRID tenant account.

You can use the Tenant Manager to monitor a tenant account's storage usage and to manage users with identity federation or by creating local groups and users. For S3 tenant accounts, you can also manage S3 keys, manage S3 buckets, and configure platform services.

How to sign in and sign out

Sign in to Tenant Manager

You access the Tenant Manager by entering the URL for the tenant into the address bar of a [supported web browser](#).

What you'll need

- You must have your login credentials.
- You must have a URL for accessing the Tenant Manager, as supplied by your grid administrator. The URL will look like one of these examples:

`https://FQDN_or_Admin_Node_IP/`

`https://FQDN_or_Admin_Node_IP:port/`

`https://FQDN_or_Admin_Node_IP/?accountId=20-digit-account-id`

`https://FQDN_or_Admin_Node_IP:port/?accountId=20-digit-account-id`

The URL always contains either the fully qualified domain name (FQDN) or the IP address used to access an Admin Node, and could optionally also include a port number, the 20-digit tenant account ID, or both.


- If the URL does not include the tenant's 20-digit account ID, you must have this account ID.
- You must be using a [supported web browser](#).
- Cookies must be enabled in your web browser.
- You must have specific access permissions.

Steps


1. Launch a [supported web browser](#).
2. In the browser's address bar, enter the URL for accessing Tenant Manager.
3. If you are prompted with a security alert, install the certificate using the browser's installation wizard.
4. Sign in to the Tenant Manager.

The sign-in screen that you see depends on the URL you entered and whether your organization is using single sign-on (SSO). You will see one of the following screens:

- The Grid Manager sign-in page. Click the **Tenant Login** link in the upper right.



- The Tenant Manager sign-in page. The **Account ID** field might already be completed, as shown below.



- i. If the tenant's 20-digit account ID is not shown, select the name of the tenant account if it appears in the list of recent accounts, or enter the account ID.
- ii. Enter your username and password.
- iii. Click **Sign in**.

The Tenant Manager Dashboard appears.

- Your organization's SSO page, if SSO is enabled on the grid. For example:

A screenshot of a sign-in page for an organizational account. The page has a white background with a thin black border. At the top, the text "Sign in with your organizational account" is centered. Below this are two input fields: the top one contains the email "someone@example.com" and the bottom one is labeled "Password". At the bottom is a blue "Sign in" button.

Enter your standard SSO credentials, and click **Sign in**.

- The Tenant Manager SSO sign-in page.



- i. If the tenant's 20-digit account ID is not shown, select the name of the tenant account if it appears in the list of recent accounts, or enter the account ID.
- ii. Click **Sign in**.
- iii. Sign in with your standard SSO credentials on your organization's SSO sign-in page.

The Tenant Manager Dashboard appears.

5. If you received an initial password from someone else, change your password to secure your account. Select **username > Change Password**.




If SSO is enabled for the StorageGRID system, you cannot change your password from the Tenant Manager.

Sign out of Tenant Manager

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

Steps

1. Locate the username drop-down in the top-right corner of the user interface.



2. Select the username and then select **Sign Out**.

- If SSO is not in use:

You are signed out of the Admin Node. The Tenant Manager sign in page is displayed.



If you signed into more than one Admin Node, you must sign out of each node.

- If SSO is enabled:

You are signed out of all Admin Nodes you were accessing. The StorageGRID Sign in page is displayed. The name of the tenant account you just accessed is listed as the default in the **Recent Accounts** drop-down, and the tenant's **Account ID** is shown.



If SSO is enabled and you are also signed in to the Grid Manager, you must also sign out of the Grid Manager to sign out of SSO.

Understand Tenant Manager Dashboard

The Tenant Manager Dashboard provides an overview of a tenant account's configuration and the amount of space used by objects in the tenant's buckets (S3) or containers (Swift). If the tenant has a quota, the Dashboard shows how much of the quota is used and how much is remaining. If there are any errors related to the tenant account, the errors are shown on the Dashboard.



The Space used values are estimates. These estimates are affected by the timing of ingest, network connectivity, and node status.

When objects have been uploaded, the Dashboard looks like the following example:

Dashboard

16 Buckets
[View buckets](#)

2 Platform services endpoints
[View endpoints](#)

0 Groups
[View groups](#)

1 User
[View users](#)

Storage usage [?](#)

6.5 TB of 7.2 TB used

0.7 TB (10.1%) remaining



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

Total objects

8,418,886
objects

Tenant details [?](#)

Name:	Tenant02
ID:	3341 1240 0546 8283 2208
<input checked="" type="checkbox"/> Platform services enabled	
<input checked="" type="checkbox"/> Can use own identity source	
<input checked="" type="checkbox"/> S3 Select enabled	

Tenant account summary

The top of the Dashboard contains the following information:

- The number of configured buckets or containers, groups, and users
- The number of platform services endpoints, if any have been configured

You can select the links to view the details.

The right side of the Dashboard contains the following information:

- The total number of objects for the tenant.

For an S3 account, if no objects have been ingested and you have the Root Access permission, getting started guidelines appear instead of the total number of objects.

- Tenant details, including the tenant account name and ID and whether the tenant can use [platform services](#), [its own identity source](#), or [S3 Select](#) (only the permissions that are enabled are listed).

Storage and quota usage

The Storage usage panel contains the following information:

- The amount of object data for the tenant.



This value indicates the total amount of object data uploaded and does not represent the space used to store copies of those objects and their metadata.

- If a quota is set, the total amount of space available for object data and the amount and percentage of space remaining. The quota limits the amount of object data that can be ingested.



Quota utilization is based on internal estimates and might be exceeded in some cases. For example, StorageGRID checks the quota when a tenant starts uploading objects and rejects new ingests if the tenant has exceeded the quota. However, StorageGRID does not take into account the size of the current upload when determining if the quota has been exceeded. If objects are deleted, a tenant might be temporarily prevented from uploading new objects until the quota utilization is recalculated. Quota utilization calculations can take 10 minutes or longer.

- A bar chart that represents the relative sizes of the largest buckets or containers.

You can place your cursor over any of the chart segments to view the total space consumed by that bucket or container.



- To correspond with the bar chart, a list of the largest buckets or containers, including the total amount of object data and the number of objects for each bucket or container.

Bucket name	Space used	Number of objects
Bucket-02	944.7 GB	7,575
Bucket-09	899.6 GB	589,677
Bucket-15	889.6 GB	623,542
Bucket-06	846.4 GB	648,619
Bucket-07	730.8 GB	808,655
Bucket-04	700.8 GB	420,493
Bucket-11	663.5 GB	993,729
Bucket-03	656.9 GB	379,329
9 other buckets	2.3 TB	5,171,588

If the tenant has more than nine buckets or containers, all other buckets or containers are combined into a single entry at the bottom of the list.

Quota usage alerts

If quota usage alerts have been enabled in the Grid Manager, they will appear in the Tenant Manager when the quota is low or exceeded, as follows:

If 90% or more of a tenant's quota has been used, the **Tenant quota usage high** alert is triggered. For more information, see the alerts reference in the instructions for monitoring and troubleshooting StorageGRID.

 Only 0.6% of the quota is remaining. If the quota is exceeded, you can no longer upload new objects.

If you exceed your quota, you cannot upload new objects.

 The quota has been met. You cannot upload new objects.



To view additional details and manage rules and notifications for alerts, see the instructions for monitoring and troubleshooting StorageGRID.

Endpoint errors

If you have used the Grid Manager to configure one or more endpoints for use with platform services, the Tenant Manager Dashboard displays an alert if any endpoint errors have occurred within the past seven days.

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

To see details about an endpoint error, select Endpoints to display the Endpoints page.

Related information

[Troubleshoot platform services endpoint errors](#)

[Monitor and troubleshoot](#)

Tenant Management API

Understand Tenant Management API

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

The Tenant Management API:

- Uses the Swagger open source API platform. Swagger provides an intuitive user interface that allows developers and non-developers to interact with the API. The Swagger user interface provides complete details and documentation for each API operation.
- Uses [versioning to support non-disruptive upgrades](#).

To access the Swagger documentation for the Tenant Management API:

Steps

1. Sign in to the Tenant Manager.
2. From the top of the Tenant Manager, select the help icon and select **API Documentation**.

API operations

The Tenant Management API organizes the available API operations into the following sections:

- **account** — Operations on the current tenant account, including getting storage usage information.
- **auth** — Operations to perform user session authentication.

The Tenant Management API supports the Bearer Token Authentication Scheme. For a tenant login, you provide a username, password, and accountId in the JSON body of the authentication request (that is, `POST /api/v3/authorize`). If the user is successfully authenticated, a security token is returned. This token must be provided in the header of subsequent API requests ("Authorization: Bearer token").

For information about improving authentication security, see [Protect against Cross-Site Request Forgery](#).



If single sign-on (SSO) is enabled for the StorageGRID system, you must perform different steps to authenticate. See the [instructions for using the Grid Management API](#).

- **config** — Operations related to the product release and versions of the Tenant Management API. You can list the product release version and the major versions of the API supported by that release.
- **containers** — Operations on S3 buckets or Swift containers, as follows:

S3

- Create bucket (with and without S3 Object Lock enabled)
- Modify bucket default retention (for buckets with S3 Object Lock enabled)
- Set the consistency control for operations performed on objects
- Create, update, and delete a bucket's CORS configuration
- Enable and disable last access time updates for objects
- Manage the configuration settings for platform services, including CloudMirror replication, notifications, and search integration (metadata-notification)
- Delete empty buckets

Swift: Set the consistency level used for containers

- **deactivated-features** — Operations to view features that might have been deactivated.
- **endpoints** — Operations to manage an endpoint. Endpoints allow an S3 bucket to use an external service for StorageGRID CloudMirror replication, notifications, or search integration.
- **groups** — Operations to manage local tenant groups and to retrieve federated tenant groups from an external identity source.
- **identity-source** — Operations to configure an external identity source and to manually synchronize federated group and user information.
- **regions** — Operations to determine which regions have been configured for the StorageGRID system.

- **s3** — Operations to manage S3 access keys for tenant users.
- **s3-object-lock** — Operations on global S3 Object Lock settings, used to support regulatory compliance.
- **users** — Operations to view and manage tenant users.

Operation details

When you expand each API operation, you can see its HTTP action, endpoint URL, a list of any required or optional parameters, an example of the request body (when required), and the possible responses.

groups Operations on groups

GET /org/groups Lists Tenant User Groups

Parameters

Try it out

Name	Description
type string (query)	filter by group type
limit integer (query)	maximum number of results
marker string (query)	marker-style pagination offset (value is Group's URN)
includeMarker boolean (query)	if set, the marker element is also returned
order string (query)	pagination order (desc requires marker)

Responses

Response content type **application/json**

Code	Description
200	Example Value/ Model <pre>{ "responseTime": "2018-02-01T16:22:31.066Z", "status": "success", "apiVersion": "2.0" }</pre>

Issue API requests



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

Steps

1. Select the HTTP action to see the request details.
2. Determine if the request requires additional parameters, such as a group or user ID. Then, obtain these values. You might need to issue a different API request first to get the information you need.
3. Determine if you need to modify the example request body. If so, you can select **Model** to learn the requirements for each field.
4. Select **Try it out**.
5. Provide any required parameters, or modify the request body as required.
6. Select **Execute**.
7. Review the response code to determine if the request was successful.

Tenant Management API versioning

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

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

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

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

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

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

Outdated requests are marked as deprecated in the following ways:

- The response header is "Deprecated: true"
- The JSON response body includes "deprecated": true

Determine which API versions are supported in 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 API version for 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.

To configure CSRF protection, use the [Grid Management API](#) or [Tenant Management API](#).



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.

Manage system access

Use identity federation

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

Configure identity federation for Tenant Manager

You can configure identity federation for the Tenant Manager if you want tenant 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 Tenant 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 OpenLDAP server](#).
- If you plan to use Transport Layer Security (TLS) for communications with the LDAP server, the identity provider must be using TLS 1.2 or 1.3. See [Supported ciphers for outgoing TLS connections](#).

About this task

Whether you can configure an identity federation service for your tenant depends on how your tenant account was set up. Your tenant might share the identity federation service that was configured for the Grid Manager. If you see this message when you access the Identity Federation page, you cannot configure a separate federated identity source for this tenant.



This tenant account uses the LDAP server that is configured for the Grid Manager.
Contact the grid administrator for information or to change this setting.

Enter configuration

Steps

1. Select **ACCESS MANAGEMENT > 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 /.

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.

Test username

The username of a federated user.

Test password

Cancel Test Connection

- 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 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 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 groups

Create groups for an S3 tenant

You can manage permissions for S3 user groups by importing federated groups or creating local groups.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must belong to a user group that has the Root Access permission. See [Tenant management 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.

For information about S3, see [Use S3](#).

Steps

1. Select **ACCESS MANAGEMENT > Groups**.

Groups

Create and manage local and federated groups. Set group permissions to control access to specific pages and features.

2 groups

Create group

Actions

<input type="checkbox"/>	Name	ID	Type	Access mode
<input type="checkbox"/>	Applications	22cc2e27-88ee-4461-a8c6-30b550beec0	Local	Read-write
<input type="checkbox"/>	Managers	8b15b131-1d21-4539-93ad-f2298347c4d8	Local	Read-write

← Previous 1 Next →

2. Select **Create group**.

3. Select the **Local group** tab to create a local group, or select the **Federated group** tab to import a group from the previously configured identity source.

If single sign-on (SSO) is enabled for your StorageGRID system, users belonging to local groups will not be able to sign in to the Tenant Manager, although they can use client applications to manage the tenant's resources, based on group permissions.

4. Enter the group's name.

- **Local group:** Enter both a display name and a unique name. You can edit the display name later.
- **Federated group:** Enter the unique name. For Active Directory, the unique name is the name associated with the `sAMAccountName` attribute. For OpenLDAP, the unique name is the name associated with the `uid` attribute.

5. Select **Continue**.

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

- **Read-write** (default): Users can log into Tenant Manager and manage the tenant configuration.
- **Read-only**: Users can only view settings and features. They cannot make any changes or perform any operations in the Tenant Manager or Tenant Management API. Local read-only users can change their own passwords.

7. Select the Group permissions for this group.

See the information about tenant management permissions.

8. Select **Continue**.

9. Select a group policy to determine which S3 access permissions the members of this group will have.

- **No S3 Access**: Default. Users in this group do not have access to S3 resources, unless access is granted with a bucket policy. If you select this option, only the root user will have access to S3 resources by default.

- **Read Only Access:** Users in this group have read-only access to S3 resources. For example, users in this group can list objects and read object data, metadata, and tags. When you select this option, the JSON string for a read-only group policy appears in the text box. You cannot edit this string.
- **Full Access:** Users in this group have full access to S3 resources, including buckets. When you select this option, the JSON string for a full-access group policy appears in the text box. You cannot edit this string.
- **Custom:** Users in the group are granted the permissions you specify in the text box. See the instructions for implementing an S3 client application for detailed information about group policies, including language syntax and examples.

10. If you selected **Custom**, enter the group policy. Each group policy has a size limit of 5,120 bytes. You must enter a valid JSON formatted string.

In this example, members of the group are only permitted to list and access a folder matching their username (key prefix) in the specified bucket. Note that access permissions from other group policies and the bucket policy should be considered when determining the privacy of these folders.



```

{
  "Statement": [
    {
      "Sid": "AllowListBucketOfASpecificUserPrefix",
      "Effect": "Allow",
      "Action": "s3>ListBucket",
      "Resource": "arn:aws:s3:::department-bucket",
      "Condition": {
        "StringLike": {
          "s3:prefix": "${aws:username}/*"
        }
      }
    },
    {
      "Sid": "AllowUserSpecificActionsOnlyInTheSpecificFolder",
      "Effect": "Allow",
      "Action": "s3:GetObject",
      "Resource": "arn:aws:s3:::department-bucket/${aws:username}/*"
    }
  ]
}
  
```

11. Select the button that appears, depending on whether you are creating a federated group or a local group:

- Federated group: **Create group**
- Local group: **Continue**

If you are creating a local group, step 4 (Add users) appears after you select **Continue**. This step does not appear for federated groups.

12. Select the check box for each user you want to add to the group, then select **Create group**.

Optionally, you can save the group without adding users. You can add users to the group later, or select the group when you add new users.

13. Select **Finish**.

The group you created appears in the list of groups. Changes might take up to 15 minutes to take effect because of caching.

Create groups for a Swift tenant

You can manage access permissions for a Swift tenant account by importing federated groups or creating local groups. At least one group must have the Swift Administrator permission, which is required to manage the containers and objects for a Swift tenant account.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must belong to a user group that has the Root Access permission.
- If you plan to import a federated group, you have configured identity federation and the federated group already exists in the configured identity source.

Steps

1. Select **ACCESS MANAGEMENT > Groups**.

<input type="checkbox"/>	Name	ID	Type	Access mode
<input type="checkbox"/>	Applications	22cc2e27-88ee-4461-a8c6-30b550beeeec0	Local	Read-write
<input type="checkbox"/>	Managers	8b15b131-1d21-4539-93ad-f2298347c4d8	Local	Read-write

2. Select **Create group**.

3. Select the **Local group** tab to create a local group, or select the **Federated group** tab to import a group from the previously configured identity source.

If single sign-on (SSO) is enabled for your StorageGRID system, users belonging to local groups will not be able to sign in to the Tenant Manager, although they can use client applications to manage the tenant's resources, based on group permissions.

4. Enter the group's name.

- **Local group:** Enter both a display name and a unique name. You can edit the display name later.
- **Federated group:** Enter the unique name. For Active Directory, the unique name is the name associated with the `sAMAccountName` attribute. For OpenLDAP, the unique name is the name associated with the `uid` attribute.

5. Select **Continue**.

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

- **Read-write** (default): Users can log into Tenant Manager and manage the tenant configuration.
- **Read-only:** Users can only view settings and features. They cannot make any changes or perform any operations in the Tenant Manager or Tenant Management API. Local read-only users can change their own passwords.

7. Set the Group permission.

- Select the **Root Access** check box if users need to sign in to the Tenant Manager or Tenant Management API. (Default)
- Unselect the **Root Access** check box if users do not need access to the Tenant Manager or Tenant Management API. For example, unselect the check box for applications that do not need to access the tenant. Then, assign the **Swift Administrator** permission to allow these users to manage containers and objects.

8. Select **Continue**.

9. Select the **Swift administrator** check box if the user needs to be able to use the Swift REST API.

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

10. Select the button that appears, depending on whether you are creating a federated group or a local group:

- Federated group: **Create group**
- Local group: **Continue**

If you are creating a local group, step 4 (Add users) appears after you select **Continue**. This step does not appear for federated groups.

11. Select the check box for each user you want to add to the group, then select **Create group**.

Optionally, you can save the group without adding users. You can add users to the group later, or select the group when you create new users.

12. Select **Finish**.

The group you created appears in the list of groups. Changes might take up to 15 minutes to take effect because of caching.

Related information

[Tenant management permissions](#)

[Use Swift](#)

Tenant management permissions

Before you create a tenant group, consider which permissions you want to assign to that group. Tenant management permissions determine which tasks users can perform using the Tenant Manager or the Tenant Management API. A user can belong to one or more groups. Permissions are cumulative if a user belongs to multiple groups.

To sign in to the Tenant Manager or to use the Tenant Management API, users must belong to a group that has at least one permission. All users who can sign in can perform the following tasks:

- View the dashboard
- Change their own password (for local users)

For all permissions, the group's Access mode setting determines whether users can change settings and perform operations or whether they can only view the related settings and features.



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

You can assign the following permissions to a group. Note that S3 tenants and Swift tenants have different group permissions. Changes might take up to 15 minutes to take effect because of caching.

Permission	Description
Root Access	Provides full access to the Tenant Manager and the Tenant Management API. Note: Swift users must have Root Access permission to sign in to the tenant account.
Administrator	Swift tenants only. Provides full access to the Swift containers and objects for this tenant account Note: Swift users must have the Swift Administrator permission to perform any operations with the Swift REST API.
Manage Your Own S3 Credentials	S3 tenants only. Allows users to create and remove their own S3 access keys. Users who do not have this permission do not see the STORAGE (S3) > My S3 access keys menu option.

Permission	Description
Manage All Buckets	<ul style="list-style-type: none"> S3 tenants: Allows users to use the Tenant Manager and the Tenant Management API to create and delete S3 buckets and to manage the settings for all S3 buckets in the tenant account, regardless of S3 bucket or group policies. <p>Users who do not have this permission do not see the Buckets menu option.</p> <ul style="list-style-type: none"> Swift tenants: Allows Swift users to control the consistency level for Swift containers using the Tenant Management API. <p>Note: You can only assign the Manage All Buckets permission to Swift groups from the Tenant Management API. You cannot assign this permission to Swift groups using the Tenant Manager.</p>
Manage Endpoints	<p>S3 tenants only. Allows users to use the Tenant Manager or the Tenant Management API to create or edit endpoints, which are used as the destination for StorageGRID platform services.</p> <p>Users who do not have this permission do not see the Platform services endpoints menu option.</p>

Related information

[Use S3](#)

[Use Swift](#)

View and edit group details

When you view the details for a group, you can change the group's display name, permissions, policies, and the users that belong to the group.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must belong to a user group that has the Root Access permission.

Steps

- Select **ACCESS MANAGEMENT > Groups**.
- Select the name of the group whose details you want to view or edit.

Alternatively, you can select **Actions > View group details**.

The group details page appears. The following example shows the S3 group details page.

Overview

Display name:	Applications
Unique name:	group/Applications
Type:	Local
Access mode:	Read-write
Permissions:	Root Access
S3 Policy:	None
Number of users in this group:	0

[Group permissions](#)[S3 group policy](#)[Users](#)

Manage group permissions

Select an access mode for this group and select one or more permissions.

Access mode [?](#)

Select whether users can change settings and perform operations or whether they can only view settings and features.

 Read-write Read-only**Group permissions** [?](#)

Select the tenant account permissions you want to assign to this group.

 Root Access

Allows users to access all Tenant Manager features. Root Access permission supersedes all other permissions.

 Manage All Buckets

Allows users to change settings of all S3 buckets (or Swift containers) in this account.

 Manage Endpoints

Allows users to configure endpoints for platform services.

 Manage Your Own S3 Credentials

Allows users to create and delete their own S3 access keys.

[Save changes](#)

3. Make changes to the group settings as needed.



To ensure your changes are saved, select **Save changes** after you make changes in each section. When your changes are saved, a confirmation message appears in the upper right corner of the page.

- a. Optionally, select the display name or edit icon  to update the display name.

You cannot change a group's unique name. You cannot edit the display name for a federated group.

- b. Optionally, update the permissions.

- c. For group policy, make the appropriate changes for your S3 or Swift tenant.

- If you are editing a group for an S3 tenant, optionally select a different S3 group policy. If you select a custom S3 policy, update the JSON string as required.
- If you are editing a group for a Swift tenant, optionally select or unselect the **Swift Administrator** check box.

For more information about the Swift Administrator permission, see the instructions for creating groups for a Swift tenant.

- d. Optionally, add or remove users.

4. Confirm that you have selected **Save changes** for each section you changed.

Changes might take up to 15 minutes to take effect because of caching.

Related information

[Create groups for S3 tenant](#)

[Create groups for Swift tenant](#)

Add users to a local group

You can add users to a local group as needed.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must belong to a user group that has the Root Access permission.

Steps

1. Select **ACCESS MANAGEMENT > Groups**.
2. Select the name of the local group you want to add users to.

Alternatively, you can select **Actions > View group details**.

The group details page appears.

Overview

Display name:	Applications
Unique name:	group/Applications
Type:	Local
Access mode:	Read-write
Permissions:	Root Access
S3 Policy:	None
Number of users in this group:	0

[Group permissions](#)[S3 group policy](#)[Users](#)

Manage group permissions

Select an access mode for this group and select one or more permissions.

Access mode [?](#)

Select whether users can change settings and perform operations or whether they can only view settings and features.

 Read-write Read-only**Group permissions** [?](#)

Select the tenant account permissions you want to assign to this group.

 Root Access

Allows users to access all Tenant Manager features. Root Access permission supersedes all other permissions.

 Manage All Buckets

Allows users to change settings of all S3 buckets (or Swift containers) in this account.

 Manage Endpoints


Allows users to configure endpoints for platform services.

 Manage Your Own S3 Credentials

Allows users to create and delete their own S3 access keys.


[Save changes](#)

3. Select **Users**, and then select **Add users**.



The screenshot shows the 'Manage users' page. At the top, there are buttons for 'Add users' (blue) and 'Remove Users' (red). Below these are search fields for 'Search Groups...' and a magnifying glass icon. To the right, it says 'Displaying 1 results'. The main area is a table with columns: 'Username' (sorted by ascending), 'Full Name' (sorted by ascending), and 'Denied' (sorted by descending). One row is visible: 'User_02' with 'User_02_Managers' in the Full Name column and 'Denied' in the Denied column.

4. Select the users you want to add to the group, and then select **Add users**.



The screenshot shows the 'Add users' dialog box. It says 'Select local users to add to the group **Applications**'. Below is a search bar for 'Search Groups...' and a magnifying glass icon. It says 'Displaying 1 results'. A table lists a single user: 'User_01' with 'User_01_Applications' in the Full Name column. A checkbox next to 'User_01' is checked. The table has columns: 'Username' (sorted by ascending), 'Full Name' (sorted by ascending), and 'Denied' (sorted by descending). At the bottom are 'Cancel' and 'Add users' buttons.

A confirmation message appears in the upper right corner of the page. Changes might take up to 15 minutes to take effect because of caching.

Edit group name

You can edit the display name for a group. You cannot edit the unique name for a group.

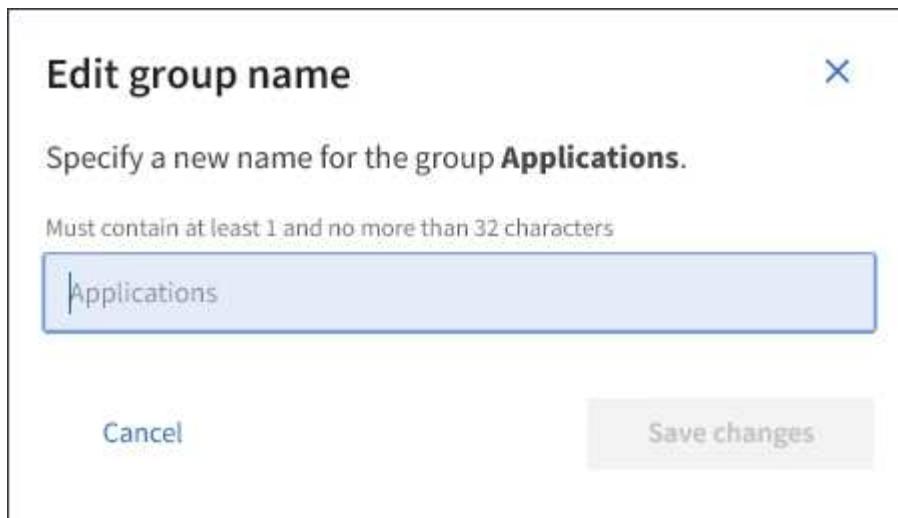
What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must belong to a user group that has the Root Access permission. See [Tenant management permissions](#).

Steps

1. Select **ACCESS MANAGEMENT > Groups**.
2. Select the check box for the group whose display name you want to edit.
3. Select **Actions > Edit group name**.

The Edit group name dialog box appears.



4. If you are editing a local group, update the display name as needed.

You cannot change a group's unique name. You cannot edit the display name for a federated group.

5. Select **Save changes**.

A confirmation message appears in the upper right corner of the page. Changes might take up to 15 minutes to take effect because of caching.

Duplicate group

You can create new groups more quickly by duplicating an existing group.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must belong to a user group that has the Root Access permission. See [Tenant management permissions](#).

Steps

1. Select **ACCESS MANAGEMENT > Groups**.
2. Select the check box for the group you want to duplicate.
3. Select **Duplicate group**. For additional details on creating a group, see the instructions for creating groups for [an S3 tenant](#) or for [a Swift tenant](#).
4. Select the **Local group** tab to create a local group, or select the **Federated group** tab to import a group from the previously configured identity source.

If single sign-on (SSO) is enabled for your StorageGRID system, users belonging to local groups will not be able to sign in to the Tenant Manager, although they can use client applications to manage the tenant's resources, [based on group permissions](#).

5. Enter the group's name.
 - **Local group**: Enter both a display name and a unique name. You can edit the display name later.
 - **Federated group**: Enter the unique name. For Active Directory, the unique name is the name associated with the `sAMAccountName` attribute. For OpenLDAP, the unique name is the name associated with the `uid` attribute.

6. Select **Continue**.
7. As needed, modify the permissions for this group.
8. Select **Continue**.
9. As needed, if you are duplicating a group for an S3 tenant, optionally select a different policy from the **Add S3 policy** radio buttons. If you selected a custom policy, update the JSON string as required.
10. Select **Create group**.

Delete group

You can delete a group from the system. Any users who belong only to that group will no longer be able to sign in to the Tenant Manager or use the tenant account.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must belong to a user group that has the Root Access permission. See [Tenant management permissions](#).

Steps

1. Select **ACCESS MANAGEMENT > Groups**.

The screenshot shows the 'Groups' management page. At the top, it says 'Create and manage local and federated groups. Set group permissions to control access to specific pages and features.' Below this, a message '2 groups' is displayed. A 'Create group' button is located in the top right corner. The main area is a table with the following data:

<input type="checkbox"/>	Name	ID	Type	Access mode
<input type="checkbox"/>	Applications	22cc2e27-88ee-4461-a8c6-30b550beec0	Local	Read-write
<input type="checkbox"/>	Managers	8b15b131-1d21-4539-93ad-f2298347c4d8	Local	Read-write

At the bottom right of the table, there are navigation links: 'Previous' and 'Next'.

2. Select the check boxes for the groups you want to delete.

3. Select **Actions > Delete group**.

A confirmation message appears.

4. Select **Delete group** to confirm you want to delete the groups indicated in the confirmation message.

A confirmation message appears in the upper right corner of the page. Changes might take up to 15 minutes to take effect because of caching.

Manage local users

You can create local users and assign them to local groups to determine which features these users can access. The Tenant Manager includes one predefined local user, named “root.” Although you can add and remove local users, you cannot remove the root user.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must belong to a read-write user group that has the Root Access permission. See [Tenant management permissions](#).



If single sign-on (SSO) is enabled for your StorageGRID system, local users will not be able to sign in to the Tenant Manager or the Tenant Management API, although they can use S3 or Swift client applications to access the tenant's resources, based on group permissions.

Access Users page

Select **ACCESS MANAGEMENT > Users**.

Actions	Username	Full Name	Denied	Type
<input type="checkbox"/>	root	Root		Local
<input type="checkbox"/>	User_01	User_01		Local
<input type="checkbox"/>	User_02	User_02		Local

Create local users

You can create local users and assign them to one or more local groups to control their access permissions.

S3 users who do not belong to any groups do not have management permissions or S3 group policies applied to them. These users might have S3 bucket access granted through a bucket policy.

Swift users who do not belong to any groups do not have management permissions or Swift container access.

Steps

1. Select **Create user**.

2. Complete the following fields.

- **Full name**: The full name for this user, for example, the first name and last name of a person or the name of an application.
- **Username**: The name this user will use to sign in. Usernames must be unique and cannot be changed.
- **Password**: A password, which is used when the user signs in.
- **Confirm password**: Type the same password you typed in the Password field.
- **Deny access**: If you select **Yes**, this user cannot sign in to the tenant account, even though the user might still belong to one or more groups.

As an example, you can use this feature to temporarily suspend a user's ability to sign in.

3. Select **Continue**.

4. Assign the user to one or more local groups.

Users who do not belong to any groups will have no management permissions. Permissions are cumulative. Users will have all permissions for all groups they belong to.

5. Select **Create user**.

Changes might take up to 15 minutes to take effect because of caching.

Edit user details

When you edit the details for a user, you can change the user's full name and password, add the user to different groups, and prevent the user from accessing the tenant.

Steps

1. In the Users list, select the name of the user whose details you want to view or edit.

Alternatively, you can select the check box for the user, and then select **Actions > View user details**.

2. Make changes to the user settings as needed.

a. Change the user's full name as needed by selecting the full name or the edit icon  in the Overview section.

You cannot change the username.

b. On the **Password** tab, change the user's password as needed.

c. On the **Access** tab, allow the user to sign in (select **No**), or prevent the user from signing in (select **Yes**) as needed.

d. On the **Groups** tab, add the user to groups or remove the user from groups as needed.

e. As necessary for each section, select **Save changes**.

Changes might take up to 15 minutes to take effect because of caching.

Duplicate local users

You can duplicate a local user to create a new user more quickly.

Steps

1. In the Users list, select the user you want to duplicate.
2. Select **Duplicate user**.
3. Modify the following fields for the new user.
 - **Full name**: The full name for this user, for example, the first name and last name of a person or the name of an application.
 - **Username**: The name this user will use to sign in. Usernames must be unique and cannot be changed.
 - **Password**: A password, which is used when the user signs in.
 - **Confirm password**: Type the same password you typed in the Password field.
 - **Deny access**: If you select **Yes**, this user cannot sign in to the tenant account, even though the user might still belong to one or more groups.

As an example, you can use this feature to temporarily suspend a user's ability to sign in.

4. Select **Continue**.
5. Select one or more local groups.

Users who do not belong to any groups will have no management permissions. Permissions are cumulative. Users will have all permissions for all groups they belong to.

6. Select **Create user**.

Changes might take up to 15 minutes to take effect because of caching.

Delete local users

You can permanently delete local users who no longer need to access the StorageGRID tenant account.

Using the Tenant Manager, you can delete local users, but not federated users. You must use the federated identity source to delete federated users.

Steps

1. In the Users list, select the check box for the local user you want to delete.
2. Select **Actions > Delete user**.
3. In the confirmation dialog box, select **Delete user** to confirm you want to delete the user from the system.

Changes might take up to 15 minutes to take effect because of caching.

Manage S3 tenant accounts

Manage S3 access keys

Each user of an S3 tenant account must have an access key to store and retrieve objects in the StorageGRID system. An access key consists of an access key ID and a secret

access key.

About this task

S3 access keys can be managed as follows:

- Users who have the **Manage Your Own S3 Credentials** permission can create or remove their own S3 access keys.
- Users who have the **Root Access** permission can manage the access keys for the S3 root account and all other users. Root access keys provide full access to all buckets and objects for the tenant unless explicitly disabled by a bucket policy.

StorageGRID supports Signature Version 2 and Signature Version 4 authentication. Cross-account access is not permitted unless explicitly enabled by a bucket policy.

Create your own S3 access keys

If you are using an S3 tenant and you have the appropriate permission, you can create your own S3 access keys. You must have an access key to access your buckets and objects in the S3 tenant account.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must have the **Manage Your Own S3 Credentials** permission. See [Tenant management permissions](#).

About this task

You can create one or more S3 access keys that allow you to create and manage buckets for your tenant account. After you create a new access key, update the application with your new access key ID and secret access key. For security, do not create more keys than you need, and delete the keys you are not using. If you have only one key and it is about to expire, create a new key before the old one expires, and then delete the old one.

Each key can have a specific expiration time or no expiration. Follow these guidelines for expiration time:

- Set an expiration time for your keys to limit your access to a certain time period. Setting a short expiration time can help reduce your risk if your access key ID and secret access key are accidentally exposed. Expired keys are removed automatically.
- If the security risk in your environment is low and you do not need to periodically create new keys, you do not have to set an expiration time for your keys. If you decide later to create new keys, delete the old keys manually.

 The S3 buckets and objects belonging to your account can be accessed using the access key ID and secret access key displayed for your account in the Tenant Manager. For this reason, protect access keys as you would a password. Rotate access keys on a regular basis, remove any unused keys from your account, and never share them with other users.

Steps


1. Select **STORAGE (S3) > My access keys**.

The My access keys page appears and lists any existing access keys.

2. Select **Create key**.

3. Do one of the following:

- Select **Do not set an expiration time** to create a key that will not expire. (Default)
- Select **Set an expiration time**, and set the expiration date and time.




4. Select **Create access key**.

The Download access key dialog box appears, listing your access key ID and secret access key.

5. Copy the access key ID and the secret access key to a safe location, or select **Download .csv** to save a spreadsheet file containing the access key ID and secret access key.



Do not close this dialog box until you have copied or downloaded this information. You cannot copy or download keys after the dialog box has been closed.



6. Select **Finish**.

The new key is listed on the My access keys page. Changes might take up to 15 minutes to take effect because of caching.

View your S3 access keys

If you are using an S3 tenant and you have the appropriate permission, you can view a list of your S3 access keys. You can sort the list by expiration time, so you can determine which keys will expire soon. As needed, you can create new keys or delete keys that you are no longer using.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must have the Manage Your Own S3 Credentials permission.



The S3 buckets and objects belonging to your account can be accessed using the access key ID and secret access key displayed for your account in the Tenant Manager. For this reason, protect access keys as you would a password. Rotate access keys on a regular basis, remove any unused keys from your account, and never share them with other users.

Steps

1. Select **STORAGE (S3) > My access keys**.

The My access keys page appears and lists any existing access keys.

My access keys

Manage your personal S3 access keys. If a key will expire soon, you can create a new key and delete the one it is replacing.

4 keys

[Create key](#)

[Delete key](#)

<input type="checkbox"/>	Access key ID	Expiration time
<input type="checkbox"/>	*****OTLS	2020-11-23 12:00:00 MST
<input type="checkbox"/>	*****0M45	2020-12-01 19:00:00 MST
<input type="checkbox"/>	*****69QJ	None
<input type="checkbox"/>	*****3R8P	None

2. Sort the keys by **Expiration time** or **Access key ID**.
3. As needed, create new keys and manually delete keys that you are no longer using.

If you create new keys before the existing keys expire, you can begin using the new keys without temporarily losing access to the objects in the account.

Expired keys are removed automatically.

Related information

[Create your own S3 access keys](#)

[Delete your own S3 access keys](#)

[Delete your own S3 access keys](#)

If you are using an S3 tenant and you have the appropriate permission, you can delete your own S3 access keys. After an access key is deleted, it can no longer be used to access the objects and buckets in the tenant account.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).

- You must have the Manage Your Own S3 Credentials permission. See [Tenant management permissions](#).



The S3 buckets and objects belonging to your account can be accessed using the access key ID and secret access key displayed for your account in the Tenant Manager. For this reason, protect access keys as you would a password. Rotate access keys on a regular basis, remove any unused keys from your account, and never share them with other users.


Steps

1. Select **STORAGE (S3) > My access keys**.

The My access keys page appears and lists any existing access keys.

2. Select the check box for each access key you want to remove.
3. Select **Delete key**.

A confirmation dialog box appears.



4. Select **Delete key**.

A confirmation message appears in the upper right corner of the page. Changes might take up to 15 minutes to take effect because of caching.

Create another user's S3 access keys

If you are using an S3 tenant and you have the appropriate permission, you can create S3 access keys for other users, such as applications that need access to buckets and objects.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must have the Root Access permission.

About this task

You can create one or more S3 access keys for other users so they can create and manage buckets for their tenant account. After you create a new access key, update the application with the new access key ID and

secret access key. For security, do not create more keys than the user needs, and delete the keys that are not being used. If you have only one key and it is about to expire, create a new key before the old one expires, and then delete the old one.

Each key can have a specific expiration time or no expiration. Follow these guidelines for expiration time:

- Set an expiration time for the keys to limit the user's access to a certain time period. Setting a short expiration time can help reduce risk if the access key ID and secret access key are accidentally exposed. Expired keys are removed automatically.
- If the security risk in your environment is low and you do not need to periodically create new keys, you do not have to set an expiration time for the keys. If you decide later to create new keys, delete the old keys manually.


 The S3 buckets and objects belonging to a user can be accessed using the access key ID and secret access key displayed for that user in the Tenant Manager. For this reason, protect access keys as you would a password. Rotate access keys on a regular basis, remove any unused keys from the account, and never share them with other users.

Steps

1. Select **ACCESS MANAGEMENT > Users**.
2. Select the user whose S3 access keys you want to manage.

The user detail page appears.

3. Select **Access keys**, then select **Create key**.
4. Do one of the following:
 - Select **Do not set an expiration time** to create a key that does not expire. (Default)
 - Select **Set an expiration time**, and set the expiration date and time.




5. Select **Create access key**.

The Download access key dialog box appears, listing the access key ID and secret access key.

6. Copy the access key ID and the secret access key to a safe location, or select **Download .csv** to save a spreadsheet file containing the access key ID and secret access key.



Do not close this dialog box until you have copied or downloaded this information. You cannot copy or download keys after the dialog box has been closed.



7. Select **Finish**.

The new key is listed on the Access keys tab of the user details page. Changes might take up to 15 minutes to take effect because of caching.

Related information

[Tenant management permissions](#)

View another user's S3 access keys

If you are using an S3 tenant and you have appropriate permissions, you can view another user's S3 access keys. You can sort the list by expiration time so you can determine which keys will expire soon. As needed, you can create new keys and delete keys that are no longer in use.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must have the Root Access permission.

 The S3 buckets and objects belonging to a user can be accessed using the access key ID and secret access key displayed for that user in the Tenant Manager. For this reason, protect access keys as you would a password. Rotate access keys on a regular basis, remove any unused keys from the account, and never share them with other users.

Steps


1. Select **ACCESS MANAGEMENT > Users**.

The Users page appears and lists the existing users.

2. Select the user whose S3 access keys you want to view.

The User details page appears.

3. Select **Access keys**.



Access key ID	Expiration time
*****WX5J	2020-11-21 12:00:00 MST
*****60HM	2020-11-23 13:00:00 MST
*****J505	None
*****4MTF	None

4. Sort the keys by **Expiration time** or **Access key ID**.
5. As needed, create new keys and manually delete keys that are no longer in use.

If you create new keys before the existing keys expire, the user can begin using the new keys without temporarily losing access to the objects in the account.

Expired keys are removed automatically.

Related information

[Create another user's S3 access keys](#)

[Delete another user's S3 access keys](#)

Delete another user's S3 access keys

If you are using an S3 tenant and you have appropriate permissions, you can delete another user's S3 access keys. After an access key is deleted, it can no longer be used to access the objects and buckets in the tenant account.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must have the Root Access permission. See [Tenant management permissions](#).



The S3 buckets and objects belonging to a user can be accessed using the access key ID and secret access key displayed for that user in the Tenant Manager. For this reason, protect access keys as you would a password. Rotate access keys on a regular basis, remove any unused keys from the account, and never share them with other users.

Steps

1. Select **ACCESS MANAGEMENT > Users**.

The Users page appears and lists the existing users.


2. Select the user whose S3 access keys you want to manage.

The User details page appears.

3. Select **Access keys**, and then select the check box for each access key you want to delete.

4. Select **Actions > Delete selected key**.

A confirmation dialog box appears.



5. Select **Delete key**.

A confirmation message appears in the upper right corner of the page. Changes might take up to 15 minutes to take effect because of caching.

Manage S3 buckets


Use S3 Object Lock with tenants

You can use the S3 Object Lock feature in StorageGRID if your objects must comply with regulatory requirements for retention.

What is S3 Object Lock?

The StorageGRID S3 Object Lock feature is an object-protection solution that is equivalent to S3 Object Lock in Amazon Simple Storage Service (Amazon S3).

As shown in the figure, when the global S3 Object Lock setting is enabled for a StorageGRID system, an S3 tenant account can create buckets with or without S3 Object Lock enabled. If a bucket has S3 Object Lock enabled, S3 client applications can optionally specify retention settings for any object version in that bucket. An object version must have retention settings specified to be protected by S3 Object Lock.



The StorageGRID S3 Object Lock feature provides a single retention mode that is equivalent to the Amazon S3 compliance mode. By default, a protected object version cannot be overwritten or deleted by any user. The StorageGRID S3 Object Lock feature does not support a governance mode, and it does not allow users with special permissions to bypass retention settings or to delete protected objects.

If a bucket has S3 Object Lock enabled, the S3 client application can optionally specify either or both of the following object-level retention settings when creating or updating an object:

- **Retain-until-date:** If an object version's retain-until-date is in the future, the object can be retrieved, but it cannot be modified or deleted. As required, an object's retain-until-date can be increased, but this date cannot be decreased.
- **Legal hold:** Applying a legal hold to an object version immediately locks that object. For example, you might need to put a legal hold on an object that is related to an investigation or legal dispute. A legal hold has no expiration date, but remains in place until it is explicitly removed. Legal holds are independent of the retain-until-date.

You can also [specify a default retention mode and default retention period for the bucket](#). These are applied to each object added to the bucket that does not specify its own retention settings.

For details on these settings, see [Use S3 object lock](#).

Manage legacy Compliant buckets

The S3 Object Lock feature replaces the Compliance feature that was available in previous StorageGRID versions. If you created compliant buckets using a previous version of StorageGRID, you can continue to manage the settings of these buckets; however, you can no longer create new compliant buckets. For instructions, see the NetApp Knowledge Base article.


[NetApp Knowledge Base: How to manage legacy Compliant buckets in StorageGRID 11.5](#)

S3 Object Lock workflow

The workflow diagram shows the high-level steps for using the S3 Object Lock feature in StorageGRID.

Before you can create buckets with S3 Object Lock enabled, the grid administrator must enable the global S3 Object Lock setting for the entire StorageGRID system. The grid administrator must also ensure that the [information lifecycle management \(ILM\) policy](#) is “compliant”; it must meet the requirements of buckets with S3 Object Lock enabled. For details, contact your grid administrator or see the instructions for managing objects with information lifecycle management.

After the global S3 Object Lock setting has been enabled, you can create buckets with S3 Object Lock enabled. You can then use the S3 client application to optionally specify retention settings for each object version.



Requirements for S3 Object Lock

Before enabling S3 Object Lock for a bucket, review the requirements for S3 Object Lock buckets and objects and the lifecycle of objects in buckets with S3 Object Lock enabled.

Requirements for buckets with S3 Object Lock enabled

- If the global S3 Object Lock setting is enabled for the StorageGRID system, you can use the Tenant Manager, the Tenant Management API, or the S3 REST API to create buckets with S3 Object Lock enabled.

This example from the Tenant Manager shows a bucket with S3 Object Lock enabled.

Buckets

Create buckets and manage bucket settings.



Actions	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

- 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 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.

Create S3 bucket

You can use the Tenant Manager to create S3 buckets for object data. When you create a bucket, you must specify the bucket's name and region. If the global S3 Object Lock setting is enabled for the StorageGRID system, you can optionally enable S3 Object Lock for the bucket.

What you'll need

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the Manage All Buckets or the Root Access permission. These permissions override the permissions settings in group or bucket policies.



Permissions to set or modify S3 Object Lock properties of buckets or objects can be granted by [bucket policy or group policy](#).

- If you plan to create a bucket with S3 Object Lock, you have enabled the global S3 Object Lock setting for the StorageGRID system, and you have reviewed the requirements for S3 Object Lock buckets and objects.

[Use S3 Object Lock](#)

Steps

1. Select **STORAGE (S3) > Buckets**.
2. Select **Create bucket**.

Create bucket


1 Enter details 2 Manage object settings
Optional

Enter bucket details
Enter the bucket's name and select the bucket's region.

Bucket name ?

Region ?
 ▼

Cancel Continue



3. Enter a unique name for the bucket.



You cannot change the bucket name after creating the bucket.

Bucket names must comply with these rules:

- Must be unique across each StorageGRID system (not just unique within the tenant account).
- Must be DNS compliant.
- Must contain at least 3 and no more than 63 characters.
- Each label must start and end with a lowercase letter or a number and can only use lowercase letters, numbers, and hyphens.
- Should not use periods in virtual hosted style requests. Periods will cause problems with server wildcard certificate verification.



For more information, see the [Amazon Web Services \(AWS\) documentation on bucket naming rules](#).

4. Select the region for this bucket.

Your StorageGRID administrator manages the available regions. A bucket's region can affect the data-protection policy applied to objects. By default, all buckets are created in the `us-east-1` region.



You cannot change the region after creating the bucket.

5. Select **Continue**.

6. Optionally, enable object versioning for the bucket.

Enable object versioning if you want to store every version of each object in this bucket. You can then retrieve previous versions of an object as needed.

7. If the S3 Object Lock section appears, optionally enable S3 Object Lock for the bucket.



You cannot enable or disable S3 Object Lock after creating the bucket.

The S3 Object Lock section appears only if the global S3 Object Lock setting is enabled.

S3 Object Lock must be enabled for the bucket before an S3 client application can specify retain-until-date and legal hold settings for the objects added to the bucket.

If you enable S3 Object Lock for a bucket, bucket versioning is enabled automatically. You can also [specify a default retention mode and default retention period for the bucket](#) that are applied to each object ingested to the bucket that does not specify its own retention settings.

8. Select **Create bucket**.

The bucket is created and added to the table on the Buckets page.

Related information

[Manage objects with ILM](#)

[Understand Tenant Management API](#)

[Use S3](#)

View S3 bucket details

You can view a list of the buckets and bucket settings in your tenant account.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).


Steps

1. Select **STORAGE (S3) > Buckets**.

The Buckets page appears and lists all buckets for the tenant account.

Buckets

Create buckets and manage bucket settings.



The screenshot shows the AWS S3 Buckets list interface. At the top, there is a header with '3 buckets' and a 'Create bucket' button. Below the header is a table with the following data:

Actions	Name	S3 Object Lock	Region	Object Count	Space Used	Date Created
<input type="checkbox"/>	bucket-01a	✓	us-east-1	0	0 bytes	2022-01-06 13:48:08 MST
<input type="checkbox"/>	bucket-02a	✓	us-east-1	0	0 bytes	2022-01-06 13:48:26 MST
<input type="checkbox"/>	bucket-03a		us-east-1	0	0 bytes	2022-01-06 13:48:38 MST

2. Review the information for each bucket.

As required, you can sort the information by any column, or you can page forward and back through the list.

- Name: The bucket's unique name, which cannot be changed.
- S3 Object Lock: Whether S3 Object Lock is enabled for this bucket.

This column is not displayed if the global S3 Object lock setting is disabled. This column also shows information for any legacy Compliant buckets.

- Region: The bucket's region, which cannot be changed.
- Object Count: The number of objects in this bucket.
- Space Used: The logical size of all objects in this bucket. The logical size does not include the actual space required for replicated or erasure-coded copies or for object metadata.
- Date Created: The date and time the bucket was created.



The Object Count and Space Used values displayed are estimates. These estimates are affected by the timing of ingests, network connectivity, and node status. If buckets have versioning enabled, deleted object versions are included in the object count.

3. To view and manage the settings for a bucket, select the bucket name.

The bucket details page allows you to view and edit the settings for bucket options, bucket access, and [platform services](#).

Buckets > bucket-01

Overview ^

Name:	bucket-01
Region:	us-east-1
Date created:	2021-11-30 09:55:55 MST

[View bucket contents in Experimental S3 Console](#) ↗

Bucket options Bucket access Platform services

Consistency level	Read-after-new-write (default)	▼
Last access time updates	Disabled	▼
Object versioning	Enabled	▼
S3 Object Lock	Disabled	▼

Change the consistency level

If you are using an S3 tenant, you can use the Tenant Manager or the Tenant Management API to change the consistency control for operations performed on the objects in S3 buckets.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must belong to a user group that has the Manage All Buckets or the Root Access permission. These permissions override the permissions settings in group or bucket policies. See [Tenant management permissions](#).

About this task

Consistency level provides a balance between the availability of the objects and the consistency of those objects across different Storage Nodes and sites. In general, you should use the **Read-after-new-write** consistency level for your buckets.

If the **Read-after-new-write** consistency level does not meet the client application's requirements, you can change the consistency level by setting the bucket consistency level or by using the **Consistency-Control** header. The **Consistency-Control** header overrides the bucket consistency level.



When you change a bucket's consistency level, only those objects that are ingested after the change are guaranteed to meet the revised level.

Steps

1. Select **STORAGE (S3) > Buckets**.

2. Select the bucket name from the list.

The bucket details page appears.

3. Select **Bucket options > Consistency level**.

4. Select a consistency level for operations performed on the objects in this bucket.

- **All**: Provides the highest level of consistency. All nodes receive the data immediately, or the request will fail.
- **Strong-global**: Guarantees read-after-write consistency for all client requests across all sites.
- **Strong-site**: Guarantees read-after-write consistency for all client requests within a site.
- **Read-after-new-write** (default): Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees. Recommended for most cases.
- **Available**: Provides eventual consistency for both new objects and object updates. For S3 buckets, use only as required (for example, for a bucket that contains log values that are rarely read, or for HEAD or GET operations on keys that do not exist). Not supported for S3 FabricPool buckets.

5. Select **Save changes**.

Enable or disable last access time updates

When grid administrators create the information lifecycle management (ILM) rules for a StorageGRID system, they can optionally specify that an object's last access time be used to determine whether to move that object to a different storage location. If you are using an S3 tenant, you can take advantage of such rules by enabling last access time updates for the objects in an S3 bucket.

These instructions only apply to StorageGRID systems that include at least one ILM rule that uses the **Last Access Time** option in its placement instructions. You can ignore these instructions if your StorageGRID system does not include such a rule.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must belong to a user group that has the Manage All Buckets or the Root Access permission. These permissions override the permissions settings in group or bucket policies. See [Tenant management permissions](#).

Last Access Time is one of the options available for the **Reference Time** placement instruction for an ILM rule. Setting the Reference Time for a rule to Last Access Time lets grid administrators specify that objects be placed in certain storage locations based on when those objects were last retrieved (read or viewed).

For example, to ensure that recently viewed objects remain on faster storage, a grid administrator can create an ILM rule specifying the following:

- Objects that have been retrieved in the past month should remain on local Storage Nodes.
- Objects that have not been retrieved in the past month should be moved to an off-site location.



See the instructions for managing objects with information lifecycle management.

By default, updates to last access time are disabled. If your StorageGRID system includes an ILM rule that uses the **Last Access Time** option and you want this option to apply to objects in this bucket, you must enable updates to last access time for the S3 buckets specified in that rule.



Updating the last access time when an object is retrieved can reduce StorageGRID performance, especially for small objects.

A performance impact occurs with last access time updates because StorageGRID must perform these additional steps every time objects are retrieved:

- Update the objects with new timestamps
- Add the objects to the ILM queue, so they can be reevaluated against current ILM rules and policy

The table summarizes the behavior applied to all objects in the bucket when last access time is disabled or enabled.

Type of request	Behavior if last access time is disabled (default)		Behavior if last access time is enabled	
	Last access time updated?	Object added to ILM evaluation queue?	Last access time updated?	Object added to ILM evaluation queue?
Request to retrieve an object, its access control list, or its metadata	No	No	Yes	Yes
Request to update an object's metadata	Yes	Yes	Yes	Yes
Request to copy an object from one bucket to another	<ul style="list-style-type: none">• No, for the source copy• Yes, for the destination copy	<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	<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	Yes, for the assembled object	Yes, for the assembled object

Steps

1. Select **STORAGE (S3) > Buckets**.
2. Select the bucket name from the list.

The bucket details page appears.

3. Select **Bucket options > Last access time updates**.
4. Select the appropriate radio button to enable or disable last access time updates.

Bucket options Bucket access Platform services

Consistency level Read-after-new-write (default) ▼

Last access time updates Disabled ▲

Enable or disable last access time updates for the objects in this bucket.

When last access time updates are disabled, the following behavior applies to objects in the bucket:

- Requests to retrieve an object, its access control list, or its metadata do not update the object's last access time. The object is not added to ILM evaluation queues.
- Requests to update an object's metadata update the object's last access time. The object is added to ILM evaluation queues.
- Requests to copy an object from one bucket to another do not update the last access time for the source copy and do not add the source object to the ILM evaluation queue. However, the last access time is updated for the destination copy, and the destination object is added to ILM evaluation queues.
- A request to complete a multipart upload causes the last access time for the assembled object to be updated. The new object is added to ILM evaluation queues.

i Updating the last access time when an object is retrieved can reduce performance, especially for small objects.

Enable last access time updates when retrieving an object

Disable last access time updates when retrieving an object

Save changes

5. Select **Save changes**.

Related information

[Tenant management permissions](#)

[Manage objects with ILM](#)

Change object versioning for a bucket

If you are using an S3 tenant, you can use the Tenant Manager or the Tenant Management API to change the versioning state for S3 buckets.

What you'll need

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the Manage All Buckets or the Root Access permission. These permissions override the permissions settings in group or bucket policies.

[Tenant management permissions](#)

About this task

You can enable or suspend object versioning for a bucket. After you enable versioning for a bucket, it cannot return to an unversioned state. However, you can suspend versioning for the bucket.


- Disabled: Versioning has never been enabled
- Enabled: Versioning is enabled
- Suspended: Versioning was previously enabled and is suspended

S3 object versioning

ILM rules and policies for S3 versioned objects (Example 4)

Steps

1. Select **STORAGE (S3) > Buckets**.
2. Select the bucket name from the list.
3. Select **Bucket options > Object versioning**.



The screenshot shows the 'Bucket options' tab selected in the navigation bar. The 'Object versioning' section is expanded, showing the status as 'Enabled'. A descriptive text explains that enabling versioning stores every version of each object in the bucket, allowing recovery from errors. Below this, two radio buttons are shown: 'Enable versioning' (selected) and 'Suspend versioning'. A 'Save changes' button is located in the bottom right corner of the section.

Bucket options Bucket access Platform services

Consistency level Read-after-new-write (default) ▾

Last access time updates Disabled ▾

Object versioning Enabled ▾

Enable object versioning if you want to store every version of each object in this bucket. You can then retrieve a previous object version to recover from an error.

After versioning is enabled, you can optionally suspend versioning for the bucket. New object versions are no longer created, but you can still retrieve any existing object versions.

Enable versioning

Suspend versioning

Save changes

4. Select a versioning state for the objects in this bucket.



If S3 Object Lock or legacy compliance is enabled, the **Object versioning** options are disabled.

Option	Description
Enable versioning	<p>Enable object versioning if you want to store every version of each object in this bucket. You can then retrieve previous versions of an object as needed.</p> <p>Objects that were already in the bucket will be versioned when they are modified by a user.</p>
Suspend versioning	Suspend object versioning if you no longer want new object versions to be created. You can still retrieve any existing object versions.

5. Select **Save changes**.

Configure Cross-Origin Resource Sharing (CORS)

You can configure Cross-Origin Resource Sharing (CORS) for an S3 bucket if you want that bucket and objects in that bucket to be accessible to web applications in other domains.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must belong to a user group that has the Manage All Buckets or the Root Access permission. These permissions override the permissions settings in group or bucket policies.

About this task

Cross-Origin Resource Sharing (CORS) is a security mechanism that allows client web applications in one domain to access resources in a different domain. For example, suppose you use an S3 bucket named Images to store graphics. By configuring CORS for the Images bucket, you can allow the images in that bucket to be displayed on the website <http://www.example.com>.

Steps

1. Use a text editor to create the XML required to enable CORS.

This example shows the XML used to enable CORS for an S3 bucket. This XML allows any domain to send GET requests to the bucket, but it only allows the <http://www.example.com> domain to send POST and DELETE requests. All request headers are allowed.

```
<CORSConfiguration
  xmlns="http://s3.amazonaws.com/doc/2020-10-22/">
  <CORSRule>
    <AllowedOrigin>*</AllowedOrigin>
    <AllowedMethod>GET</AllowedMethod>
    <AllowedHeader>*</AllowedHeader>
  </CORSRule>
  <CORSRule>
    <AllowedOrigin>http://www.example.com</AllowedOrigin>
    <AllowedMethod>GET</AllowedMethod>
    <AllowedMethod>POST</AllowedMethod>
    <AllowedMethod>DELETE</AllowedMethod>
    <AllowedHeader>*</AllowedHeader>
  </CORSRule>
</CORSConfiguration>
```

For more information about the CORS configuration XML, see [Amazon Web Services \(AWS\) Documentation: Amazon Simple Storage Service Developer Guide](#).

2. In the Tenant Manager, select **STORAGE (S3) > Buckets**.

3. Select the bucket name from the list.

The bucket details page appears.

4. Select **Bucket access > Cross-Origin Resource Sharing (CORS)**.

5. Select the **Enable CORS** check box.

6. Paste the CORS configuration XML into the text box, and select **Save changes**.

The screenshot shows the 'Bucket access' tab in the AWS Tenant Manager. Under 'Cross-Origin Resource Sharing (CORS)', the status is 'Disabled'. A checkbox labeled 'Enable CORS' is checked. Below it is a text area containing the following XML configuration:

```
<corsConfiguration xmlns="http://s3.amazonaws.com/doc/2020-10-22/">
  <corsRule>
    <allowedOrigin>*</allowedOrigin>
    <allowedMethod>GET</allowedMethod>
    <allowedHeader>*</allowedHeader>
  </corsRule>
  <corsRule>
    <allowedOrigin>http://www.example.com</allowedOrigin>
    <allowedMethod>GET</allowedMethod>
    <allowedMethod>POST</allowedMethod>
    <allowedMethod>DELETE</allowedMethod>
  </corsRule>

```

At the bottom right of the text area is a 'Clear' button. At the bottom right of the entire screen is a 'Save changes' button.

7. To modify the CORS setting for the bucket, update the CORS configuration XML in the text box or select **Clear** to start over. Then select **Save changes**.
8. To disable CORS for the bucket, unselect the **Enable CORS** check box, and then select **Save changes**.

Delete S3 bucket

You can use the Tenant Manager to delete one or more S3 buckets that are empty.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must belong to a user group that has the Manage All Buckets or the Root Access permission. These permissions override the permissions settings in group or bucket policies. See [Tenant management permissions](#).
- The buckets you want to delete are empty.

About this task

These instructions describe how to delete an S3 bucket using the Tenant Manager. You can also delete S3 buckets using the [Tenant Management API](#) or the [S3 REST API](#).


You cannot delete an S3 bucket if it contains objects or noncurrent object versions. For information about how

S3 versioned objects are deleted, see the [instructions for managing objects with information lifecycle management](#).

Steps

1. Select **STORAGE (S3) > Buckets**.

The Buckets page appears and shows all existing S3 buckets.




The screenshot shows the AWS S3 Buckets page. At the top, it says "Create buckets and manage bucket settings." Below that, it shows "3 buckets" and a "Create bucket" button. The main area is a table with columns: Actions, Name, S3 Object Lock, Region, Object Count, Space Used, and Date Created. The table contains three rows for buckets:

Actions	Name	S3 Object Lock	Region	Object Count	Space Used	Date Created
<input type="checkbox"/>	bucket-01a	✓	us-east-1	0	0 bytes	2022-01-06 13:48:08 MST
<input type="checkbox"/>	bucket-02a	✓	us-east-1	0	0 bytes	2022-01-06 13:48:26 MST
<input type="checkbox"/>	bucket-03a		us-east-1	0	0 bytes	2022-01-06 13:48:38 MST

2. Select the check box for the empty bucket you want to delete. You can select more than one bucket at a time.

The Actions menu is enabled.

3. From the Actions menu, select **Delete bucket** (or **Delete buckets** if you have chosen more than one).



The screenshot shows the AWS S3 Buckets page with the "Delete bucket" option selected in the Actions menu. The table below shows two buckets: bucket-01 and bucket-02. bucket-01 has a checked checkbox in the Actions column, while bucket-02 does not.

Actions	Name	Region	Object Count	Space Used	Date Created
<input checked="" type="checkbox"/>	bucket-01	us-east-1	0	0 bytes	2021-12-02 11:14:26 MST
<input type="checkbox"/>	bucket-02	us-east-1	0	0 bytes	2021-12-02 11:14:49 MST

4. When the confirmation dialog box appears, select **Yes** to delete all of the buckets you have chosen.

StorageGRID confirms that each bucket is empty and then deletes each bucket. This operation might take a few minutes.

If a bucket is not empty, an error message appears. You must delete all objects before you can delete a bucket.

Use Experimental S3 Console

You can use S3 Console to view the objects in an S3 bucket.

You can also use S3 Console to do the following:

- Add and delete objects, object versions, and folders
- Rename objects
- Move and copy objects between buckets and folders
- Manage object tags
- View object metadata
- Download objects

 S3 Console has not been fully tested and is marked as “experimental.” It is not intended for bulk management of objects or for use in a production environment. Tenants should only use S3 Console when performing functions for a small number of objects, such as when uploading objects to simulate a new ILM policy, troubleshooting ingest issues, or using proof-of-concept or non-production grids.

What you'll need

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You have the Manage Your Own S3 Credentials permission.
- You have created a bucket.
- You know the user's access key ID and secret access key. Optionally, you have a `.csv` file containing this information. See the [instructions for creating access keys](#).

Steps

1. Select **Buckets**.
2. Select [Experimental S3 Console](#) . You can also access this link from the bucket details page.
3. On the Experimental S3 Console sign-in page, paste the access key ID and secret access key into the fields. Otherwise, select **Upload access keys** and select your `.csv` file.
4. Select **Sign in**.
5. Manage objects as needed.

Buckets > bucket-01

bucket-01

Upload New folder Refresh Actions ▾ Search by prefix 🔍

Name	Logical space used	Last modified on
03_Grid_Primer_11.5.pdf	2.73 MB	2021-12-03 09:43:26 MST
04_Tenant_Users_Guide_11.5.pdf	1.07 MB	2021-12-03 09:44:24 MST
06_Tenant_Users_Guide_11.5.pdf	1.25 MB	2021-12-03 09:44:27 MST
08_Tenant_Users_Guide_11.5.pdf	1.25 MB	2021-12-03 09:44:27 MST
09_Tenant_Users_Guide_11.5.pdf	1.25 MB	2021-12-03 09:44:26 MST
10_Grid_Primer_11.5.pdf	2.8 MB	2021-12-03 09:43:27 MST

Select an object or folder to view its details.

Displaying 16 objects Selected 0 objects

Manage S3 platform services

What are platform services?

StorageGRID platform services can help you implement a hybrid cloud strategy.

If the use of platform services is allowed for your tenant account, you can configure the following services for any S3 bucket:

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

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



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

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

For example, you could configure alerts to be sent to administrators about each object added to a bucket,

where the objects represent log files associated with a critical system event.



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

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

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



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

Because the target location for platform services is typically external to your StorageGRID deployment, platform services give you the power and flexibility that comes from using external storage resources, notification services, and search or analysis services for your data.

Any combination of platform services can be configured for a single S3 bucket. For example, you could configure both the CloudMirror service and notifications on a StorageGRID S3 bucket so that you can mirror specific objects to the Amazon Simple Storage Service, while sending a notification about each such object to a third party monitoring application to help you track your AWS expenses.



The use of platform services must be enabled for each tenant account by a StorageGRID administrator using the Grid Manager or the Grid Management API.

How platform services are configured

Platform services communicate with external endpoints that you configure using the Tenant Manager or the Tenant Management API. Each endpoint represents an external destination, such as a StorageGRID S3 bucket, an Amazon Web Services bucket, a Simple Notification Service (SNS) topic, or an Elasticsearch cluster hosted locally, on AWS, or elsewhere.

After you create an endpoint, you can enable a platform service for a bucket by adding XML configuration to the bucket. The XML configuration identifies the objects that the bucket should act on, the action that the bucket should take, and the endpoint that the bucket should use for the service.

You must add separate XML configurations for each platform service that you want to configure. For example:

1. If you want all objects whose keys start with `/images` to be replicated to an Amazon S3 bucket, you must add a replication configuration to the source bucket.
2. If you also want to send notifications when these objects are stored to the bucket, you must add a notifications configuration.
3. Finally, if you want to index the metadata for these objects, you must add the metadata notification configuration that is used to implement search integration.

The format for the configuration XML is governed by the S3 REST APIs used to implement StorageGRID platform services:

Platform service	S3 REST API
CloudMirror replication	<ul style="list-style-type: none"> • GET Bucket replication • PUT Bucket replication
Notifications	<ul style="list-style-type: none"> • GET Bucket notification • PUT Bucket notification
Search integration	<ul style="list-style-type: none"> • GET Bucket metadata notification configuration • PUT Bucket metadata notification configuration <p>These operations are custom to StorageGRID.</p>

See the instructions for implementing S3 client applications for details on how StorageGRID implements these APIs.

Related information

[Considerations for using platform services](#)

[Use S3](#)

CloudMirror replication service

You can enable CloudMirror replication for an S3 bucket if you want StorageGRID to replicate specified objects added to the bucket to one or more destination buckets.

CloudMirror replication operates independently of the grid's active ILM policy. The CloudMirror service replicates objects as they are stored to the source bucket and delivers them to the destination bucket as soon as possible. Delivery of replicated objects is triggered when object ingest succeeds.

If you enable CloudMirror replication for an existing bucket, only the new objects added to that bucket are replicated. Any existing objects in the bucket are not replicated. To force the replication of existing objects, you can update the existing object's metadata by performing an object copy.



If you are using CloudMirror replication to copy objects to an AWS S3 destination, be aware that Amazon S3 limits the size of user-defined metadata within each PUT request header to 2 KB. If an object has user-defined metadata greater than 2 KB, that object will not be replicated.

In StorageGRID, you can replicate the objects in a single bucket to multiple destination buckets. To do so, specify the destination for each rule in the replication configuration XML. You cannot replicate an object to more than one bucket at the same time.

Additionally, you can configure CloudMirror replication on versioned or unversioned buckets, and you can specify a versioned or unversioned bucket as the destination. You can use any combination of versioned and unversioned buckets. For example, you could specify a versioned bucket as the destination for an unversioned source bucket, or vice versa. You can also replicate between unversioned buckets.

Deletion behavior for the CloudMirror replication service is the same as the deletion behavior of the Cross Region Replication (CRR) service provided by Amazon S3—deleting an object in a source bucket never deletes a replicated object in the destination. If both source and destination buckets are versioned, the delete marker is replicated. If the destination bucket is not versioned, deleting an object in the source bucket does not

replicate the delete marker to the destination bucket or delete the destination object.

As objects are replicated to the destination bucket, StorageGRID marks them as “replicas.” A destination StorageGRID bucket will not replicate objects marked as replicas again, protecting you from accidental replication loops. This replica marking is internal to StorageGRID and does not prevent you from leveraging AWS CRR when using an Amazon S3 bucket as the destination.



The custom header used to mark a replica is `x-ntap-sg-replica`. This marking prevents a cascading mirror. StorageGRID does support a bi-directional CloudMirror between two grids.

The uniqueness and ordering of events in the destination bucket are not guaranteed. More than one identical copy of a source object might be delivered to the destination as a result of operations taken to guarantee delivery success. In rare cases, when the same object is updated simultaneously from two or more different StorageGRID sites, the ordering of operations on the destination bucket might not match the ordering of events on the source bucket.

CloudMirror replication is typically configured to use an external S3 bucket as a destination. However, you can also configure replication to use another StorageGRID deployment or any S3-compatible service.

Understand notifications for buckets

You can enable event notification for an S3 bucket if you want StorageGRID to send notifications about specified events to a destination Amazon Simple Notification Service (SNS).

You can [configure event notifications](#) by associating notification configuration XML with a source bucket. The notification configuration XML follows S3 conventions for configuring bucket notifications, with the destination SNS topic specified as the URN of an endpoint.

Event notifications are created at the source bucket as specified in the notification configuration and are delivered to the destination. If an event associated with an object succeeds, a notification about that event is created and queued for delivery.

The uniqueness and ordering of notifications are not guaranteed. More than one notification of an event might be delivered to the destination as a result of operations taken to guarantee delivery success. And because delivery is asynchronous, the time ordering of notifications at the destination is not guaranteed to match the ordering of events on the source bucket, particularly for operations that originate from different StorageGRID sites. You can use the `sequencer` key in the event message to determine the order of events for a particular object, as described in Amazon S3 documentation.

Supported notifications and messages

StorageGRID event notification follows the Amazon S3 API with the following limitations:

- You cannot configure a notification for the following event types. These event types are **not** supported.
 - `s3:ReducedRedundancyLostObject`
 - `s3:ObjectRestore:Completed`
- Event notifications sent from StorageGRID use the standard JSON format except that they do not include some keys and use specific values for others, as shown in the table:

Key name	StorageGRID value
eventSource	sgws:s3
awsRegion	not included
x-amz-id-2	not included
arn	urn:sgws:s3:::bucket_name

Understand search integration service

You can enable search integration for an S3 bucket if you want to use an external search and data analysis service for your object metadata.

The search integration service is a custom StorageGRID service that automatically and asynchronously sends S3 object metadata to a destination endpoint whenever an object or its metadata is updated. You can then use sophisticated search, data analysis, visualization, or machine learning tools provided by the destination service to search, analyze, and gain insights from your object data.

You can enable the search integration service for any versioned or unversioned bucket. Search integration is configured by associating metadata notification configuration XML with the bucket that specifies which objects to act on and the destination for the object metadata.

Notifications are generated in the form of a JSON document named with the bucket name, object name, and version ID, if any. Each metadata notification contains a standard set of system metadata for the object in addition to all of the object's tags and user metadata.

 For tags and user metadata, StorageGRID passes dates and numbers to Elasticsearch as strings or as S3 event notifications. To configure Elasticsearch to interpret these strings as dates or numbers, follow the Elasticsearch instructions for dynamic field mapping and for mapping date formats. You must enable the dynamic field mappings on the index before you configure the search integration service. After a document is indexed, you cannot edit the document's field types in the index.

Notifications are generated and queued for delivery whenever:

- An object is created.
- An object is deleted, including when objects are deleted as a result of the operation of the grid's ILM policy.
- Object metadata or tags are added, updated, or deleted. The complete set of metadata and tags is always sent on update — not just the changed values.

After you add metadata notification configuration XML to a bucket, notifications are sent for any new objects that you create and for any objects that you modify by updating its data, user metadata, or tags. However, notifications are not sent for any objects that were already in the bucket. To ensure that object metadata for all objects in the bucket is sent to the destination, you should do either of the following:

- Configure the search integration service immediately after creating the bucket and before adding any objects.
- Perform an action on all objects already in the bucket that will trigger a metadata notification message to be

sent to the destination.

The StorageGRID search integration service supports an Elasticsearch cluster as a destination. As with the other platform services, the destination is specified in the endpoint whose URN is used in the configuration XML for the service. Use the [NetApp Interoperability Matrix Tool](#) to determine the supported versions of Elasticsearch.

Related information

[Configuration XML for search integration](#)

[Object metadata included in metadata notifications](#)

[JSON generated by search integration service](#)

[Configure search integration service](#)

Considerations for using platform services

Before implementing platform services, review the recommendations and considerations for using these services.

For information about S3, see [Use S3](#).

Considerations for using platform services

Consideration	Details
Destination endpoint monitoring	You must monitor the availability of each destination endpoint. If connectivity to the destination endpoint is lost for an extended period of time and a large backlog of requests exists, additional client requests (such as PUT requests) to StorageGRID will fail. You must retry these failed requests when the endpoint becomes reachable.
Destination endpoint throttling	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.

Consideration	Details
Ordering guarantees	<p>StorageGRID guarantees ordering of operations on an object within a site. As long as all operations against an object are within the same site, the final object state (for replication) will always equal the state in StorageGRID.</p> <p>StorageGRID makes a best effort attempt to order requests when operations are made across StorageGRID sites. For example, if you write an object initially to site A and then later overwrite the same object at site B, the final object replicated by CloudMirror to the destination bucket is not guaranteed to be the newer object.</p>
ILM-driven object deletions	<p>To match the deletion behavior of the AWS CRR and SNS services, CloudMirror and event notification requests are not sent when an object in the source bucket is deleted because of StorageGRID ILM rules. For example, no CloudMirror or event notifications requests are sent if an ILM rule deletes an object after 14 days.</p> <p>In contrast, search integration requests are sent when objects are deleted because of ILM.</p>

Considerations for using CloudMirror replication service

Consideration	Details
Replication status	StorageGRID does not support the <code>x-amz-replication-status</code> header.
Object size	<p>The maximum size for objects that can be replicated to a destination bucket by the CloudMirror replication service is 5 TiB, which is the same as the maximum <i>supported</i> object size.</p> <p>Note: The maximum <i>recommended</i> size for a single PUT Object operation is 5 GiB (5,368,709,120 bytes). If you have objects that are larger than 5 GiB, use multipart upload instead.</p>
Bucket versioning and version IDs	<p>If the source S3 bucket in StorageGRID has versioning enabled, you should also enable versioning for the destination bucket.</p> <p>When using versioning, note that the ordering of object versions in the destination bucket is best effort and not guaranteed by the CloudMirror service, due to limitations in the S3 protocol.</p> <p>Note: Version IDs for the source bucket in StorageGRID are not related to the version IDs for the destination bucket.</p>

Consideration	Details
Tagging for object versions	<p>The CloudMirror service does not replicate any PUT Object tagging or DELETE Object tagging requests that supply a version ID, due to limitations in the S3 protocol. Because version IDs for the source and destination are not related, there is no way to ensure that a tag update to a specific version ID will be replicated.</p> <p>In contrast, the CloudMirror service does replicate PUT Object tagging requests or DELETE Object tagging requests that do not specify a version ID. These requests update the tags for the latest key (or the latest version if the bucket is versioned). Normal ingests with tags (not tagging updates) are also replicated.</p>
Multipart uploads and ETag values	<p>When mirroring objects that were uploaded using a multipart upload, the CloudMirror service does not preserve the parts. As a result, the ETag value for the mirrored object will be different than the ETag value of the original object.</p>
Objects encrypted with SSE-C (server-side encryption with customer-provided keys)	<p>The CloudMirror service does not support objects that are encrypted with SSE-C. If you attempt to ingest an object into the source bucket for CloudMirror replication and the request includes the SSE-C request headers, the operation fails.</p>
Bucket with S3 Object Lock enabled	<p>If the destination S3 bucket for CloudMirror replication has S3 Object Lock enabled, the attempt to configure bucket replication (PUT Bucket replication) will fail with an AccessDenied error.</p>

Configure platform services endpoints

Before you can configure a platform service for a bucket, you must configure at least one endpoint to be the destination for the platform service.

Access to platform services is enabled on a per-tenant basis by a StorageGRID administrator. To create or use a platform services endpoint, you must be a tenant user with Manage Endpoints or Root Access permission, in a grid whose networking has been configured to allow Storage Nodes to access external endpoint resources. Contact your StorageGRID administrator for more information.

What is a platform services endpoint?

When you create a platform services endpoint, you specify the information that StorageGRID needs to access the external destination.

For example, if you want to replicate objects from a StorageGRID bucket to an AWS S3 bucket, you create a platform services endpoint that includes the information and credentials StorageGRID needs to access the destination bucket on AWS.

Each type of platform service requires its own endpoint, so you must configure at least one endpoint for each platform service you plan to use. After defining a platform services endpoint, you use the endpoint's URN as the destination in the configuration XML used to enable the service.

You can use the same endpoint as the destination for more than one source bucket. For example, you could configure several source buckets to send object metadata to the same search integration endpoint so that you can perform searches across multiple buckets. You can also configure a source bucket to use more than one

endpoint as a target, which enables you to do things like send notifications about object creation to one SNS topic and notifications about object deletion to a second SNS topic.

Endpoints for CloudMirror replication

StorageGRID supports replication endpoints that represent S3 buckets. These buckets might be hosted on Amazon Web Services, the same or a remote StorageGRID deployment, or another service.

Endpoints for notifications

StorageGRID supports Simple Notification Service (SNS) endpoints. Simple Queue Service (SQS) or AWS Lambda endpoints are not supported.

Endpoints for the search integration service

StorageGRID supports search integration endpoints that represent Elasticsearch clusters. These Elasticsearch clusters can be in a local datacenter or hosted in an AWS cloud or elsewhere.

The search integration endpoint refers to a specific Elasticsearch index and type. You must create the index in Elasticsearch before creating the endpoint in StorageGRID, or endpoint creation will fail. You do not need to create the type before creating the endpoint. StorageGRID will create the type if required when it sends object metadata to the endpoint.

Related information

[Administer StorageGRID](#)

Specify URN for platform services endpoint

When you create a platform services endpoint, you must specify a Unique Resource Name (URN). You will use the URN to reference the endpoint when you create configuration XML for the platform service. The URN for each endpoint must be unique.

StorageGRID validates platform services endpoints as you create them. Before you create a platform services endpoint, confirm that the resource specified in the endpoint exists and that it can be reached.

URN elements

The URN for a platform services endpoint must start with either `arn:aws` or `urn:mysite`, as follows:

- If the service is hosted on Amazon Web Services (AWS), use `arn:aws`.
- If the service is hosted on Google Cloud Platform (GCP), use `arn:aws`.
- If the service is hosted locally, use `urn:mysite`

For example, if you are specifying the URN for a CloudMirror endpoint hosted on StorageGRID, the URN might begin with `urn:sgws`.

The next element of the URN specifies the type of platform service, as follows:

Service	Type
CloudMirror replication	s3

Service	Type
Notifications	sns
Search integration	es

For example, to continue specifying the URN for a CloudMirror endpoint hosted on StorageGRID, you would add s3 to get `urn:sgws:s3`.

The final element of the URN identifies the specific target resource at the destination URI.

Service	Specific resource
CloudMirror replication	bucket-name
Notifications	sns-topic-name
Search integration	domain-name/index-name/type-name

Note: If the Elasticsearch cluster is **not** configured to create indexes automatically, you must create the index manually before you create the endpoint.

URNs for services hosted on AWS and GCP

For AWS and GCP entities, the complete URN is a valid AWS ARN. For example:

- CloudMirror replication:

```
arn:aws:s3:::bucket-name
```

- Notifications:

```
arn:aws:sns:region:account-id:topic-name
```

- Search integration:

```
arn:aws:es:region:account-id:domain/domain-name/index-name/type-name
```



For an AWS search integration endpoint, the `domain-name` must include the literal string `domain/`, as shown here.

URNs for locally-hosted services

When using locally-hosted services instead of cloud services, you can specify the URN in any way that creates

a valid and unique URN, as long as the URN includes the required elements in the third and final positions. You can leave the elements indicated by optional blank, or you can specify them in any way that helps you identify the resource and make the URN unique. For example:

- CloudMirror replication:

```
urn:mysite:s3:optional:optional:bucket-name
```

For a CloudMirror endpoint hosted on StorageGRID, you can specify a valid URN that begins with `urn:sgws:`

```
urn:sgws:s3:optional:optional:bucket-name
```

- Notifications:

```
urn:mysite:sns:optional:optional:sns-topic-name
```

- Search integration:

```
urn:mysite:es:optional:optional:domain-name/index-name/type-name
```



For locally-hosted search integration endpoints, the `domain-name` element can be any string as long as the URN of the endpoint is unique.

Create platform services endpoint

You must create at least one endpoint of the correct type before you can enable a platform service.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- Platform services must be enabled for your tenant account by a StorageGRID administrator.
- You must belong to a user group that has the Manage Endpoints permission.
- The resource referenced by the platform services endpoint must have been created:
 - CloudMirror replication: S3 bucket
 - Event notification: SNS topic
 - Search notification: Elasticsearch index, if the destination cluster is not configured to automatically create indexes.
- You must have the information about the destination resource:
 - Host and port for the Uniform Resource Identifier (URI)



If you plan to use a bucket hosted on a StorageGRID system as an endpoint for CloudMirror replication, contact the grid administrator to determine the values you need to enter.

- Unique Resource Name (URN)

[Specify URN for platform services endpoint](#)

- Authentication credentials (if required):
 - Access Key: Access key ID and secret access key
 - Basic HTTP: Username and password
 - CAP (C2S Access Portal): Temporary credentials URL, server and client certificates, client keys, and an optional client private key passphrase.
- Security certificate (if using a custom CA certificate)

Steps

1. Select **STORAGE (S3) > Platform services endpoints**.

The Platform services endpoints page appears.

Platform services endpoints

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

0 endpoints

[Create endpoint](#)

[Delete endpoint](#)

Display name

Last error

Type

URI

URN

No endpoints found

[Create endpoint](#)

2. Select **Create endpoint**.

Create endpoint

1 Enter details 2 Select authentication type 3 Verify server

Optional Optional

Enter endpoint details

Enter the endpoint's display name, URI, and URN.

Display name [?](#)

URI [?](#)

URN [?](#)

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

3. Enter a display name to briefly describe the endpoint and its purpose.

The type of platform service that the endpoint supports is shown beside the endpoint name when it is listed on the Endpoints page, so you do not need to include that information in the name.

4. In the **URI** field, specify the Unique Resource Identifier (URI) of the endpoint.

Use one of the following formats:

```
https://host:port  
http://host:port
```

If you do not specify a port, port 443 is used for HTTPS URLs and port 80 is used for HTTP URLs.

For example, the URI for a bucket hosted on StorageGRID might be:

```
https://s3.example.com:10443
```

In this example, `s3.example.com` represents the DNS entry for the virtual IP (VIP) of the StorageGRID high availability (HA) group, and 10443 represents the port defined in the load balancer endpoint.



Whenever possible, you should connect to an HA group of load-balancing nodes to avoid a single point of failure.

Similarly, the URI for a bucket hosted on AWS might be:

```
https://s3-aws-region.amazonaws.com
```



If the endpoint is used for the CloudMirror replication service, do not include the bucket name in the URI. You include the bucket name in the **URN** field.


5. Enter the Unique Resource Name (URN) for the endpoint.



You cannot change an endpoint's URN after the endpoint has been created.

6. Select **Continue**.

7. Select a value for **Authentication type**, and then enter or upload the required credentials.



The credentials that you supply must have write permissions for the destination resource.

Authentication type	Description	Credentials
Anonymous	Provides anonymous access to the destination. Only works for endpoints that have security disabled.	No authentication.
Access Key	Uses AWS-style credentials to authenticate connections with the destination.	<ul style="list-style-type: none"> Access key ID Secret access key
Basic HTTP	Uses a username and password to authenticate connections to the destination.	<ul style="list-style-type: none"> Username Password
CAP (C2S Access Portal)	Uses certificates and keys to authenticate connections to the destination.	<ul style="list-style-type: none"> Temporary credentials URL Server CA certificate (PEM file upload) Client certificate (PEM file upload) Client private key (PEM file upload, OpenSSL encrypted format or unencrypted private key format) Client private key passphrase (optional)

8. Select **Continue**.

9. Select a radio button for **Verify server** to choose how TLS connection to the endpoint is verified.

Create endpoint

Enter details Select authentication type Verify server

Verify server

Use this method to validate the certificate for TLS connections to the endpoint resource. If you select "Use custom CA certificate," copy and paste the custom security certificate in the text box.

Use custom CA certificate

Use operating system CA certificate

Do not verify certificate

```
-----BEGIN CERTIFICATE-----
abcdefhijkl123456780ABCDEFGHIJKLM
123456/7890ABCDEFabcdefhijkl1ABCD
-----END CERTIFICATE-----
```

[Previous](#) [Test and create endpoint](#)

Type of certificate verification	Description
Use custom CA certificate	Use a custom security certificate. If you select this setting, copy and paste the custom security certificate in the CA Certificate text box.
Use operating system CA certificate	Use the default Grid CA certificate installed on the operating system to secure connections.
Do not verify certificate	The certificate used for the TLS connection is not verified. This option is not secure.

10. Select **Test and create endpoint**.

- A success message appears if the endpoint can be reached using the specified credentials. The connection to the endpoint is validated from one node at each site.
- An error message appears if endpoint validation fails. If you need to modify the endpoint to correct the error, select **Return to endpoint details** and update the information. Then, select **Test and create endpoint**.



Endpoint creation fails if platform services are not enabled for your tenant account. Contact your StorageGRID administrator.

After you have configured an endpoint, you can use its URN to configure a platform service.

Related information

[Specify URN for platform services endpoint](#)

[Configure CloudMirror replication](#)

[Configure event notifications](#)

[Configure search integration service](#)

Test connection for platform services endpoint

If the connection to a platform service has changed, you can test the connection for the endpoint to validate that the destination resource exists and that it can be reached using the credentials you specified.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must belong to a user group that has the Manage Endpoints permission.

About this task

StorageGRID does not validate that the credentials have the correct permissions.

Steps

1. Select **STORAGE (S3) > Platform services endpoints**.

The Platform services endpoints page appears and shows the list of platform services endpoints that have already been configured.

Platform services endpoints

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

4 endpoints


[Create endpoint](#)

[Delete endpoint](#)

<input type="checkbox"/>	Display name	Last error	Type	URI	URN
<input type="checkbox"/>	my-endpoint-1		S3 Bucket	http://10.96.104.167:10443	urn:sgws:s3:::bucket1
<input type="checkbox"/>	my-endpoint-2	✗ 2 hours ago	Search	http://10.96.104.30:9200	urn:sgws:es:::mydomain/sveloso/_doc
<input type="checkbox"/>	my-endpoint-3		Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example1
<input type="checkbox"/>	my-endpoint-4		S3 Bucket	http://10.96.104.167:10443	urn:sgws:s3:::bucket2

2. Select the endpoint whose connection you want to test.

The endpoint details page appears.



Display name: **my-endpoint-1** 

Type: **S3 Bucket**

URI: **http://10.96.104.167:10443**

URN: **urn:sgws:s3:::bucket1**

Connection **Configuration**

Verify connection 

Some errors might continue to appear after they are resolved. To see if an error is current or to force the removal of a resolved error, select **Test connection**.

Test connection

3. Select **Test connection**.

- A success message appears if the endpoint can be reached using the specified credentials. The connection to the endpoint is validated from one node at each site.
- An error message appears if endpoint validation fails. If you need to modify the endpoint to correct the error, select **Configuration** and update the information. Then, select **Test and save changes**.

Edit platform services endpoint

You can edit the configuration for a platform services endpoint to change its name, URI, or other details. For example, you might need to update expired credentials or change the URI to point to a backup Elasticsearch index for failover. You cannot change the URN for a platform services endpoint.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must belong to a user group that has the Manage Endpoints permission. See [Tenant management permissions](#).

Steps

1. Select **STORAGE (S3) > Platform services endpoints**.

The Platform services endpoints page appears and shows the list of platform services endpoints that have

already been configured.

Platform services endpoints

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

4 endpoints

[Create endpoint](#)

[Delete endpoint](#)

<input type="checkbox"/>	Display name <small>?</small>	Last error <small>?</small>	Type <small>?</small>	URI <small>?</small>	URN <small>?</small>
<input type="checkbox"/>	my-endpoint-1		S3 Bucket	http://10.96.104.167:10443	urn:sgws:s3:::bucket1
<input type="checkbox"/>	my-endpoint-2	✖ 2 hours ago	Search	http://10.96.104.30:9200	urn:sgws:es:::mydomain/sveloso/_doc
<input type="checkbox"/>	my-endpoint-3		Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example1
<input type="checkbox"/>	my-endpoint-4		S3 Bucket	http://10.96.104.167:10443	urn:sgws:s3:::bucket2

2. Select the endpoint you want to edit.

The endpoint details page appears.

3. Select **Configuration**.

Overview

Display name:	my-endpoint-3 
Type:	Notifications
URI:	http://10.96.104.202:8080/
URN:	arn:aws:sns:us-west-2::example1

[Connection](#)

[Configuration](#)

Edit configuration

Endpoint details

URI 

http://10.96.104.202:8080/

URN 

arn:aws:sns:us-west-2::example1

Authentication type

Basic HTTP

Username 

testme

Password 

[Edit password](#)

Verify server

Use custom CA certificate

Use operating system CA certificate

Do not verify certificate

```
-----BEGIN CERTIFICATE-----  
abcdefghijklmnopqrstuvwxyz0ABCDEFGHIJKLMNOPQRSTUVWXYZ  
-----END CERTIFICATE-----
```

[Test and save changes](#)

4. As needed, change the configuration of the endpoint.



You cannot change an endpoint's URN after the endpoint has been created.

- a. To change the display name for the endpoint, select the edit icon
- b. As needed, change the URI.
- c. As needed, change the authentication type.
 - For Access Key authentication, change the key as necessary by selecting **Edit S3 key** and pasting a new access key ID and secret access key. If you need to cancel your changes, select **Revert S3 key edit**.
 - For Basic HTTP authentication, change the username as needed. Change the password as needed by selecting **Edit password** and entering the new password. If you need to cancel your changes, select **Revert password edit**.
 - For CAP (C2S Access Portal) authentication, change the temporary credentials URL or optional client private key passphrase and upload new certificate and key files as needed.



The Client private key must be in OpenSSL encrypted format or unencrypted private key format.

- d. As needed, change the method for verifying the server.

5. Select **Test and save changes**.

- A success message appears if the endpoint can be reached using the specified credentials. The connection to the endpoint is verified from one node at each site.
- An error message appears if endpoint validation fails. Modify the endpoint to correct the error, and then select **Test and save changes**.

Delete platform services endpoint

You can delete an endpoint if you no longer want to use the associated platform service.

What you'll need

- You must be signed in to the Tenant Manager using a [supported web browser](#).
- You must belong to a user group that has the **Manage Endpoints** permission. See [Tenant management permissions](#).

Steps

1. Select **STORAGE (S3) > Platform services endpoints**.

The Platform services endpoints page appears and shows the list of platform services endpoints that have already been configured.

Platform services endpoints

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

4 endpoints

[Create endpoint](#)

[Delete endpoint](#)

<input type="checkbox"/>	Display name	Last error	Type	URI	URN
<input type="checkbox"/>	my-endpoint-1		S3 Bucket	http://10.96.104.167:10443	urn:sgws:s3:::bucket1
<input type="checkbox"/>	my-endpoint-2	✖ 2 hours ago	Search	http://10.96.104.30:9200	urn:sgws:es:::mydomain/sveloso/_doc
<input type="checkbox"/>	my-endpoint-3		Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example1
<input type="checkbox"/>	my-endpoint-4		S3 Bucket	http://10.96.104.167:10443	urn:sgws:s3:::bucket2

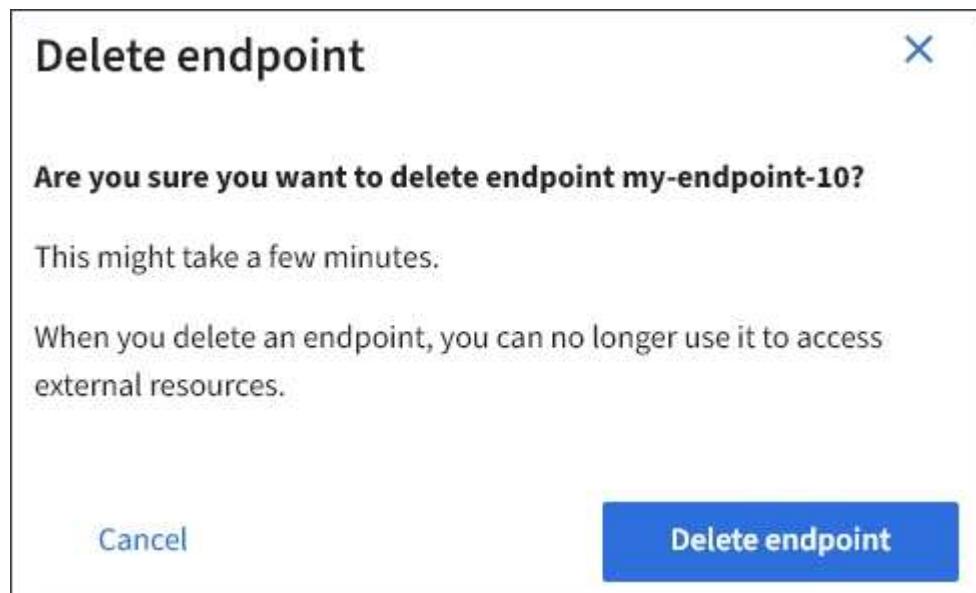
2. Select the check box for each endpoint you want to delete.



If you delete a platform services endpoint that is in use, the associated platform service will be disabled for any buckets that use the endpoint. Any requests that have not yet been completed will be dropped. Any new requests will continue to be generated until you change your bucket configuration to no longer reference the deleted URN. StorageGRID will report these requests as unrecoverable errors.

3. Select **Actions > Delete endpoint**.

A confirmation message appears.



4. Select **Delete endpoint**.

Troubleshoot platform services endpoint errors

If an error occurs when StorageGRID attempts to communicate with a platform services endpoint, a message is displayed on the Dashboard. On the Platform services endpoints page, the Last error column indicates how long ago the error occurred. No error is displayed if the permissions associated with an endpoint's credentials are incorrect.

Determine if error has occurred

If any platform services endpoint errors have occurred within the past 7 days, the Tenant Manager Dashboard displays an alert message. You can go the Platform services endpoints page to see more details about the error.

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

The same error that appears on the Dashboard also appears at the top of the Platform services endpoints page. To view a more detailed error message:

Steps

1. From the list of endpoints, select the endpoint that has the error.
2. On the endpoint details page, select **Connection**. This tab displays only the most recent error for an endpoint and indicates how long ago the error occurred. Errors that include the red X icon  occurred within the past 7 days.

Overview

Display name:	my-endpoint-2 
Type:	Search
URI:	http://10.96.104.30:9200
URN:	urn:sgws:es:::mydomain/sveloso/_doc

Connection

Configuration

Verify connection

Some errors might continue to appear after they are resolved. To see if an error is current or to force the removal of a resolved error, select **Test connection**.

[Test connection](#)

Last error details

 2 hours ago

Endpoint failure: Endpont has an AWS failure: RequestError: send request failed; caused by: url.Error; caused by: net:OpError; caused by: os.SyscallError (logID: 143H5UDUUKMGDRWJ)

Check if error is still current

Some errors might continue to be shown in the **Last error** column even after they are resolved. To see if an error is current or to force the removal of a resolved error from the table:

Steps

1. Select the endpoint.

The endpoint details page appears.

2. Select **Connection > Test connection**.

Selecting **Test connection** causes StorageGRID to validate that the platform services endpoint exists and that it can be reached with the current credentials. The connection to the endpoint is validated from one node at each site.

Resolve endpoint errors

You can use the **Last error** message on the endpoint details page to help determine what is causing the error. Some errors might require you to edit the endpoint to resolve the issue. For example, a CloudMirroring error can occur if StorageGRID is unable to access the destination S3 bucket because it does not have the correct

access permissions or the access key has expired. The message is “Either the endpoint credentials or the destination access needs to be updated,” and the details are “AccessDenied” or “InvalidAccessKeyId.”

If you need to edit the endpoint to resolve an error, selecting **Test and save changes** causes StorageGRID to validate the updated endpoint and confirm that it can be reached with the current credentials. The connection to the endpoint is validated from one node at each site.

Steps

1. Select the endpoint.
2. On the endpoint details page, select **Configuration**.
3. Edit the endpoint configuration as needed.
4. Select **Connection > Test connection**.

Endpoint credentials with insufficient permissions

When StorageGRID validates a platform services endpoint, it confirms that the endpoint’s credentials can be used to contact the destination resource and it does a basic permissions check. However, StorageGRID does not validate all of the permissions required for certain platform services operations. For this reason, if you receive an error when attempting to use a platform service (such as “403 Forbidden”), check the permissions associated with the endpoint’s credentials.

Additional platform services troubleshooting

For additional information about troubleshooting platform services, see the instructions for administering StorageGRID.

[Administer StorageGRID](#)

Related information

[Create platform services endpoint](#)

[Test connection for platform services endpoint](#)

[Edit platform services endpoint](#)

Configure CloudMirror replication

The [CloudMirror replication service](#) is one of the three StorageGRID platform services. You can use CloudMirror replication to automatically replicate objects to an external S3 bucket.

What you’ll need

- Platform services must be enabled for your tenant account by a StorageGRID administrator.
- You must have already created a bucket to act as the replication source.
- The endpoint that you intend to use as a destination for CloudMirror replication must already exist, and you must have its URN.
- You must belong to a user group that has the Manage All Buckets or the Root Access permission, which allows you to manage the settings for all S3 buckets in your tenant account. These permissions override the permission settings in group or bucket policies when configuring the bucket using the Tenant Manager.

About this task

CloudMirror replication copies objects from a source bucket to a destination bucket that is specified in an endpoint. To enable CloudMirror replication for a bucket, you must create and apply valid bucket replication configuration XML. The replication configuration XML must use the URN of an S3 bucket endpoint for each destination.



Replication is not supported for source or destination buckets with S3 Object Lock enabled.

For general information on bucket replication and how to configure it, see the Amazon Simple Storage Service (S3) documentation on cross-region replication (CRR). For information on how StorageGRID implements the S3 bucket replication configuration API, see the [instructions for implementing S3 client applications](#).

If you enable CloudMirror replication on a bucket that contains objects, new objects added to the bucket are replicated, but the existing objects in the bucket are not. You must update existing objects to trigger replication.

If you specify a storage class in the replication configuration XML, StorageGRID uses that class when performing operations against the destination S3 endpoint. The destination endpoint must also support the specified storage class. Be sure to follow any recommendations provided by the destination system vendor.

Steps

1. Enable replication for your source bucket:

Use a text editor to create the replication configuration XML required to enable replication, as specified in the S3 replication API. When configuring the XML:

- Note that StorageGRID only supports V1 of the replication configuration. This means that StorageGRID does not support the use of the `Filter` element for rules, and follows V1 conventions for deletion of object versions. See the Amazon documentation on replication configuration for details.
- Use the URN of an S3 bucket endpoint as the destination.
- Optionally add the `<StorageClass>` element, and specify one of the following:
 - **STANDARD**: The default storage class. If you do not specify a storage class when you upload an object, the STANDARD storage class is used.
 - **STANDARD_IA**: (Standard - infrequent access.) Use this storage class for data that is accessed less frequently, but that still requires rapid access when needed.
 - **REDUCED_REDUNDANCY**: Use this storage class for noncritical, reproducible data that can be stored with less redundancy than the STANDARD storage class.
- If you specify a `Role` in the configuration XML it will be ignored. This value is not used by StorageGRID.

```
<ReplicationConfiguration>
  <Role></Role>
  <Rule>
    <Status>Enabled</Status>
    <Prefix>2020</Prefix>
    <Destination>
      <Bucket>urn:sgws:s3:::2017-records</Bucket>
      <StorageClass>STANDARD</StorageClass>
    </Destination>
  </Rule>
</ReplicationConfiguration>
```

2. In the Tenant Manager select **STORAGE (S3) > Buckets**.

3. Select the name of the source bucket.

The bucket details page appears.

4. Select **Platform services > Replication**.

5. Select the **Enable replication** check box.

6. Paste the replication configuration XML into the text box, and select **Save changes**.

Bucket options Bucket access Platform services

Replication Disabled ^

Enable the CloudMirror replication service to copy objects from a source bucket to a destination bucket that is specified in an endpoint.

- Platform services must be enabled for your tenant account by a StorageGRID administrator.
- You must have already configured an endpoint for each destination bucket.
- You must specify the URN of each endpoint in the replication configuration XML for the source bucket.

Enable replication

Clear

```
<ReplicationConfiguration>
  <Role></Role>
  <Rule>
    <Status>Enabled</Status>
    <Prefix>2020</Prefix>
    <Destination>
      <Bucket>urn:sgws:s3:::2017-records</Bucket>
      <StorageClass>STANDARD</StorageClass>
    </Destination>
  </Rule>
</ReplicationConfiguration>
```

Save changes



Platform services must be enabled for each tenant account by a StorageGRID administrator using the Grid Manager or Grid Management API. Contact your StorageGRID administrator if an error occurs when you save the configuration XML.

7. Verify that replication is configured correctly:

- Add an object to the source bucket that meets the requirements for replication as specified in the replication configuration.

In the example shown earlier, objects that match the prefix “2020” are replicated.

- Confirm that the object has been replicated to the destination bucket.

For small objects, replication happens quickly.

Related information

[Use S3](#)

[Create platform services endpoint](#)

Configure event notifications

The notifications service is one of the three StorageGRID platform services. You can enable notifications for a bucket to send information about specified events to a destination service that supports the AWS Simple Notification Service™ (SNS).

What you'll need

- Platform services must be enabled for your tenant account by a StorageGRID administrator.
- You must have already created a bucket to act as the source of notifications.
- The endpoint that you intend to use as a destination for event notifications must already exist, and you must have its URN.
- You must belong to a user group that has the Manage All Buckets or the Root Access permission, which allows you to manage the settings for all S3 buckets in your tenant account. These permissions override the permission settings in group or bucket policies when configuring the bucket using the Tenant Manager.

About this task

After you configure event notifications, whenever a specified event occurs for an object in the source bucket, a notification is generated and sent to the Simple Notification Service (SNS) topic used as the destination endpoint. To enable notifications for a bucket, you must create and apply valid notification configuration XML. The notification configuration XML must use the URN of an event notifications endpoint for each destination.

For general information on event notifications and how to configure them, see Amazon documentation. For information on how StorageGRID implements the S3 bucket notification configuration API, see the instructions for implementing S3 client applications.

If you enable event notifications for a bucket that contains objects, notifications are sent only for actions that are performed after the notification configuration is saved.

Steps

1. Enable notifications for your source bucket:

- Use a text editor to create the notification configuration XML required to enable event notifications, as specified in the S3 notification API.
- When configuring the XML, use the URN of an event notifications endpoint as the destination topic.

```
<NotificationConfiguration>
  <TopicConfiguration>
    <Id>Image-created</Id>
    <Filter>
      <S3Key>
        <FilterRule>
          <Name>prefix</Name>
          <Value>images/</Value>
        </FilterRule>
      </S3Key>
    </Filter>
    <Topic>arn:aws:sns:us-east-1:050340950352:sgws-topic</Topic>
    <Event>s3:ObjectCreated:*</Event>
  </TopicConfiguration>
</NotificationConfiguration>
```

2. In the Tenant Manager select **STORAGE (S3) > Buckets**.
3. Select the name of the source bucket.

The bucket details page appears.

4. Select **Platform services > Event notifications**.
5. Select the **Enable event notifications** check box.
6. Paste the notification configuration XML into the text box, and select **Save changes**.

Bucket options Bucket access Platform services

Replication Disabled

Event notifications Disabled

Enable the event notification service for an S3 bucket if you want StorageGRID to send notifications about specified events to a destination Amazon Simple Notification Service (SNS).

- Platform services must be enabled for your tenant account by a StorageGRID administrator.
- You must have already configured an endpoint for the destination of event notifications.
- You must specify the URN of that endpoint in the notification configuration XML for the source bucket.

Enable event notifications

Clear

```
<NotificationConfiguration>
  <TopicConfiguration>
    <Id>Image-created</Id>
    <Filter>
      <S3Key>
        <FilterRule>
          <Name>prefix</Name>
          <Value>images/</Value>
        </FilterRule>
      </S3Key>
    </Filter>
    <Topic>arn:aws:sns:us-east-1:050340950352:sgws-topic</Topic>
  </TopicConfiguration>
</NotificationConfiguration>
```

Save changes



Platform services must be enabled for each tenant account by a StorageGRID administrator using the Grid Manager or Grid Management API. Contact your StorageGRID administrator if an error occurs when you save the configuration XML.

7. Verify that event notifications are configured correctly:

- Perform an action on an object in the source bucket that meets the requirements for triggering a notification as configured in the configuration XML.

In the example, an event notification is sent whenever an object is created with the `images/` prefix.

b. Confirm that a notification has been delivered to the destination SNS topic.

For example, if your destination topic is hosted on the AWS Simple Notification Service (SNS), you could configure the service to send you an email when the notification is delivered.

```
{
  "Records": [
    {
      "eventVersion": "2.0",
      "eventSource": "sgws:s3",
      "eventTime": "2017-08-08T23:52:38Z",
      "eventName": "ObjectCreated:Put",
      "userIdentity": {
        "principalId": "11111111111111111111"
      },
      "requestParameters": {
        "sourceIPAddress": "193.51.100.20"
      },
      "responseElements": {
        "x-amz-request-id": "122047343"
      },
      "s3": {
        "s3SchemaVersion": "1.0",
        "configurationId": "Image-created",
        "bucket": {
          "name": "test1",
          "ownerIdentity": {
            "principalId": "11111111111111111111"
          },
          "arn": "arn:sgws:s3:::test1"
        },
        "object": {
          "key": "images/cat.jpg",
          "size": 0,
          "eTag": "d41d8cd98f00b204e9800998ecf8427e",
          "sequencer": "14D90402421461C7"
        }
      }
    }
  ]
}
```

If the notification is received at the destination topic, you have successfully configured your source bucket for StorageGRID notifications.

Related information

[Understand notifications for buckets](#)

[Use S3](#)

[Create platform services endpoint](#)

Use search integration service

The search integration service is one of the three StorageGRID platform services. You can enable this service to send object metadata to a destination search index whenever an object is created, deleted, or its metadata or tags are updated.

You can configure search integration by using the Tenant Manager to apply custom StorageGRID configuration XML to a bucket.

 Because the search integration service causes object metadata to be sent to a destination, its configuration XML is referred to as *metadata notification configuration XML*. This configuration XML is different than the *notification configuration XML* used to enable event notifications.

See the [instructions for implementing S3 client applications](#) for details about the following custom StorageGRID S3 REST API operations:

- `DELETE` Bucket metadata notification configuration request
- `GET` Bucket metadata notification configuration request
- `PUT` Bucket metadata notification configuration request

Related information

[Configuration XML for search integration](#)

[Object metadata included in metadata notifications](#)

[JSON generated by search integration service](#)

[Configure search integration service](#)

[Use S3](#)

Configuration XML for search integration

The search integration service is configured using a set of rules contained within `<MetadataNotificationConfiguration>` and `</MetadataNotificationConfiguration>` tags. Each rule specifies the objects that the rule applies to, and the destination where StorageGRID should send those objects' metadata.

Objects can be filtered on the prefix of the object name. For example, you could send metadata for objects with the prefix `images` to one destination, and metadata for objects with the prefix `videos` to another. Configurations that have overlapping prefixes are not valid, and are rejected when they are submitted. For example, a configuration that includes one rule for objects with the prefix `test` and a second rule for objects with the prefix `test2` is not allowed.

Destinations must be specified using the URN of a StorageGRID endpoint that has been created for the search integration service. These endpoints refer to an index and type defined on an Elasticsearch cluster.

```
<MetadataNotificationConfiguration>
  <Rule>
    <ID>Rule-1</ID>
    <Status>rule-status</Status>
    <Prefix>key-prefix</Prefix>
    <Destination>
      <Urn>arn:aws:es:region:account-
ID:domain/mydomain/myindex/mytype</Urn>
    </Destination>
  </Rule>
  <Rule>
    <ID>Rule-2</ID>
    ...
  </Rule>
  ...
</MetadataNotificationConfiguration>
```

The table describes the elements in the metadata notification configuration XML.

Name	Description	Required
MetadataNotificationConfiguration	Container tag for rules used to specify the objects and destination for metadata notifications. Contains one or more Rule elements.	Yes
Rule	Container tag for a rule that identifies the objects whose metadata should be added to a specified index. Rules with overlapping prefixes are rejected. Included in the MetadataNotificationConfiguration element.	Yes
ID	Unique identifier for the rule. Included in the Rule element.	No
Status	Status can be 'Enabled' or 'Disabled'. No action is taken for rules that are disabled. Included in the Rule element.	Yes

Name	Description	Required
Prefix	<p>Objects that match the prefix are affected by the rule, and their metadata is sent to the specified destination.</p> <p>To match all objects, specify an empty prefix.</p> <p>Included in the Rule element.</p>	Yes
Destination	<p>Container tag for the destination of a rule.</p> <p>Included in the Rule element.</p>	Yes
Urn	<p>URN of the destination where object metadata is sent. Must be the URN of a StorageGRID endpoint with the following properties:</p> <ul style="list-style-type: none"> • es must be the third element. • The URN must end with the index and type where the metadata is stored, in the form domain-name/myindex/mytype. <p>Endpoints are configured using the Tenant Manager or Tenant Management API. They take the following form:</p> <ul style="list-style-type: none"> • arn:aws:es:region:account-ID:domain/mydomain/myindex/mytype • urn:mysite:es:::mydomain/myindex/mytype <p>The endpoint must be configured before the configuration XML is submitted, or configuration will fail with a 404 error.</p> <p>URN is included in the Destination element.</p>	Yes

Use the sample metadata notification configuration XML to learn how to construct your own XML.

Metadata notification configuration that applies to all objects

In this example, object metadata for all objects is sent to the same destination.

```

<MetadataNotificationConfiguration>
  <Rule>
    <ID>Rule-1</ID>
    <Status>Enabled</Status>
    <Prefix></Prefix>
    <Destination>
      <Urn>urn:myes:es:::sgws-notifications/test1/all</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>

```

Metadata notification configuration with two rules

In this example, object metadata for objects that match the prefix `/images` is sent to one destination, while object metadata for objects that match the prefix `/videos` is sent to a second destination.

```

<MetadataNotificationConfiguration>
  <Rule>
    <ID>Images-rule</ID>
    <Status>Enabled</Status>
    <Prefix>/images</Prefix>
    <Destination>
      <Urn>arn:aws:es:us-east-1:3333333:domain/es-
domain/graphics/imagetype</Urn>
    </Destination>
  </Rule>
  <Rule>
    <ID>Videos-rule</ID>
    <Status>Enabled</Status>
    <Prefix>/videos</Prefix>
    <Destination>
      <Urn>arn:aws:es:us-west-1:2222222:domain/es-
domain/graphics/videotype</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>

```

Related information

[Use S3](#)

[Object metadata included in metadata notifications](#)

[JSON generated by search integration service](#)

[Configure search integration service](#)

Configure the search integration service

The search integration service sends object metadata to a destination search index whenever an object is created, deleted, or its metadata or tags are updated.

What you'll need

- Platform services must be enabled for your tenant account by a StorageGRID administrator.
- You must have already created an S3 bucket whose contents you want to index.
- The endpoint that you intend to use as a destination for the search integration service must already exist, and you must have its URN.
- You must belong to a user group that has the Manage All Buckets or the Root Access permission, which allows you to manage the settings for all S3 buckets in your tenant account. These permissions override the permission settings in group or bucket policies when configuring the bucket using the Tenant Manager.

About this task

After you configure the search integration service for a source bucket, creating an object or updating an object's metadata or tags triggers object metadata to be sent to the destination endpoint. If you enable the search integration service for a bucket that already contains objects, metadata notifications are not automatically sent for existing objects. You must update these existing objects to ensure that their metadata is added to the destination search index.

Steps

1. Use a text editor to create the metadata notification XML required to enable search integration.
 - See the information about configuration XML for search integration.
 - When configuring the XML, use the URN of a search integration endpoint as the destination.

```
<MetadataNotificationConfiguration>
  <Rule>
    <Status>Enabled</Status>
    <Prefix></Prefix>
    <Destination>
      <Urn>arn:aws:es:us-east-
1:11111111111111:domain/mydomain/myindex/mytype</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>
```

2. In the Tenant Manager select **STORAGE (S3) > Buckets**.

3. Select the name of the source bucket.

The bucket details page appears.

4. Select **Platform services > Search integration**
5. Select the **Enable search integration** check box.
6. Paste the metadata notification configuration into the text box, and select **Save changes**.

Bucket options Bucket access Platform services

Replication Disabled

Event notifications Disabled

Search integration Disabled

Enable the search integration service to send object metadata to a destination search index whenever an object is created, deleted, or its metadata or tags are updated.

- Platform services must be enabled for your tenant account by a StorageGRID administrator.
- You must have already configured an endpoint for the search integration service.
- You must specify the URN of that endpoint in the search integration configuration XML for the bucket you want to index.

Enable search integration

Clear

```
<MetadataNotificationConfiguration>
  <Rule>
    <Status>Enabled</Status>
    <Prefix></Prefix>
    <Destination>
      <Urn>arn:aws:es:us-east-1:111111111111:domain/mydomain/myindex/mytype</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>
```

Save changes



Platform services must be enabled for each tenant account by a StorageGRID administrator using the Grid Manager or Management API. Contact your StorageGRID administrator if an error occurs when you save the configuration XML.

7. Verify that the search integration service is configured correctly:

- Add an object to the source bucket that meets the requirements for triggering a metadata notification as specified in the configuration XML.

In the example shown earlier, all objects added to the bucket trigger a metadata notification.

- Confirm that a JSON document that contains the object's metadata and tags was added to the search index specified in the endpoint.

After you finish

As necessary, you can disable search integration for a bucket using either of the following methods:

- Select **STORAGE (S3) > Buckets** and unselect the **Enable search integration** check box.
- If you are using the S3 API directly, use a **DELETE** Bucket metadata notification request. See the instructions for implementing S3 client applications.

Related information

[Understand search integration service](#)

[Configuration XML for search integration](#)

[Use S3](#)

[Create platform services endpoint](#)

JSON generated by search integration service

When you enable the search integration service for a bucket, a JSON document is generated and sent to the destination endpoint each time object metadata or tags are added, updated, or deleted.

This example shows an example of the JSON that could be generated when an object with the key `SGWS/Tagging.txt` is created in a bucket named `test`. The `test` bucket is not versioned, so the `versionId` tag is empty.

```
{  
  "bucket": "test",  
  "key": "SGWS/Tagging.txt",  
  "versionId": "",  
  "accountId": "86928401983529626822",  
  "size": 38,  
  "md5": "3d6c7634a85436eeee06d43415012855",  
  "region": "us-east-1"  
  "metadata": {  
    "age": "25"  
  },  
  "tags": {  
    "color": "yellow"  
  }  
}
```

Object metadata included in metadata notifications

The table lists all the fields that are included in the JSON document that is sent to the destination endpoint when search integration is enabled.

The document name includes the bucket name, object name, and version ID if present.

Type	Item name and description
Bucket and object information	bucket: Name of the bucket
	key: Object key name
	versionID: Object version, for objects in versioned buckets
	region: Bucket region, for example us-east-1
System metadata	size: Object size (in bytes) as visible to an HTTP client
	md5: Object hash
User metadata	metadata: All user metadata for the object, as key-value pairs
	key:value
Tags	tags: All object tags defined for the object, as key-value pairs
	key:value

 For tags and user metadata, StorageGRID passes dates and numbers to Elasticsearch as strings or as S3 event notifications. To configure Elasticsearch to interpret these strings as dates or numbers, follow the Elasticsearch instructions for dynamic field mapping and for mapping date formats. You must enable the dynamic field mappings on the index before you configure the search integration service. After a document is indexed, you cannot edit the document's field types in the index.

Copyright information

Copyright © 2025 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—with 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.