



StorageGRID® 11.2

Tenant User's Guide

August 2019 | 215-13587_2019-08_en-us
doccomments@netapp.com

 **NetApp®**

Contents

Using the Tenant Manager	5
Using a StorageGRID tenant account	5
Web browser requirements	6
Signing in to the Tenant Manager	7
Signing out of the Tenant Manager	9
Understanding the Dashboard	9
Understanding the Tenant Management API	12
Tenant Management API versioning	14
Protecting against Cross-Site Request Forgery (CSRF)	15
Managing system access for tenant users	16
Using identity federation	16
Configuring a federated identity source	16
Forcing synchronization with the identity source	21
Disabling identity federation	21
Managing groups	22
Creating groups for an S3 tenant	22
Creating groups for a Swift tenant	25
Tenant management permissions	27
Cloning a group	28
Editing a group	30
Removing a group	31
Managing local users	31
Creating local users	32
Cloning local users	33
Editing local users	34
Changing a local user's password	35
Removing local users	36
Managing S3 tenant accounts	38
Managing S3 access keys	38
Creating your own S3 access keys	38
Removing your own S3 access keys	40
Creating another user's S3 access keys	41
Removing another user's S3 access keys	42
Managing S3 buckets	43
Managing compliant buckets and objects	43
Creating an S3 bucket	47
Configuring compliance settings	50
Changing the consistency level	53
Configuring Cross-Origin Resource Sharing (CORS)	54
Enabling or disabling last access time updates	56
Managing S3 platform services	59

What platform services are	59
Understanding the CloudMirror replication service	60
Understanding notifications for buckets	61
Understanding the search integration service	62
Considerations for using platform services	63
Configuring endpoints	64
What an endpoint is	65
Specifying the URN for an endpoint	65
Creating an endpoint	67
Editing or removing an endpoint	70
Troubleshooting endpoint errors	71
Configuring CloudMirror replication	72
Configuring event notifications	74
Configuring the search integration service	77
Configuration XML for search integration	78
Configuring the search integration service	81
JSON generated by the search integration service	83
Object metadata included in metadata notifications	83
Copyright	84
Trademark	85
How to send comments about documentation and receive update notifications	86

Using the Tenant Manager

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

You can use the Tenant Manager to monitor your storage usage and manage user accounts 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.

Using a StorageGRID tenant account

Tenant accounts allow you to use the Simple Storage Service (S3) REST API or the Swift REST API to store and retrieve objects on a StorageGRID system.

Each tenant account supports the use of a single protocol, which is specified by the grid administrator when the account is created. To store and retrieve objects to a StorageGRID system with both protocols, you must have two tenant accounts: one for S3 buckets and objects, and one for Swift containers and objects. Each tenant account has its own federated or local groups, users, containers (buckets for S3), 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.

Note: 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 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.

Creating tenant accounts

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 account
- Which client protocol will be used by the tenant account (S3 or Swift)
- For S3 tenant accounts: Whether the tenant account has permission to use platform services. If the use of platform services is permitted, 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 single sign-on (SSO) is not in use for the StorageGRID system, whether the tenant account will use its own identity source or share the grid's identity source, and the initial password for the tenant's local root user.
- If SSO is enabled, which federated group has Root Access permission to configure the tenant account.

Configuring S3 tenants

After an S3 tenant account is created, 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
- Using platform services (if enabled)
- Monitoring storage usage

Attention: 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.

Configuring Swift tenants

After a Swift tenant account is created, users with the Root Access permission can access the Tenant Manager to perform tasks such as the following:

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

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

Related information

[Administering StorageGRID](#)

[Implementing S3 client applications](#)

[Implementing Swift client applications](#)

Web browser requirements

You must use a supported web browser.

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

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

Signing in to the Tenant Manager

You access the sign in page by entering the URL for the tenant into the address bar of a supported web browser.

Before you begin

- You know your login credentials.
- You know the tenant's 20-digit account ID, or you have the URL for the tenant, which is the IP address or fully qualified domain name of a StorageGRID Admin Node, followed by `/?accountId=20-digit-account-id`.
- You are using a supported web browser.
- Cookies are enabled in your web browser.
- You have specific access permissions.

Steps

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

StorageGRID® Tenant Manager

Recent -- Optional --


Account ID 39105156032765926037

Username

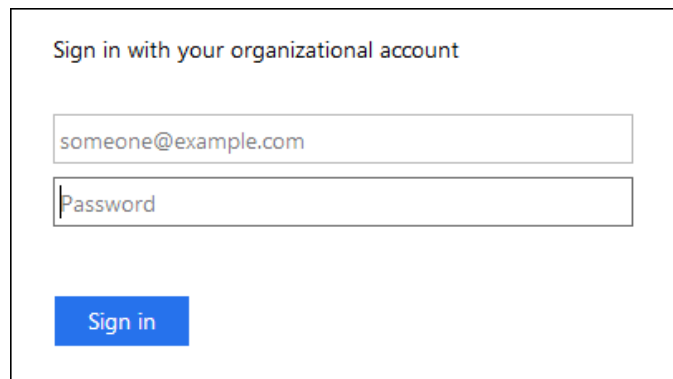
Password

Sign in

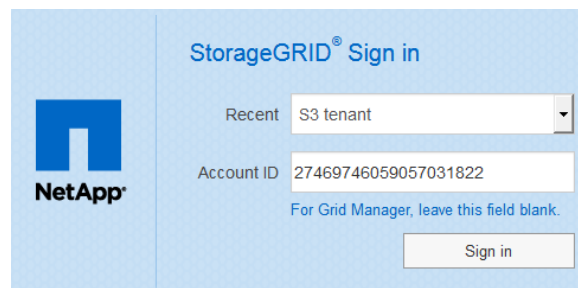
- c. If you received the password from someone at another company, such as a service provider, consider changing the password to ensure it is secure. Select **Access Control > Users**.
- If single sign-on (SSO) is not being used for the StorageGRID system, and you entered the IP address or fully qualified domain name of a StorageGRID Admin Node:
 - a. The Grid Manager sign-in page appears. Click the **Tenant Login** link in the upper right.



- b. On the Tenant Manager sign-in page, select the name of the tenant account if it appears in the list of recent accounts, or enter the tenant's 20-digit account ID.
 - c. Enter your username and password.
 - d. Click **Sign in**.
- If SSO is enabled for the StorageGRID system, and you entered the complete URL for the tenant, your organization's SSO page appears. Enter your standard SSO credentials. For example:

A screenshot of a sign-in form titled 'Sign in with your organizational account'. It contains two input fields: the first for an email address with the placeholder 'someone@example.com', and the second for a password with the placeholder 'Password'. Below the fields is a blue 'Sign in' button.

- If SSO is enabled for the StorageGRID system, and you entered the IP address or fully qualified domain name of a StorageGRID Admin Node without the tenant Account ID:
 - a. Do either of the following:
 - Enter the Account ID, and click **Sign in**.
 - Select the name of the tenant account if it appears in the list of recent accounts, and click **Sign in**.

A screenshot of the StorageGRID Sign in page. On the left is the NetApp logo. The main area is titled 'StorageGRID® Sign in'. It features a 'Recent' dropdown menu with 'S3 tenant' selected, an 'Account ID' input field containing '27469746059057031822', and a note: 'For Grid Manager, leave this field blank.' Below these is a 'Sign in' button.

- b. Sign in with your standard SSO credentials on your organization's SSO sign-in page.

Related tasks

[Changing a local user's password](#) on page 35

Related references


[Web browser requirements](#) on page 6

Signing out of the 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 **Sign Out** link in the top-right corner of the user interface.



2. Click **Sign Out**.

Option	Description
SSO not in use	<p>You are signed out of the Admin Node.</p> <p>The Tenant Manager sign in page is displayed.</p> <p>Note: If you signed into more than one Admin Node, you must sign out of each node.</p>
SSO enabled	<p>You are signed out of all Admin Nodes you were accessing.</p> <p>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.</p> <p>Note: 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.</p>

Understanding the Dashboard

When you first sign in to the Tenant Manager, the Dashboard shows how much storage the tenant account is using. If you have configured one or more endpoints for use with platform services, the dashboard also indicates if there are any recent endpoint errors.

Storage Usage

The Storage Usage panel shows which containers (buckets for S3) are consuming the most storage. Up to eight containers can be shown. The Other segment combines all other containers, including any containers that consume less than 1% of the total storage.

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

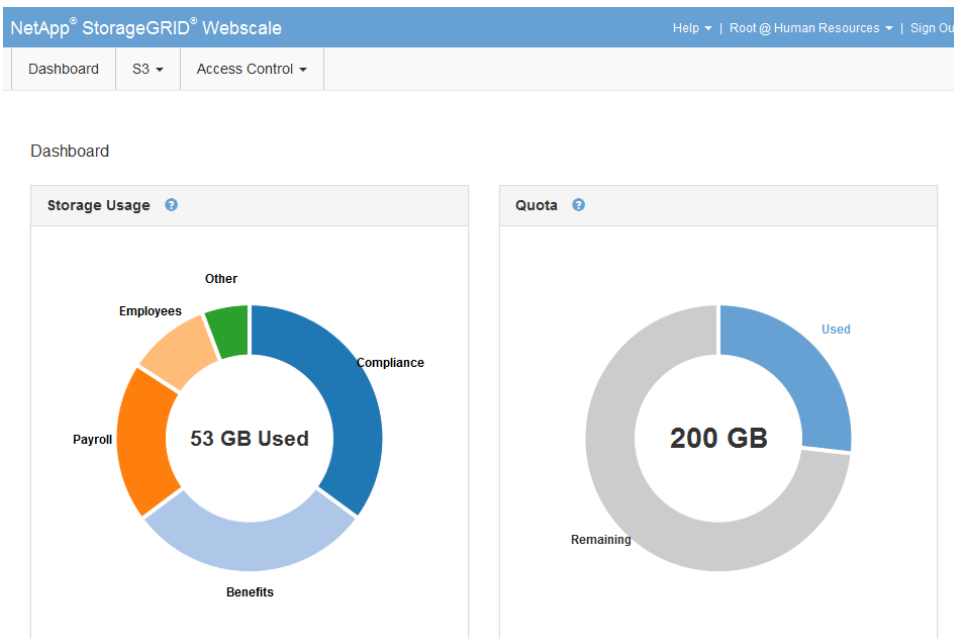
Quota

If the maximum number of gigabytes, terabytes, or petabytes available for the tenant was specified when the account was created, the Quota panel shows how much of that quota has been used and how much is still available.


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

If the quota is exceeded, the tenant account cannot create new objects.

If no quota was set, the tenant has an unlimited quota, and an informational message is displayed.



You can place your cursor over any of the chart segments to obtain more information, including the number of stored objects and total bytes for each container or bucket.



Endpoint error

If you have configured one or more endpoints for use with platform services, the Dashboard displays a message if any endpoint errors have occurred within the past 7 days.

Dashboard

One or more endpoints have recently encountered an error and might not be functioning properly. Go to the [Endpoints](#) page to view the error details. The last error occurred 5 days ago.

To see details about this or any other endpoint errors, click **Endpoints** to display the Endpoints page.

Related concepts

[Troubleshooting endpoint errors](#) on page 71


Understanding the 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.

To access the Swagger documentation for the Tenant Management API:

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



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*").
 See “Protecting against Cross-Site Request Forgery” for information on improving authentication security.
Note: If single sign-on (SSO) is enabled for the StorageGRID system, you must perform different steps to authenticate. See “Authenticating in to the API if single sign-on is enabled” in the instructions for administering StorageGRID.
- **compliance** – Operations to determine how global compliance is configured for the StorageGRID system.
- **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. For S3, you can create compliant and non-compliant buckets; modify compliance settings; 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; and manage the configuration settings for platform services, including CloudMirror replication, notifications, and search integration (metadata-notification). For Swift, you can 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.
- **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 <small>(query)</small>	filter by group type
limit integer <small>(query)</small>	maximum number of results
marker string <small>(query)</small>	marker-style pagination offset (value is Group's URN)
includeMarker boolean <small>(query)</small>	if set, the marker element is also returned
order string <small>(query)</small>	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": "2.2" }</pre>

Issuing API requests

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

Steps

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

Related concepts

[Protecting against Cross-Site Request Forgery \(CSRF\)](#) on page 15

Related information

[Administering StorageGRID](#)

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
```

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

When 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

Determining which API versions are supported in the current release

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

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

```

    2,
    3
  ]
}
```

Specifying an API version for a request

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

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

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

Protecting against Cross-Site Request Forgery (CSRF)

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

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

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

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

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

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

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

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

See the online API documentation for additional examples and details.

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

Managing system access for tenant users

You grant users access to a tenant account by importing groups from a federated identity source and assigning management permissions. You can also create local tenant groups and users, unless single sign-on (SSO) is in effect for the entire StorageGRID system.

Using 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.

Steps

1. [Configuring a federated identity source](#) on page 16
2. [Forcing synchronization with the identity source](#) on page 21
3. [Disabling identity federation](#) on page 21

Configuring a federated identity source

You can configure identity federation if you want tenant groups and users to be managed in another system such as Active Directory, OpenLDAP, or Oracle Directory Server.

Before you begin

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

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.

Identity Federation

Configure Active Directory or OpenLDAP as a federated identity source, so you can grant management permissions to federated groups. See the documentation to configure other identity sources such as Oracle Directory Server.

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.

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

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

Steps

1. Select **Access Control > Identity Federation**.
2. Select **Enable Identity Federation**.

The fields for configuring the LDAP server appear.

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

You can select **Active Directory**, **OpenLDAP**, or **Other**.

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

4. If you selected **Other**, complete the fields in the **LDAP Attributes** section.
 - **Unique User Name:** The name of the attribute that contains the unique identifier of an LDAP user. This attribute is equivalent to `sAMAccountName` for Active Directory and `uid` for OpenLDAP. If you are configuring Oracle Directory Server, enter `uid`.
 - **User UUID:** The name of the attribute that contains the permanent unique identifier of an LDAP user. This attribute is equivalent to `objectGUID` for Active Directory and `entryUUID` for OpenLDAP. If you are configuring Oracle Directory Server, enter `nsuniqueid`. Each user’s value for the specified attribute must be a 32-digit hexadecimal number in either 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. Enter the required LDAP server and network connection information in the **LDAP Server** section:
 - **Hostname:** The server host name or IP address of the LDAP server.
 - **Port:** The port used to connect to the LDAP server. Enter 389.
 - **Username:** The full path of the distinguished name (DN) for the user that will connect to the LDAP server.

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

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

- `sAMAccountName` or `uid`
 - `objectGUID`, `entryUUID`, or `nsuniqueid`
 - `cn`
 - `memberOf` or `isMemberOf`
- **Password:** The password associated with the username.
 - **Group Base DN:** The full path of the distinguished name (DN) for an LDAP subtree you want to search for groups. In the Active Directory example (below), all groups whose Distinguished Name is relative to the base DN (`DC=storagegrid,DC=example,DC=com`) can be used as federated groups.

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

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


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

6. Select a security setting from the **Transport Layer Security (TLS)** drop-down list to specify if TLS is used to secure communications with the LDAP server:

- **Use operating system CA certificate:** Use the default CA certificate installed on the operating system to secure connections.
- **Use custom CA certificate:** Use a custom security certificate.
If you select this setting, copy and paste the custom security certificate in the CA Certificate text box.
- **Do not use TLS:** The network traffic between the StorageGRID system and the LDAP server will not be secured.

7. Optionally, click **Test Connection** to validate your connection settings for the LDAP server.

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



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

Example

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

Enable Identity Federation

LDAP Service Type Active Directory

LDAP Server

Hostname

Port

Username

Password

Group Base DN

User Base DN

Transport Layer Security (TLS) Use custom CA certificate

CA Certificate

```
-----BEGIN CERTIFICATE-----
MIIFmzCCA4OgAwIBAgIJAM5MuRrbdKo/MA0GCSqGSIb3
DQEBDQUAMGMxCzAJBgNV
BAYTAiVTMRcwFQYDVQQIDA5Ob3J0aCBDYXJvYmEYDTE
MMAoGA1UEBwwDUiRQM8w
DQYDVQQKDAZOZXRlcHAxHDAaBgNVBAsME1N0b3Jh
```

Example

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

Enable Identity Federation

LDAP Service Type

LDAP Attributes

User Unique Name

User UUID

Group Unique Name

Group UUID

LDAP Server

Hostname

Port

Username

Password

Group Base DN

User Base DN

Transport Layer Security (TLS)

CA Certificate

Related concepts

[Tenant management permissions](#) on page 27

[Guidelines for configuring an OpenLDAP server](#) on page 20

Guidelines for configuring an OpenLDAP server

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

Memberof and refint overlays

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

Indexing

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

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

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

See the information about reverse group membership maintenance in the Administrator's Guide for OpenLDAP.

Related information

[OpenLDAP documentation: Version 2.4 Administrator's Guide](#)

Forcing synchronization with the identity source

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

Before you begin

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

Steps

1. Select **Access Control > Identity Federation**.

The Identity Federation page appears. The **Synchronize** button is at the bottom of the page.

2. Click **Synchronize**.

A confirmation message is displayed indicating that synchronization started successfully.

Related concepts

[Tenant management permissions](#) on page 27

Disabling identity federation

If you configured an identity federation service for this tenant, you can temporarily or permanently disable identity federation for tenant groups and users. When identity federation is disabled, there is no communication between the StorageGRID system and the identity source. However, any settings you have configured are retained, allowing you to easily re-enable identity federation in the future.

Before you begin

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

About this task

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

- Federated users will be unable to sign in.
- Federated users who are currently signed in will retain access to the tenant account 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.

Steps

1. Select **Access Control > Identity Federation**.
2. Deselect the **Enable Identity Federation** checkbox.
3. Click **Save**.

Related concepts

[Tenant management permissions](#) on page 27

Managing groups

You assign permissions to user groups to control which tasks tenant users can perform. You can import federated groups from an identity source, such as Active Directory or OpenLDAP, or you can create local groups.

Attention: 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 access S3 and Swift resources, based on group permissions.

Choices

- [Creating groups for an S3 tenant](#) on page 22
- [Creating groups for a Swift tenant](#) on page 25
- [Tenant management permissions](#) on page 27
- [Cloning a group](#) on page 28
- [Editing a group](#) on page 30
- [Removing a group](#) on page 31

Creating groups for an S3 tenant

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

Before you begin

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

Steps

1. Select **Access Control > Groups**.

Groups

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

<input type="button" value="+ Add"/> <input type="button" value="Clone"/> <input type="button" value="Edit"/> <input type="button" value="Remove"/>			
	Name	ID	Federated
<input type="radio"/>	Applications	5832bfe-9337-4877-9b87-4b20b8018ee1	
<input checked="" type="radio"/>	Managers	fcecbaac-1994-4f9e-a875-200a9d64f89c	

Group Type: Show rows per page

2. Click Add.

The Add Group page appears.

Add Group

Create a new local group or import a group from the external identity source.

Type Local Federated

Display Name

Unique Name

Management Permissions

Root Access Manage Your Own S3 Credentials

Manage All Containers Manage Endpoints

S3 Policy ?

Add an S3 group policy to control user access permissions for specific S3 resources, including buckets. Non-root users have no access by default. See the S3 Implementation Guide for details and examples.

Changes to a group policy might not take effect for up to 15 minutes due to caching.

Group Policy

This group has no access to S3 resources unless access is granted by a bucket policy. To allow group access, select a predefined policy or enter a custom policy.

3. For the group's type, select **Local to create a local group, or select **Federated** to import a group from the previously configured identity source.**

Attention: 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.

If you selected...	Enter...
Local	Both a display name and a unique name for this group. You can edit the display name later.
Federated	The unique name of the federated group. Note: For Active Directory, the unique name is the name associated with the <code>sAMAccountName</code> attribute. For OpenLDAP, the unique name is the name associated with the <code>uid</code> attribute.

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

See “Tenant management permissions.”

6. From the **Group Policy** drop-down, select how you want to create the group policy that defines which S3 access permissions members of this group will have.

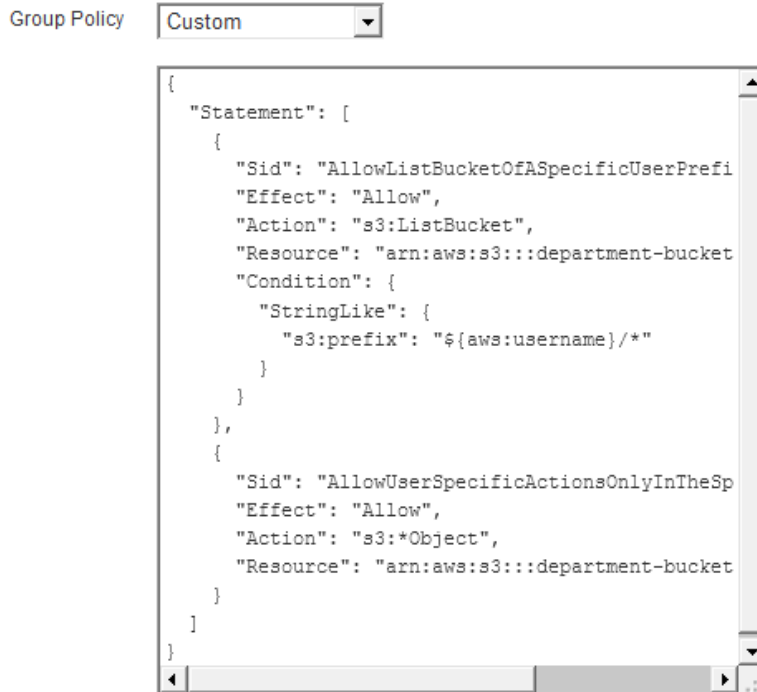
Option	Description
No S3 Access	Default. Users in this group do not have access to S3 resources, unless access is granted with a bucket policy. If you select this option, only the root user will have access to S3 resources by default.
Read Only Access	Users in this group have read-only access to S3 resources.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

7. If you selected **Custom**, enter the group policy.

Note: Each group policy has a size limit of 5,120 bytes. You must enter a valid JSON formatted string.

Example

In this example, members of the group are only permitted to list and access their specific folder (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.



8. Click **Save**.

New group policies might take up to 15 minutes to take effect because of caching.

Related concepts

[Tenant management permissions](#) on page 27

Related information

[Implementing S3 client applications](#)

Creating 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 Administrator permission, which is required to manage the containers and objects for a Swift tenant account.

Before you begin


- You must be signed in to the Tenant Manager using a supported browser.
- You must be the root user or have 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 Control > Groups**.

Groups

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



2. Click **Add**.
3. For the group's type, select **Local** to create a local group, or select **Federated** to import a group from the previously configured identity source.

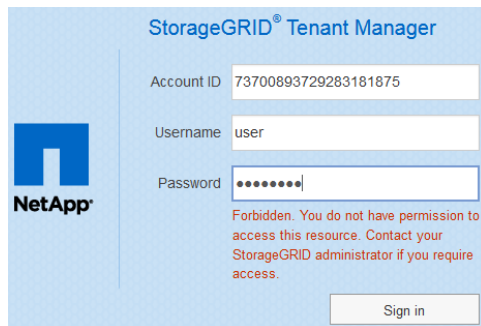
Attention: If single sign-on (SSO) is enabled for your StorageGRID system, users belonging to local groups will not be able to sign in to the Tenant Manager, although they can use client applications to manage the tenant's resources, based on group permissions.

4. Enter the group's name.

If you selected...	Enter...
Local	Both a display name and a unique name for this group. You can edit the display name later.
Federated	The unique name of the federated group. Note: For Active Directory, the unique name is the name associated with the <code>sAMAccountName</code> attribute. For OpenLDAP, the unique name is the name associated with the <code>uid</code> attribute.

5. In the **Management Permissions** section, select **Root Access** if you want users in this group to be able to sign in to Tenant Manager or to the Tenant Management API.

Attention: Users belonging to groups that do not have the Root Access permission receive an error if they try to sign in to the tenant account.



6. In the **Swift Permissions** section, select **Administrator** if you want users in this group to be able to use the Swift REST API to create and manage Swift containers and objects.

Add Group

Create a new local group or import a group from the external identity source.

Type Local Federated

Display Name

Unique Name

Management Permissions

Root Access

Swift Permissions

Administrator

Attention: Users must have the Administrator permission to perform operations with the Swift REST API. The Root Access permission does not allow Swift users to use the Swift REST API.

7. Click **Save**.

New group policies might take up to 15 minutes to take effect because of caching.

Related concepts

[Tenant management permissions](#) on page 27

Related information

[Implementing Swift client applications](#)

Tenant management permissions

Tenant management permissions are assigned to groups and determine which tasks users can perform using the Tenant Manager or the Tenant Management API. A user can belong to one or more 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)

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 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 S3 > My Credentials menu option.
Manage All Containers	<ul style="list-style-type: none"> • S3 tenants: Allows users to use the Tenant Manager and the Tenant Management API to manage the settings for all S3 buckets in the tenant account, regardless of S3 bucket or group policies. Users who do not have this permission do not see the S3 > Buckets menu option. • 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 Containers permission to Swift groups from the Tenant Management API. You cannot assign this permission to Swift groups using the Tenant Manager.</p>
Manage Endpoints	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. Users who do not have this permission do not see the S3 > Endpoints menu option.

Related information

[Implementing S3 client applications](#)

[Implementing Swift client applications](#)

Cloning a group

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

Before you begin

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

Steps

1. Select **Access Control > Groups**.

Groups

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

Name	ID	Federated
Applications	5832bfbe-9337-4877-9b87-4b20b8018ee1	
Managers	fcecbaac-1994-4f9e-a875-200a9d64f89c	

Group Type: All

Show 20 rows per page

2. Select the group you want to clone.

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

3. Click **Clone**.

4. For the group's type, select **Local** to create a local group, or select **Federated** to import a group from the previously configured identity source.

Attention: 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.

If you selected...	Enter...
Local	Both a display name and a unique name for this group. You can edit the display name later.
Federated	The unique name of the federated group. Note: For Active Directory, the unique name is the name associated with the <code>sAMAccountName</code> attribute. For OpenLDAP, the unique name is the name associated with the <code>uid</code> attribute.

6. Assign permissions to this group.

7. If you are cloning a group for an S3 tenant, optionally select a different option from the **Group Policy** drop-down.

8. If you selected a custom policy, update the JSON string as required.

9. Click **Save**.

New group policies might take up to 15 minutes to take effect because of caching.

Related concepts

[Tenant management permissions](#) on page 27

Editing a group

You can edit a group to change the display name of a local group or to update permissions.

Before you begin

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

Steps

1. Select **Access Control > Groups**.

Groups

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

<input type="button" value="+ Add"/> <input type="button" value="Clone"/> <input type="button" value="Edit"/> <input type="button" value="Remove"/>			
	Name	ID	Federated
<input type="radio"/>	Applications	5832bfbе-9337-4877-9b87-4b20b8018ee1	
<input checked="" type="radio"/>	Managers	fcecbaac-1994-4f9e-a875-200a9d64f89c	

Group Type Show rows per page

2. Select the group you want to edit.

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

3. Click **Edit**.

4. 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. Update the permissions as needed.

6. If you are editing a group for an S3 tenant, optionally select a different option from the **Group Policy** drop-down.

7. If you selected a custom policy, update the JSON string as required.

8. Click **Save**.

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

Related concepts

[Tenant management permissions](#) on page 27

Removing a group

You can remove a group. 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.

Before you begin

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

Steps

1. Select **Access Control > Groups**.

Groups

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

<input type="button" value="+ Add"/> <input type="button" value="Clone"/> <input type="button" value="Edit"/> <input type="button" value="Remove"/>			
	Name	ID	Federated
<input type="radio"/>	Applications	5832bfbе-9337-4877-9b87-4b20b8018ee1	
<input checked="" type="radio"/>	Managers	fcecbaac-1994-4f9e-a875-200a9d64f89c	

Group Type Show rows per page

2. Select the group you want to remove.

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

3. Click **Remove**.

A confirmation dialog appears.

4. Click **OK** to confirm you want to remove the group.

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

Related concepts

[Tenant management permissions](#) on page 27

Managing local users

You can create local users and assign them to local admin groups to determine which Tenant Manager 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.

Attention: 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.

Choices

- [Creating local users](#) on page 32
- [Cloning local users](#) on page 33

- [Editing local users](#) on page 34
- [Changing a local user's password](#) on page 35
- [Removing local users](#) on page 36

Creating local users

You can create local users and assign them to one or more local groups to control their access permissions. Because local users must be assigned to local groups, you should create the groups before creating the users.

Before you begin

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

Steps

1. Select **Access Control > Users**.

Users

View local and federated users. Edit properties and group membership of local users.

<input type="button" value="+ Create"/> <input type="button" value="Clone"/> <input type="button" value="Edit"/> <input type="button" value="Edit S3 Keys"/> <input type="button" value="Change Password"/> <input type="button" value="Remove"/>				
	Username	Full Name	Denied	Federated
<input type="radio"/>	root	Root		
<input type="radio"/>	User_01	User_01		
<input checked="" type="radio"/>	User_02	User_02		

User Type: Show rows per page

2. Click **Create**.
3. 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.
 - **Unique name:** A unique username, which is used when the user signs in.
 - **Deny access:** If selected, this user cannot sign in to the tenant account, even though the user might still belong to one or more groups.

Note: You can use this check box to temporarily suspend a user's ability to sign in.
 - **Password:** A password, which is used when the user signs in.

Create User

Create a local user.

Full Name

Unique Name

Deny Access

Password

Password

Confirm Password

Group Membership

Group Name
<input type="checkbox"/> Managers

4. In the **Group Membership** section, select one or more local groups.
Permissions are cumulative. Users will have all permissions for all groups they belong to.
5. Click **Save**.

Related concepts

[Tenant management permissions](#) on page 27

Cloning local users

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

Before you begin

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

Steps

1. Select **Access Control > Users**.

Users

View local and federated users. Edit properties and group membership of local users.

<input type="button" value="+ Create"/> <input type="button" value="Clone"/> <input type="button" value="Edit"/> <input type="button" value="Edit S3 Keys"/> <input type="button" value="Change Password"/> <input type="button" value="Remove"/>				
	Username	Full Name	Denied	Federated
<input type="radio"/>	root	Root		
<input type="radio"/>	User_01	User_01		
<input checked="" type="radio"/>	User_02	User_02		

User Type: Show rows per page

2. Select the user you want to clone.

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

3. Click **Clone**.

4. 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.
- **Unique name:** A unique username, which is used when the user signs in.
- **Deny access:** If selected, this user cannot sign in to the tenant account, even though the user might still belong to one or more groups.

Note: You can use this check box to temporarily suspend a user's ability to sign in.

- **Password:** A password, which is used when the user signs in.

5. In the **Group Membership** section, select one or more local groups.

Permissions are cumulative. Users will have all permissions for all groups they belong to.

6. Click **Save**

Related concepts

[Tenant management permissions](#) on page 27

Editing local users

You can edit local users to change names, prevent them from being able to access the tenant, or assign them to different groups.

Before you begin

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

Steps

1. Select **Access Control > Users**.

Users

View local and federated users. Edit properties and group membership of local users.

<input type="button" value="+ Create"/> <input type="button" value="Clone"/> <input type="button" value="Edit"/> <input type="button" value="Edit S3 Keys"/> <input type="button" value="Change Password"/> <input type="button" value="Remove"/>				
	Username	Full Name	Denied	Federated
<input type="radio"/>	root	Root		
<input type="radio"/>	User_01	User_01		
<input checked="" type="radio"/>	User_02	User_02		

User Type: Show rows per page

2. Select the user you want to edit.

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

3. Click **Edit**.

4. Update the following fields as required:
 - **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.
 - **Deny access:** If selected, this user cannot sign in to the tenant, even though the user might still belong to one or more groups.

Note: You can use this check box to temporarily suspend a user's ability to sign in.

5. In the **Group Membership** section, select one or more local groups.

Permissions are cumulative. Users will have all permissions for all groups they belong to.

6. Click **Save**

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

Related concepts

[Tenant management permissions](#) on page 27

Changing a local user's password

A tenant administrator can change passwords for local tenant users.

Before you begin

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

Steps

1. Select **Access Control > Users**.

Users

View local and federated users. Edit properties and group membership of local users.

The screenshot shows a web interface for managing users. At the top, there are several action buttons: '+ Create', 'Clone', 'Edit', 'Edit S3 Keys', 'Change Password', and 'Remove'. Below these is a table with the following columns: Username, Full Name, Denied, and Federated. The table contains three rows: 'root' (Full Name: Root), 'User_01' (Full Name: User_01), and 'User_02' (Full Name: User_02). The 'User_02' row is highlighted in blue, indicating it is selected. Below the table, there is a 'User Type' dropdown menu set to 'All', a 'Show 20 rows per page' indicator, and navigation arrows.

2. Select the user, and click **Change Password**.
3. Enter the new password, and click **Save**.

The screenshot shows a dialog box titled 'Change Password - user2'. It contains two input fields: 'New Password' and 'Confirm New Password', both filled with dots to represent masked text. At the bottom right of the dialog, there are two buttons: 'Cancel' and 'Save'.

Related concepts

[Tenant management permissions](#) on page 27

Removing local users

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

Before you begin

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

Steps

1. Select **Access Control > Users**.

Users

View local and federated users. Edit properties and group membership of local users.

This screenshot is identical to the one above, showing the 'Users' management interface with 'User_02' selected in the table.

2. Select the user you want to remove.

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

3. Click **Remove**.

A confirmation dialog appears.

4. Click **OK** to confirm you want to remove the user.

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

Related concepts

[Tenant management permissions](#) on page 27

Managing S3 tenant accounts

You can use the Tenant Manager to manage S3 access keys and to create and manage S3 buckets.

Choices

- [Managing S3 access keys](#) on page 38
- [Managing S3 buckets](#) on page 43

Managing S3 access keys

Each user of an S3 tenant account must have an access key to store and retrieve objects on 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 also provide full access to the tenant's buckets and objects 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.

Choices

- [Creating your own S3 access keys](#) on page 38
- [Removing your own S3 access keys](#) on page 40
- [Creating another user's S3 access keys](#) on page 41
- [Removing another user's S3 access keys](#) on page 42

Creating 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.

Before you begin

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

About this task

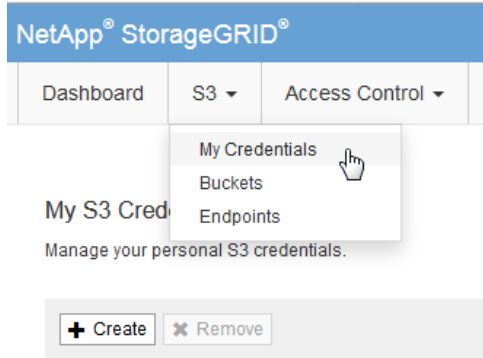
You can create one or more S3 access keys. Multiple access keys allow you to begin using a new key without temporarily losing access to the objects in the account. You simply create the new access key, update the application with your new access key ID and secret key, and then remove the old access key from StorageGRID.

Attention: The S3 account can be accessed using the Access Key ID and Secret Key for any currently displayed key. For this reason, protect your access keys as you would a password. Rotate

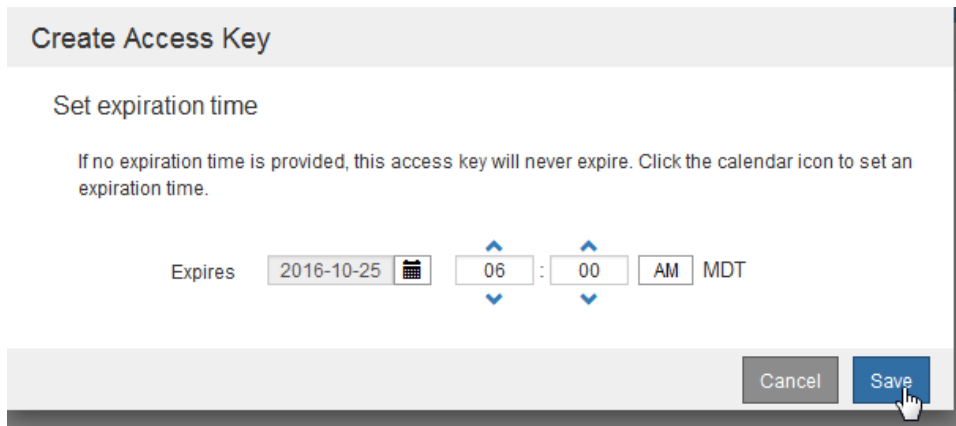
access keys on a regular basis, remove any unused keys from your account, and never share them with other users.

Steps

1. Click **S3 > My Credentials**.



2. Click **Create**.
3. Use the calendar control to select the expiration date and then set the time, or leave the default value of Never, and click **Save**.



You can set an expiration date and time to limit your access to a certain time period or to cause old keys to be removed automatically. Setting a short expiration time can help reduce your risk if your access key ID and secret key are accidentally exposed.

The Save Keys dialog box is displayed, listing your Access Key ID and Secret Access Key.

4. Copy the Access Key ID and the Secret Access Key to a safe location, or click **Download** to save a spreadsheet file (.csv) containing the Access Key ID and Secret Access Key.

Save Keys

You will not be able to view the Access Key ID and Secret Access Key after you close this dialog. To save the keys for future reference, click the Download button or copy and paste the values to another location.

Access Key ID	<input style="width: 90%;" type="text" value="9PELXW0KAZVP1QCNMWGC"/>
Secret Access Key	<input style="width: 90%;" type="text" value="F30BjSpVKSI9pt6FxGwKGJ1Q47AbxrTVh21qcuQY"/>

Download
Finish

Attention: Do not close this dialog until you have copied or downloaded this information.

5. Click **Finish**.

Related concepts

[Tenant management permissions](#) on page 27

Removing your own S3 access keys

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

Before you begin

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

About this task

You should remove any access keys from your StorageGRID user account that you are no longer using.

Steps

1. Click **S3 > My Credentials**.
2. Select the entry you want to remove.
3. Click **Remove**.
4. Click **OK**.

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

Related concepts

[Tenant management permissions](#) on page 27

Creating 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.

Before you begin

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

About this task

The S3 account can be accessed using the Access Key ID and Secret Access Key for any currently displayed key. 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. Click **Access Control > Users**.
2. Select the user whose S3 access keys you want to manage, and click **Edit S3 Keys**.
The Managing S3 Access Key dialog appears, showing any S3 access keys previously defined for the user.
3. Click **Create**.
4. Use the calendar control to select the expiration date and then set the time, or leave the default value of Never, and click **Save**.

Create Access Key

Set expiration time

If no expiration time is provided, this access key will never expire. Click the calendar icon to set an expiration time.

Expires 2016-10-25 06 : 00 AM MDT

Cancel Save

You can set an expiration date and time to limit the user's access to a certain time period or to cause old access keys to be removed automatically. Setting a short expiration time can help reduce the risk if the Access Key ID and Secret Access Key are accidentally exposed.

The Save Keys dialog box is displayed, listing the Access Key ID and Secret Access Key.

5. Copy the Access Key ID and the Secret Access Key to a safe location, or click **Download** to save a spreadsheet file (.csv) containing the Access Key ID and Secret Access Key.

Save Keys

You will not be able to view the Access Key ID and Secret Access Key after you close this dialog. To save the keys for future reference, click the Download button or copy and paste the values to another location.

Access Key ID

Secret Access Key

Download
Finish

Attention: Do not close this dialog until you have copied or downloaded this information.

6. Click **Finish**.

Related concepts

[Tenant management permissions](#) on page 27

Removing another user's S3 access keys

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

Before you begin

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

Steps

1. Click **Access Control > Users**.
2. Select the user whose S3 access keys you want to manage, and click **Edit S3 Keys**.
The Managing S3 Access Key dialog appears, showing any S3 access keys previously defined for the user.
3. Select the entry you want to remove.
4. Click **Remove**.
5. Click **OK**.
Changes might take up to 15 minutes to take effect because of caching.

Related concepts

[Tenant management permissions](#) on page 27

Managing S3 buckets

If you are using an S3 tenant and you have the appropriate permissions, you can create S3 buckets, specify compliance settings, update consistency level settings, configure cross-origin resource sharing (CORS), and enable and disable last access time update settings for the S3 buckets that belong to this tenant account.

Choices

- [Managing compliant buckets and objects](#) on page 43
- [Creating an S3 bucket](#) on page 47
- [Configuring compliance settings](#) on page 50
- [Changing the consistency level](#) on page 53
- [Configuring Cross-Origin Resource Sharing \(CORS\)](#) on page 54
- [Enabling or disabling last access time updates](#) on page 56

Managing compliant buckets and objects

If you are using an S3 tenant account and you need to comply with regulatory requirements when saving object data, you can enable the compliance setting when creating an S3 bucket. If you plan to create compliant buckets, review the steps to create and manage compliant buckets and the considerations for compliant buckets and objects.

How StorageGRID protects compliant data

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

StorageGRID meets the relevant storage requirements of these regulations:

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

Compliance and retention

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

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

Compliance and the storage of duplicate data

StorageGRID ensures that duplicate copies of each compliant object are stored on the grid for the entire retention period. When the global Compliance setting is enabled, grid administrators must use

a compliant rule as the default rule in the ILM policy. At least two copies of each object must exist from the time the object is ingested until the object is deleted.

Compliance and security features

StorageGRID protects compliant objects with the following platform security features:

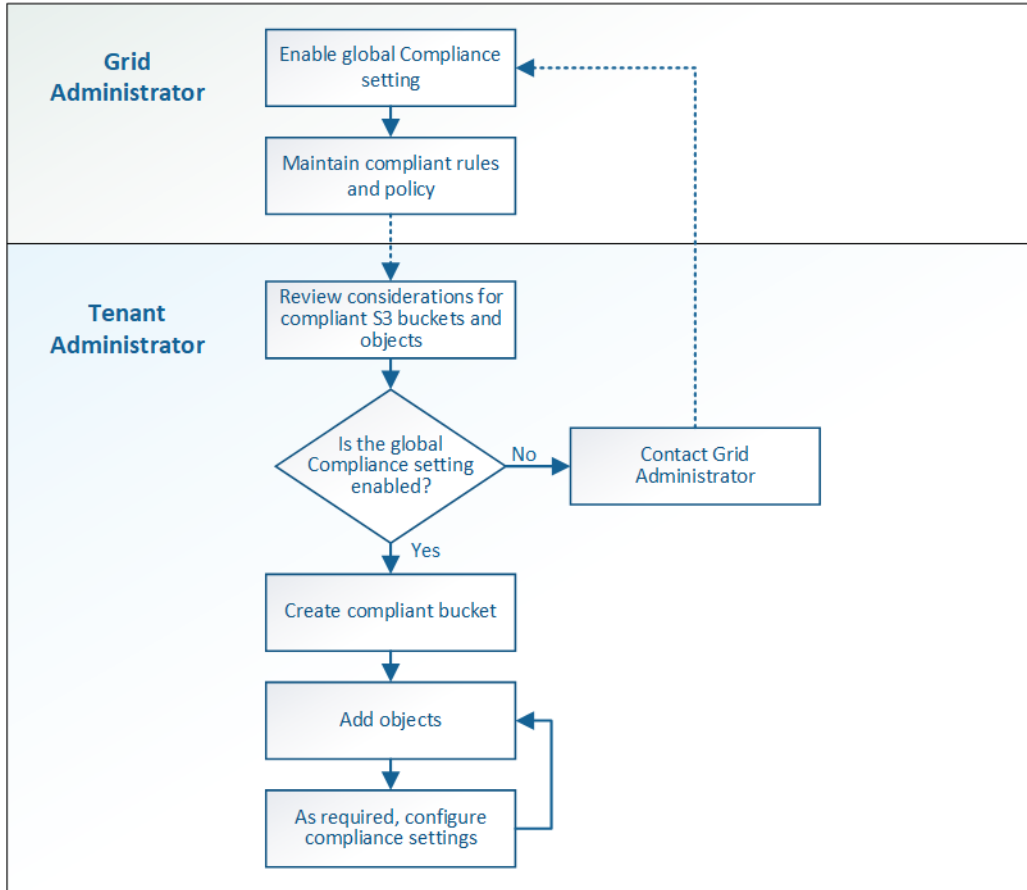
- Internal public key infrastructure and node certificates are used to authenticate and encrypt internode communication. Internode communication is secured by TLS.
- Rules for firewalls and iptables are automatically configured to control incoming and outgoing network traffic, as well as closing unused ports.
- The base operating system of StorageGRID appliances and virtual nodes is hardened; unrelated software packages are removed.
- Root login over SSH is disabled on all grid nodes. SSH access between nodes uses certificate authentication.
- Separate networks are available for Client, Admin, and internal Grid traffic

Compliance workflow

The workflow diagram shows the high-level steps for creating and managing S3 buckets that are compliant. You must coordinate closely with the grid administrator if you plan to create compliant buckets in order to ensure that the objects in those buckets are protected in a manner that meets regulatory requirements.

Before any compliant buckets can be created, the grid administrator must enable the global Compliance setting for the entire StorageGRID system. The grid administrator is also responsible for maintaining the information lifecycle management (ILM) rules that will be used to protect the object data in your compliant buckets. For details, see information about managing S3 buckets and objects in the instructions for administering StorageGRID.

Once the global Compliance setting has been enabled, you can create compliant buckets using the Tenant Manager, the Tenant Management API, or the S3 REST API.



Related concepts

Understanding the Tenant Management API on page 12

Related information

Administering StorageGRID
Implementing S3 client applications

Considerations for compliant buckets and objects

Make sure you understand the restrictions StorageGRID places on compliant buckets and how compliance affects the life of an object.

Restrictions for using compliant buckets

- The global Compliance setting must be enabled before you can create a compliant bucket. If the global setting is disabled, you will not see the Compliance fields in the Tenant Manager, and errors will occur if you try to create a compliant bucket with the Tenant Management API or the S3 REST API.
- If you need to create compliant buckets, you must enable compliance and specify compliance settings when you create the bucket. After a bucket has been saved, compliance cannot be enabled or disabled for the bucket.
- When you specify the retention period for the bucket, you are specifying the minimum amount of time each object in that bucket must be retained (stored) within StorageGRID.

- You can edit bucket settings to increase the retention period, but you can never decrease this value.
- If your organization is notified of a pending legal action or regulatory investigation, you can preserve relevant information by placing a legal hold on the bucket. When a bucket is under a legal hold, no object in that bucket can be deleted even if its retention period has ended. As soon as the legal hold is lifted, objects in the bucket can be deleted when their retention periods end.
- You can add new objects to a compliant bucket at any time, regardless of the bucket's compliance settings.
- You can retrieve objects from a compliant bucket at any time, regardless of the bucket's compliance settings.
- Object versioning is not supported for compliant buckets.

Restrictions for objects in compliant buckets

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

1. Object ingest

- When an object is ingested, the system generates metadata for the object that includes a unique object identifier (UUID) and the ingest date and time. The object inherits the compliance settings from the bucket.
- After an object is ingested into a compliant bucket, its data, S3 user-defined metadata, or S3 object tags cannot be modified, even after the retention period expires.
- StorageGRID maintains three copies of all object metadata at each site to provide redundancy and protect object metadata from loss. Metadata is stored independently of object data.

2. Retention period

- The retention period for an object starts when the object is ingested into the bucket.
- Each time the object is accessed or looked up, the compliance settings for the bucket are also looked up. The system uses the object's ingest time and date and the bucket's retention period setting to calculate when the object's retention period will expire.
- During an object's retention period, multiple copies of the object are stored by StorageGRID. The exact number and type of copies and the storage locations are determined by rules in the active ILM policy.

Note: Contact your StorageGRID administrator to understand how objects will be managed or to request a new ILM rule be added to manage the objects in a particular bucket.

- During an object's retention period, or when legal hold is enabled for the bucket, you cannot delete the object.

3. Object deletion

- When an object's retention period ends, all copies of the object can be deleted, unless legal hold is enabled for the bucket.
- When an object's retention period ends, a bucket-level compliance setting allows you to control how objects are deleted: by users when required or automatically by the system.
- If the bucket setting is to delete objects automatically, all copies of the object are removed by the scanning ILM process in StorageGRID. When an object's retention period ends, the object is scheduled for deletion. The actual amount of time needed to delete all object copies can vary, depending on the number of objects in the grid and how busy the grid processes are.

Creating an 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. Optionally, you can also specify compliance settings for the objects in the bucket.

Before you begin

- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a user group that has the **Manage All Containers** or the **Root Access** permission. These permissions override the permissions settings in group or bucket policies.
- If you need to specify a non-default region for the bucket, the region must have been defined for the StorageGRID system.
- If you plan to create a compliant bucket, the global Compliance setting must have been enabled for the entire StorageGRID system. Contact your StorageGRID administrator if you need to have Compliance enabled.
- If you plan to create a compliant bucket, you must have reviewed the workflow for managing compliant buckets, and you must understand the restrictions for compliant buckets.

About this task

You can also create S3 buckets using the Tenant Management API or the S3 REST API. If you need to delete a bucket, you must use the S3 API.

Steps

1. Select **S3 > Buckets**.

The Buckets page appears, listing any buckets that have already been created.

2. Select **Create Bucket**.

The Create Bucket dialog box appears.

- If the global Compliance setting is disabled, this dialog box includes **Name** and **Region** fields.

The screenshot shows a 'Create Bucket' dialog box. At the top, there is a title bar with the text 'Create Bucket'. Below the title bar, there is a section titled 'Bucket Details' with a question mark icon to its right. Under 'Bucket Details', there are two input fields. The first is labeled 'Name' and is an empty text box. The second is labeled 'Region' and is a dropdown menu with 'us-east-1' selected. At the bottom right of the dialog, there are two buttons: 'Cancel' and 'Save'.

- If the global Compliance setting is enabled, this dialog box also includes a Compliance section and an **Enable Compliance** check box.

Create Bucket

Bucket Details ⓘ

Name

Region

Compliance ⓘ

Compliance settings apply to all objects in the bucket. You cannot disable compliance after the bucket is saved.

Enable Compliance

Cancel Save

3. Enter a unique name for 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 between 3 and 63 characters.
- Can be a series of one or more labels, with adjacent labels separated by a period. Each label must start and end with a lowercase letter or a number and can only use lowercase letters, numbers, and hyphens.
- Must not look like a text-formatted IP address.
- Should not use periods in virtual hosted-style requests because periods will cause problems with server wildcard certificate verification.

Note: See the Amazon Web Services (AWS) Documentation for more information.

4. From the **Region** drop-down list, select the region for this bucket.

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. Contact your StorageGRID administrator to learn which region you should select.

Attention: You cannot change the region after saving the bucket.

5. If the page does not include a Compliance section or you do not want to create a compliant bucket, click **Save**.

Attention: You cannot enable compliance after saving the bucket.

The new bucket is saved and added to the list on the Buckets page.

6. If the page includes a Compliance section and you want to create a compliant bucket, specify the compliance settings.

- a. Select the **Enable compliance** check box.

Attention: You cannot disable compliance after saving the bucket.

Additional compliance fields appear.

Create Bucket

Bucket Details ?

Name

Region

Compliance ?

Compliance settings apply to all objects in the bucket. You cannot disable compliance after the bucket is saved.

Enable Compliance

Retention Period days

After Retention Period

allow users to delete objects

delete objects automatically

Legal Hold

- b. In the **Retention Period** box, specify the length of the retention period for objects added to this bucket, in days.

The retention period for an object starts when that object is ingested into the grid. After you save the bucket, you can increase the retention period, but you cannot decrease this value.

You can enter any number greater than or equal to 0.001. Fractional numbers of minutes are rounded to the nearest integer. For example, if you enter 0.001 (1.44 minutes), the system sets the retention period to 1 minute.

Note: If you want to test compliance settings before implementing this feature in a production-level grid, you can specify a very short duration for the retention period. However, be aware that it will take longer than an hour to automatically delete objects in compliant buckets, even if the retention period is very short. This is because of time required for ILM scanning.

- c. For **After Retention Period**, select what should happen to each object at the end of its retention period.

Option	Description
allow users to delete objects	Users can delete an object manually as soon as its retention period expires, unless the bucket is under a legal hold.
delete objects automatically	An object will be deleted automatically when its retention period expires, unless the bucket is under a legal hold.

- d. If you need to put all objects in this bucket under a legal hold, select the **Legal Hold** check box.

When **Legal Hold** is selected, objects in this bucket cannot be deleted, even if their retention period has expired.

- e. Click **Save**.

The Confirm Retention Period dialog box appears.

- f. If you are sure you want to save this bucket with its current settings, re-enter the number of days each object must be retained.

⚠ Confirm Retention Period

After you save this bucket, you cannot change the region or disable compliance. You can increase the retention period and change the auto-delete and legal hold settings.

To continue, type the retention period again.

Objects will be retained for days

Cancel
OK

- g. Click **OK**.

The compliant bucket is saved and added to the list on the Buckets page.

Related concepts

[Compliance workflow](#) on page 44

[Understanding the Tenant Management API](#) on page 12

Related information

[Administering StorageGRID](#)

[Implementing S3 client applications](#)

[Amazon Web Services \(AWS\) Documentation: Bucket Restrictions and Limitations](#)

Configuring compliance settings

If you created a compliant S3 bucket, you can change the bucket's compliance settings. You can increase, but not decrease, the retention period; place the bucket under a legal hold; remove a legal hold; or change the auto-delete setting. You cannot disable compliance or edit the bucket's name or region.

Before you begin

- A grid administrator has enabled the global Compliance setting for the entire StorageGRID system. The Buckets page does not include the **Configure Compliance** button if this setting is disabled.
- You want to edit a compliant bucket. You cannot modify compliance settings if the bucket was initially created with compliance disabled.
- You must be signed in to the Tenant Manager using a supported browser.
- You belong to a user group that has the **Manage All Containers** or the **Root Access** permission. These permissions override the permissions settings in group or bucket policies.

Steps

1. Select S3 > Buckets.

The Buckets page appears and shows all existing S3 buckets. The table includes each bucket's name, creation time, and region. The **Compliant** column includes a checkmark for compliant buckets. If a compliant bucket is currently under a legal hold, this column also includes a lock icon.

Buckets
Manage settings of your buckets.

Bucket Name	Creation Time	Region	Compliant
<input type="radio"/> compliant-01	2018-02-27 08:39:29 MST	us-east-1	✓
<input type="radio"/> compliant-02	2018-02-27 08:40:53 MST	us-west-1	✓
<input type="radio"/> compliant-03	2018-02-27 10:13:43 MST	eu-west-1	✓
<input type="radio"/> compliant-04	2018-02-27 12:08:50 MST	us-east-1	✓
<input type="radio"/> non-compliant-01	2018-02-27 08:19:34 MST	us-east-1	
<input type="radio"/> non-compliant-02	2018-02-27 08:41:10 MST	eu-west-1	

Displaying 6 buckets.

2. Select a compliant bucket from the list.

The **Configure Compliance** button is disabled if you select a non-compliant bucket.

3. Click Configure Compliance.

The Edit Bucket Settings dialog box appears.

Edit Bucket Settings - compliant-bucket

Compliance ⓘ

Compliance settings apply to all objects in the bucket. You cannot disable compliance after the bucket is saved.

Retention Period days

After Retention Period

allow users to delete objects

delete objects automatically

Legal Hold

4. Modify one or more of the compliance settings for this bucket, as required.

- a. In the **Retention Period** box, specify the length of the retention period for objects added to this bucket, in days.

The retention period for an object starts when that object is ingested into the grid. You can enter any number greater than or equal to the bucket's current retention period.

- b. For **After Retention Period**, select what should happen to each object at the end of its retention period.

Option	Description
allow users to delete objects	Users can delete an object as soon as its retention period expires, unless the bucket is under a legal hold.
delete objects automatically	An object will be deleted automatically when its retention period expires, unless the bucket is under a legal hold.

- c. To put this bucket under a legal hold, select the **Legal Hold** check box. To remove a legal hold, unselect the check box.

When **Legal Hold** is selected, objects in this bucket cannot be deleted, even if their retention period has expired.

5. Click **Save**.

If you did not increase the retention period, the bucket is saved without an additional confirmation step.

Note: If a Service Unavailable error message appears, go to step 7.

6. If you increased the retention period, re-enter the number of days each object must be retained, and click **OK**.

⚠ Confirm Retention Period

After you increase the retention period, you cannot decrease it.

To continue, type the retention period again.

Objects will be retained for days

The bucket is saved with the updated settings.

7. If a Service Unavailable error message appears when you try to save the bucket:
 - a. Click **OK** to close the message.

! Error

503: Service Unavailable

Unable to update compliance settings because the changes cannot be consistently applied on enough storage services. Contact your grid administrator for assistance.

- b. Click **Cancel** to cancel your changes.
- c. Contact your grid administrator to ensure that the required storage services are made available as soon as possible.

- d. When the services are available again, repeat this procedure to update the compliance settings.

Changing 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.

Before you begin

- You must be signed in to the Tenant Manager using a supported browser.
- Users must belong to a user group that has the **Manage All Containers** or the **Root Access** permission. These permissions override the permissions settings in group or bucket policies.

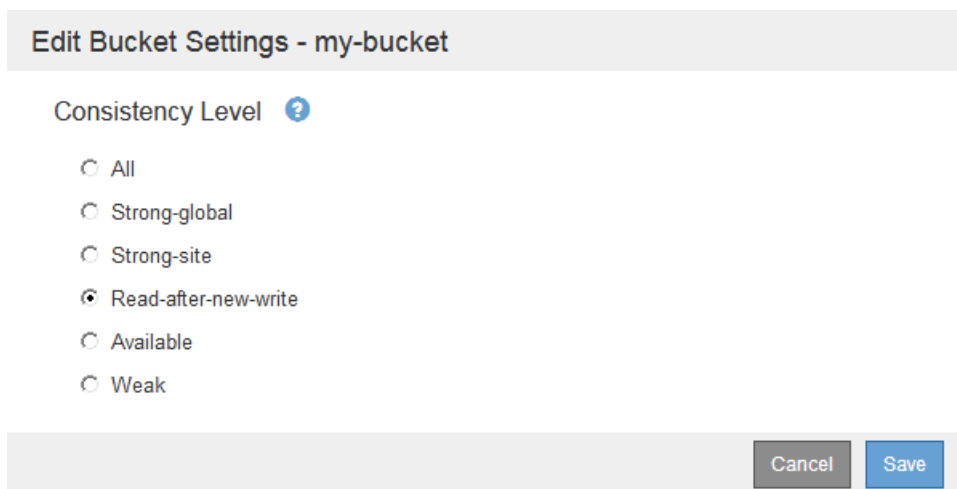
About this task

Consistency level makes a trade-off 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.

Note: 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. Click **S3 > Buckets**.
2. Select a bucket from the list.
3. Click **Configure Consistency Level**.
4. Select a consistency level for operations performed on the objects in this bucket.



Consistency level	Description
All	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.

Consistency level	Description
Strong-site	Guarantees read-after-write consistency for all client requests within a site.
Read-after-new-write	Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability, and data protection guarantees. Matches AWS S3 consistency guarantees. Note: If your application attempts HEAD operations on keys that do not exist, set the Consistency Level to Available , unless you require AWS S3 consistency guarantees. Otherwise, a high number of 500 Internal Server errors can result if one or more Storage Nodes are unavailable.
Available (eventual consistency for HEAD operations)	Behaves the same as the Read-after-new-write consistency level, but only provides eventual consistency for HEAD operations. Offers higher availability for HEAD operations than Read-after-new-write if Storage Nodes are unavailable. Differs from AWS S3 consistency guarantees for HEAD operations only.
Weak	Provides eventual consistency and high availability, with minimal data protection guarantees, especially if a Storage Node fails or is unavailable. Suitable only for write-heavy workloads that require high availability, do not require read-after-write consistency, and can tolerate the potential loss of data if a node fails.

5. Click **Save**.

Related concepts

[Tenant management permissions](#) on page 27

Configuring 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.

Before you begin

- You must be signed in to the Tenant Manager using a supported browser.
- Users must belong to a user group that has the **Manage All Containers** 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.

Example

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/2006-03-01/">
  <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>
```

See *Amazon Web Services (AWS) Documentation: Amazon Simple Storage Service Developer Guide* for more information about the CORS configuration XML.

2. In the Tenant Manager, go to **S3 > Buckets**.
3. Select the bucket from the list, and click **Configure CORS**.
4. Paste the CORS configuration XML into the text box, and click **Save**.

Example**Edit Bucket Settings - cors****Cross-Origin Resource Sharing (CORS) Configuration XML**

To configure CORS for objects in this bucket, enter the CORS configuration XML and click Save. To disable CORS, clear the field and click Save.

```
<CORSConfiguration
  xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <CORSRule>
    <AllowedOrigin>http://www.example.com</AllowedOrigin>
    <AllowedMethod>GET</AllowedMethod>
    <AllowedMethod>POST</AllowedMethod>
    <AllowedMethod>PUT</AllowedMethod>
    <AllowedHeader>*</AllowedHeader>
  </CORSRule>
</CORSConfiguration>
```

Cancel

Save

5. To modify the CORS setting for the bucket, update the CORS configuration XML in the text box, and click **Save**.
6. To disable CORS for the bucket, delete the CORS configuration XML from the text box, and click **Save**.

Related information

[Amazon Web Service \(AWS\) Documentation: Amazon Simple Storage Service Developer Guide](#)

Enabling or disabling 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.

Before you begin

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.

- You must be signed in to the Tenant Manager using a supported browser.
- Users must belong to a user group that has the **Manage All Containers** or the **Root Access** permission. These permissions override the permissions settings in group or bucket policies.

About this task

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.

Note: See the instructions for administering StorageGRID for more information.

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

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

Before enabling last access time for a bucket, be aware that the enabled setting can reduce StorageGRID performance, especially in systems with small objects. The performance impact occurs because StorageGRID must perform these additional steps every time objects are retrieved:

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

Steps

1. Click **S3 > Buckets**.
2. Select a bucket from the list.
3. Click **Configure Last Access Time**.
4. Select the check box if you to want to update the last access time when retrieving an object in this bucket, or unselect the box to if you want to disable last access time updates.

Edit Bucket Settings - test2

Last Access Time Updates [?](#)

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

Update access time when retrieving an object

Cancel Save

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

5. Click **Save**.

Related concepts

[Tenant management permissions](#) on page 27

Related information

[Administering StorageGRID](#)

Managing S3 platform services

If the use of platform services is allowed for your S3 tenant account, you can use platform services to leverage external services and configure CloudMirror replication, notifications, and search integration for S3 buckets.

What platform services are

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 AWS S3 and then leverage AWS services to perform analytics on your data.
- **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.
- **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.

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 AWS Simple Storage Service, while sending a notification about each such object to a third party monitoring application to help you track your AWS expenses.

Attention: 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. Tenant accounts that were created in versions of StorageGRID earlier than 11.0 have platform services disabled by default. Contact your grid administrator for more information.

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 or on AWS.

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 AWS 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 concepts

[Understanding the CloudMirror replication service](#) on page 60

[Understanding notifications for buckets](#) on page 61

[Understanding the search integration service](#) on page 62

[Considerations for using platform services](#) on page 63

Related information

[Implementing S3 client applications](#)

Understanding the 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.

In StorageGRID, you can configure more than one bucket as a destination for replication by specifying different destinations for each rule in the configuration XML. You can also 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 AWS 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 AWS S3 bucket as the destination.

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.

Related tasks

[Configuring CloudMirror replication](#) on page 72

Related information

[Amazon Web Services \(AWS\) Documentation: Cross-Region Replication](#)

Understanding 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 AWS S3 documentation.

Supported notifications and messages

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

- You cannot configure a notification for the `s3:ReducedRedundancyLostObject` event. This event type is not supported.
- 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

Related tasks

[Configuring event notifications](#) on page 74

Related information

[Amazon Web Services \(AWS\) Documentation: Configuring Amazon S3 Event Notifications](#)

Understanding the 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.

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 *Interoperability Matrix Tool* to determine the supported versions of Elasticsearch.

Related tasks

[Configuring the search integration service](#) on page 81

Related references

[Configuration XML for search integration](#) on page 78

[Object metadata included in metadata notifications](#) on page 83

[JSON generated by the search integration service](#) on page 83

Related information

[NetApp Interoperability Matrix Tool](#)

Considerations for using platform services

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

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.
Ordering guarantees	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. 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.

Consideration	Details
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 the CloudMirror replication service

Consideration	Details
Replication status	StorageGRID does not support the <code>x-amz-replication-status</code> header.
Object size	The maximum size for objects that can be replicated to a destination bucket by the CloudMirror replication service is 625 GiB (671,088,640,000 bytes). Objects that match the filtering criteria in the CloudMirror replication configuration XML will not be ingested if they are larger than this size.
Bucket versioning	<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. Version IDs for the source bucket in StorageGRID are not related to the version IDs for the destination bucket.</p>
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>

Configuring 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 an endpoint, you must be a tenant user who has been granted the appropriate permissions, in a grid whose networking has been configured to allow Storage Nodes to access external endpoint resources. Contact your grid administrator for more information.

Steps

1. [What an endpoint is](#) on page 65
2. [Specifying the URN for an endpoint](#) on page 65

3. [Creating an endpoint](#) on page 67
4. [Editing or removing an endpoint](#) on page 70
5. [Troubleshooting endpoint errors](#) on page 71

What an endpoint is

An endpoint stores the information that StorageGRID needs to use an external resource as a target for a platform service.

For example, to replicate objects from a StorageGRID bucket to an S3 bucket, the system needs to supply AWS with a bucket URI and credentials. This information is stored in an endpoint.

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 an 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.

To use an S3 bucket hosted on a StorageGRID system as an endpoint, specify an API Gateway Node in the endpoint definition.

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 hosted in an AWS cloud or in a local datacenter.

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.

Specifying the URN for an endpoint

When you create an 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 endpoints as you create them. Before you create an endpoint, confirm that the resource specified in the endpoint exists and that it can be reached.

URN elements

The URN for an endpoint must start with either `arn:aws` or `urn:mystore`, as follows:

- If the service is hosted on AWS, 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
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	<i>bucket-name</i>
Notifications	<i>sns-topic-name</i>
Search integration	<i>domain-name/index-name/type-name</i> Note: You must create the Elasticsearch index before you create the endpoint. Endpoint validation is done using the Elasticsearch index. The type will be dynamically created when object metadata is first sent to the destination.

URNs for services hosted on AWS

For AWS 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
```

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

URNs for locally-hosted services

For locally-hosted 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
```

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

Creating an endpoint

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

Before you begin

- You must be signed in to the Tenant Manager using a supported browser.
- Platform services have been enabled for your tenant account by a StorageGRID grid administrator.
- You belong to a user group that has the **Manage Endpoints** permission.
- The resource referenced by the endpoint has already been created:
 - **CloudMirror replication:** S3 bucket
 - **Event notification:** SNS topic
 - **Search integration:** Elasticsearch index

Note: You must create the Elasticsearch index before you create the endpoint. Endpoint validation is done using the Elasticsearch index. The type will be dynamically created when object metadata is first sent to the destination.
- You have the information needed to create the endpoint, including the endpoint's Uniform Resource Identifier (URI) and Unique Resource Name (URN); the access key ID and the secret access key for the destination service, if required; and the security certificate that will be used to validate TLS connections to the endpoint.

Steps

1. Select **S3 > Endpoints**.

The Endpoints page opens and shows the list of endpoints that have already been configured.

Endpoints

Endpoints enable platform services such as CloudMirror to direct their output to an external destination. You must configure an endpoint for each platform service you plan to use.

Display Name	Type	URI	URN	Access Key ID	Last Error
No endpoints found.					

2. Click **Create** to create a new endpoint.

Create Endpoint

Display Name

URI

URN

Access Key Id (optional)

Secret Access Key (optional)

Certificate Validation

3. Enter the following information:

Field	Description
Display Name	A name that briefly describes 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.
URI	The Unique Resource Identifier (URI) of the endpoint. Specify the endpoint URI in one of the following formats: <ul style="list-style-type: none"> https://host:port http://host:port <p>If you do not specify a port, port 443 is used for HTTPS URIs and port 80 is used for HTTP URIs.</p> <p>For example, an endpoint for a bucket hosted on StorageGRID might have a URI of the form <code>https://api-gateway-node.storagegrid.example.com:8082</code>, while the URI for a bucket hosted on AWS might be <code>https://s3-aws-region.amazonaws.com</code></p>
URN	See “Specifying the URN for an endpoint.” You cannot change this value after the endpoint is saved.
Access Key ID	Optionally, the Access Key ID for the destination service, formatted as an AWS access key. For anonymous access to the destination, omit both the Access Key ID and the Secret Access Key.

Field	Description
Secret Access Key	If you specified an Access Key ID, the Secret Access Key for the destination service, formatted as an AWS secret access key. A Secret Access Key is required when you specify an Access Key ID.
Certificate Validation	The method used to validate the certificate for TLS connections to the endpoint resource: <ul style="list-style-type: none"> Do not verify certificate: The certificate used for the TLS connection is not verified. 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 CA certificate installed on the operating system to secure connections.

4. Click **Save**.


When you save an endpoint, StorageGRID validates that the endpoint that you configured exists and that it can be reached using the credentials that you specified.

If endpoint validation fails, you receive an error message that explains why endpoint validation failed. Resolve the issue, then try creating the endpoint again.

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

5. If you need to test an existing endpoint, select the endpoint, and click **Test**.

- 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 it, click **Edit**, and update the information. Then, click **Save** to validate your changes.

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

After you finish

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

Related concepts

[Specifying the URN for an endpoint](#) on page 65

Related tasks

[Configuring CloudMirror replication](#) on page 72

Configuring event notifications on page 74

Configuring the search integration service on page 81

Editing or removing an endpoint

You can edit an endpoint to change its name, URI, or other details; however, you cannot change the URN for an endpoint. You can remove an endpoint if you no longer want to use the associated platform service.

Before you begin

- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a user group that has the **Manage Endpoints** permission.

Steps

1. Select **S3 > Endpoints**.

The Endpoints page opens and shows the list of endpoints that have already been configured.

2. Optionally, edit an endpoint.

- a. Select the radio button for the endpoint.
- b. Click **Edit**.
- c. As required, change the display name, URI, credentials, or certificate.

For example, you might need to update expired credentials or change the URI to point to a backup Elasticsearch index for failover.

Attention: You cannot change the URN for an endpoint.

- d. Click **Save**.

When you save an endpoint, StorageGRID validates that the endpoint exists and that it can be reached with the current credentials. The connection to the endpoint is validated from one node at each site.

If endpoint validation fails, you receive an error message that explains why validation failed.

3. Optionally, remove an endpoint.

Attention: If you remove an 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.

- a. Select the radio button for the endpoint.
- b. Click **Remove**.
A confirmation warning is displayed.
- c. Click **OK**.

Troubleshooting endpoint errors

If an error occurs when StorageGRID attempts to communicate with a specific endpoint, a message is displayed on the Dashboard and the error message is shown in the Last Error column on the Endpoints page.

Determining if an error has occurred

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

Dashboard

One or more endpoints have recently encountered an error and might not be functioning properly. Go to the [Endpoints](#) page to view the error details. The last error occurred 5 days ago.

When you go to the Endpoints page, you can review the more detailed error message in the **Last Error** column. This column displays only the most recent error message for each endpoint, and it indicates how long ago the error occurred. Errors in red occurred within the past 7 days.

Endpoints

Endpoints enable platform services such as CloudMirror to direct their output to an external destination. You must configure an endpoint for each platform service you plan to use.

⚠ One or more endpoints have experienced an error. The Last Error column shows how long ago the error occurred. When the error is resolved, the platform service request will be retried automatically--you do not need to perform additional steps. To see if an error is current or to force the removal of a resolved error from the table, click Test or edit and resave the endpoint.

<input type="button" value="+ Create"/> <input type="button" value="Edit"/> <input type="button" value="Remove"/> <input type="button" value="Test"/>						
Display Name	Type	URI	URN	Access Key ID	Last Error	
<input type="radio"/> My endpoint	S3 Bucket	https://10.96.98.99:18082	arn:aws:s3:::dest	*****AEDI	Endpoint had an AWS failure: RequestError: send request failed caused by: Put https://10.96.98.99:18082/dest/replicatedtest3: dial tcp 10.96.98.99:18082: connect: connection refused 38 minutes ago	

Displaying 1 endpoint.

Checking if an 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, select the radio button for the endpoint, and click **Test**.

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

Resolving endpoint errors

If an endpoint error occurs, you can use the message in the **Last Error** column 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, clicking **Save** 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.

Configuring CloudMirror replication

The StorageGRID CloudMirror replication service enables a tenant to automatically replicate objects to an external S3 bucket. You can enable replication using the Tenant Manager.

Before you begin

- You must have already created a bucket to act as the replication source. Buckets are created using the StorageGRID S3 API.
- Platform services must be enabled for your tenant account by a StorageGRID grid administrator.
- 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 Containers** 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.

For general information on bucket replication and how to configure it, see the Amazon 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:
 - a. 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.
- 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.

Example

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

- In the Tenant Manager go to **S3 > Buckets**.
- Select the source bucket, then click **Configure Replication**.
- Paste the replication configuration into the text box, and click **Save**.

Edit Bucket Settings - test1

Replication Configuration XML

To configure replication for objects in this bucket, enter the replication configuration XML and click Save. To disable replication, clear the field and click Save.

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

Cancel
Save

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

2. Verify that replication is configured correctly:
 - a. 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 “2017” are replicated.

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

For small objects, replication happens quickly.

You have configured your source bucket for StorageGRID bucket replication.

Related concepts

[Understanding the CloudMirror replication service](#) on page 60

Related tasks

[Creating an endpoint](#) on page 67

Related information

[Implementing S3 client applications](#)

[Amazon Web Services \(AWS\) Documentation: Cross-Region Replication](#)

Configuring event notifications

Enabling S3 event notifications for a bucket enables a tenant to send notifications about specified events to a destination service that supports the AWS Simple Notification Service™ (SNS). You can configure notifications for a bucket using the Tenant Manager.

Before you begin

- You must have already created a bucket to act as the source of notifications. Buckets are created using the StorageGRID S3 API.
- Platform services must be enabled for your tenant account by a StorageGRID grid administrator.
- 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 Containers** 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:
 - a. 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.

Example

```
<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>
```

- b. In the Tenant Manager go to **S3 > Buckets**.
- c. Select the source bucket, then click **Configure Notifications**.
- d. Paste the notification configuration XML into the text box, and click **Save**.

Edit Bucket Settings - test1

Event Notification Configuration XML

To configure notifications for events in this bucket, enter the notification configuration XML and click Save. To disable notifications, clear the field and click Save.

```
<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>
```

Cancel

Save

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

2. Verify that event notifications are configured correctly:
 - a. 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.

Example

```
{
  "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"
      }
    }
  ]
}

```

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.

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

Related concepts

[Understanding notifications for buckets](#) on page 61

Related tasks

[Creating an endpoint](#) on page 67

Related information

[Implementing S3 client applications](#)

[Amazon Web Services \(AWS\) Documentation: Configuring Amazon S3 Event Notifications](#)

Configuring the search integration service

Search integration is a custom StorageGRID service that sends 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.

Note: 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 tasks

[Configuring the search integration service](#) on page 81

Related references

[Configuration XML for search integration](#) on page 78

[Object metadata included in metadata notifications](#) on page 83

[JSON generated by the search integration service](#) on page 83

Related information

[Implementing S3 client applications](#)

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

Name	Description	Required
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
Prefix	Objects that match the prefix are affected by the rule, and their metadata is sent to the specified destination. To match all objects, specify an empty prefix. Included in the Rule element.	Yes
Destination	Container tag for the destination of a rule. Included in the Rule element.	Yes
Urn	URN of the destination where object metadata is sent. Must be the URN of a StorageGRID endpoint with the following properties: <ul style="list-style-type: none"> • <code>es</code> must be the third element. • The URN must end with the index and type where the metadata is stored, in the form <code>domain-name/myindex/mytype</code>. <p>Endpoints are configured using the Tenant Manager or Tenant Management API. They take the following form:</p> <ul style="list-style-type: none"> • <code>arn:aws:es:region:account-ID:domain/mydomain/myindex/mytype</code> • <code>urn:mysite:es:::mydomain/myindex/mytype</code> <p>The endpoint must be configured before the configuration XML is submitted, or configuration will fail with a 404 error. 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:myses:es:::sgws-notifications/test1/all</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>
```

```

    </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-1:33333333: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-1:22222222:domain/es-domain/
graphics/videotype</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>

```

Disabling search integration for a bucket

There are two ways to disable search integration for a bucket. If you are using the S3 API directly, you can use a DELETE Bucket metadata notification request (see the instructions for implementing S3 client applications). If you are using the Tenant Manager, you can go to **S3 > Buckets** and delete the search integration configuration XML.

Related tasks

[Configuring the search integration service](#) on page 81

Related references

[Object metadata included in metadata notifications](#) on page 83

[JSON generated by the search integration service](#) on page 83

Related information

[Implementing S3 client applications](#)

Configuring the search integration service

Configuring the search integration service for an S3 bucket enables the grid to send object metadata to a destination Elasticsearch index. You can configure the search integration service using the Tenant Manager.

Before you begin

- You must have already created an S3 bucket whose contents you want to index. Buckets are created using the StorageGRID S3 API.
- Platform services must be enabled for your tenant account by a StorageGRID grid administrator.
- 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 Containers** 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. Enable search integration for your source bucket:
 - a. Use a text editor to create the metadata notification XML required to enable search integration.

See “Configuration XML for search integration” for more information. When configuring the XML, use the URN of a search integration endpoint as the destination.

Example

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

- b. In the Tenant Manager go to **S3 > Buckets**.
- c. Select the source bucket, then click **Configure Search Integration**.
- d. Paste the metadata notification configuration into the text box, and click **Save**.

Edit Bucket Settings - test1

Search Configuration XML

To send object metadata and tags to a search endpoint, enter the metadata notification configuration XML and click Save. To disable search integration, clear the field and click Save.

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

Cancel

Save

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

2. Verify that the search integration service is configured correctly:
 - a. 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.
 - b. Confirm that a JSON document that contains the object's metadata and tags was added to the search index specified in the endpoint.

You have configured your source bucket to support the search integration service.

Related concepts

[Understanding the search integration service](#) on page 62

Related tasks

[Creating an endpoint](#) on page 67

Related references

[Configuration XML for search integration](#) on page 78

Related information

[Implementing S3 client applications](#)

JSON generated by the 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": "3d6c7634a85436eee06d43415012855",
  "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	Description
Bucket and object information	bucket	Name of the bucket
	key	Object key name
	versionID	Object version, for objects in versioned buckets
System metadata	size	Object size (in bytes) as visible to an HTTP client
	md5	Object hash
User metadata	metadata <i>key:value</i>	All user metadata for the object, as key-value pairs
Tags	tags <i>key:value</i>	All object tags defined for the object, as key-value pairs

Copyright

Copyright © 2019 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[*doccomments@netapp.com*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277