



# **Manage tenants**

StorageGRID 11.7

NetApp  
April 12, 2024

# Table of Contents

- Manage tenants ..... 1
  - Manage tenants: Overview ..... 1
  - Create a tenant account ..... 2
  - Edit tenant account ..... 7
  - Change password for tenant's local root user ..... 9
  - Delete tenant account ..... 9
  - Manage platform services ..... 10
  - Manage S3 Select for tenant accounts ..... 19

# Manage tenants

## Manage tenants: Overview

As a grid administrator, you create and manage the tenant accounts that S3 and Swift clients use to store and retrieve objects.



Support for Swift client applications has been deprecated and will be removed in a future release.

### What are tenant accounts?

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

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

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

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



If you use the S3 client protocol, you can use S3 buckets and bucket policies to segregate objects between the departments in an enterprise. You don't need to use tenant accounts. See instructions for implementing [S3 buckets and bucket policies](#) for more information.

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

For more information, see [Use a tenant account](#).

### How do I create a tenant account?

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

- Basic information including the tenant name, client type (S3 or Swift) and optional storage quota.
- Permissions for the tenant account, such as whether the tenant account can use S3 platform services, configure its own identity source, use S3 Select, or use a grid federation connection.
- The initial root access for the tenant, based on whether the StorageGRID system uses local groups and users, identity federation, or single sign-on (SSO).

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

## What is Tenant Manager used for?

After you create the tenant account, tenant users can sign in to the Tenant Manager to perform tasks such as the following:

- Set up identity federation (unless the identity source is shared with the grid)
- Manage groups and users
- Use grid federation for account clone and cross-grid replication
- Manage S3 access keys
- Create and manage S3 buckets
- Use S3 platform services
- Use S3 Select
- Monitor storage usage



While S3 tenant users can create and manage S3 access key and buckets with the Tenant Manager, they must use an S3 client application to ingest and manage objects. See [Use S3 REST API](#) for details.



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

## Create a tenant account

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

The steps for creating a tenant account vary based on whether [identity federation](#) and [single sign-on](#) are configured and whether the Grid Manager account you use to create the tenant account belongs to an admin group with the Root access permission.

### Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the Root access or Tenant accounts permission.
- If the tenant account will use the identity source that was configured for the Grid Manager, and you want to grant Root access permission for the tenant account to a federated group, you have imported that federated group into the Grid Manager. You don't need to assign any Grid Manager permissions to this admin group. See [Manage admin groups](#).
- If you want to allow an S3 tenant to clone account data and replicate bucket objects to another grid using a grid federation connection:
  - You have [configured the grid federation connection](#).
  - The status of the connection is **Connected**.
  - You have Root access permission.
  - You have reviewed the considerations for [managing the permitted tenants for grid federation](#).

- If the tenant account will use the identity source that was configured for Grid Manager, you have imported the same federated group into Grid Manager on both grids.

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



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

## Access the wizard

### Steps

1. Select **TENANTS**.
2. Select **Create**.

## Enter details

### Steps

1. Enter details for the tenant.

Field	Description
Name	A name for the tenant account. Tenant names don't need to be unique. When the tenant account is created, it receives a unique, 20-digit account ID.
Description (optional)	<p>A description to help identify the tenant.</p> <p>If you are creating a tenant that will use a grid federation connection, optionally, use this field to help identify which is the source tenant and which is the destination tenant. For example, this description for a tenant created on Grid 1 will also appear for the tenant replicated to Grid 2: "This tenant was created on Grid 1."</p>
Client type	<p>The type of client protocol this tenant will use, either <b>S3</b> or <b>Swift</b>.</p> <p><b>Note:</b> Support for Swift client applications has been deprecated and will be removed in a future release.</p>
Storage quota (optional)	If you want this tenant to have a storage quota, a numerical value for the quota and the units.

2. Select **Continue**.

## Select permissions

### Steps

1. Optionally, select any permissions you want this tenant to have.



Some of these permissions have additional requirements. For details, select the help icon for each permission.

Permission	If selected...
Allow platform services	The tenant can use S3 platform services such as CloudMirror. See <a href="#">Manage platform services for S3 tenant accounts</a> .
Use own identity source	The tenant can configure and manage its own identity source for federated groups and users. This option is disabled if you have <a href="#">configured SSO</a> for your StorageGRID system.
Allow S3 Select	<p>The tenant can issue S3 SelectObjectContent API requests to filter and retrieve object data. See <a href="#">Manage S3 Select for tenant accounts</a>.</p> <p><b>Important:</b> SelectObjectContent requests can decrease load-balancer performance for all S3 clients and all tenants. Enable this feature only when required and only for trusted tenants.</p>
Use grid federation connection	<p>The tenant can use a grid federation connection.</p> <p>Selecting this option:</p> <ul style="list-style-type: none"><li>• Causes this tenant and all tenant groups and users added to the account to be cloned from this grid (the <i>source grid</i>) to the other grid in the selected connection (the <i>destination grid</i>).</li><li>• Allows this tenant to configure cross-grid replication between corresponding buckets on each grid.</li></ul> <p>See <a href="#">Manage the permitted tenants for grid federation</a>.</p> <p><b>Note:</b> You can only select <b>Use grid federation connection</b> when you are creating a new S3 tenant; you can't select this permission for an existing tenant.</p>

2. If you selected **Use grid federation connection**, select one of the available grid federation connections.

☒ Use grid federation connection ?

Connection name ?	Remote grid hostname ?	Connection status ?
Grid A-Grid B	10.96.104.230	Connected

3. Select **Continue**.

## Define root access and create tenant

### Steps

1. Define root access for the tenant account, based on whether your StorageGRID system uses identity federation, single sign-on (SSO), or both.

Option	Do this
If identity federation is not enabled	Specify the password to use when signing into the tenant as the local root user.
If identity federation is enabled	<ol style="list-style-type: none"><li>1. Select an existing federated group to have Root access permission for the tenant.</li><li>2. Optionally, specify the password to use when signing in to the tenant as the local root user.</li></ol>
If both identity federation and single sign-on (SSO) are enabled	Select an existing federated group to have Root access permission for the tenant. No local users can sign in.

2. Select **Create tenant**.

A success message appears, and the new tenant is listed on the Tenants page. To learn how to view tenant details and monitor tenant activity, see [Monitor tenant activity](#).

3. If you selected the **Use grid federation connection** permission for the tenant:
  - a. Confirm that an identical tenant was replicated to the other grid in the connection. The tenants on both grids will have the same 20-digit account ID, name, description, quota, and permissions.



If you see the error message “Tenant created without a clone,” refer to the instructions in [Troubleshoot grid federation errors](#).

- b. If you provided a local root user password when defining root access, [change the password for the local root user](#) for the replicated tenant.



A local root user can't sign in to Tenant Manager on the destination grid until the password is changed.

## Sign in to tenant (optional)

As required, you can sign in to the new tenant now to complete the configuration, or you can sign in to the tenant later. The sign-in steps depend on whether you are signed in to the Grid Manager using the default port (443) or a restricted port. See [Control access at external firewall](#).

### Sign in now

If you are using...	Do this...
Port 443 and you set a password for the local root user	<ol style="list-style-type: none"> <li>1. Select <b>Sign in as root</b>.  When you sign in, links appear for configuring buckets, identity federation, groups, and users.</li> <li>2. Select the links to configure the tenant account.  Each link opens the corresponding page in the Tenant Manager. To complete the page, see the <a href="#">instructions for using tenant accounts</a>.</li> </ol>
Port 443 and you did not set a password for the local root user	Select <b>Sign in</b> , and enter the credentials for a user in the Root access federated group.
A restricted port	<ol style="list-style-type: none"> <li>1. Select <b>Finish</b></li> <li>2. Select <b>Restricted</b> in the Tenant table to learn more about accessing this tenant account.  The URL for the Tenant Manager has this format:   <code>https://FQDN_or_Admin_Node_IP:port/?accountId=20-digit-account-id/</code> <ul style="list-style-type: none"> <li>◦ <i>FQDN_or_Admin_Node_IP</i> is a fully qualified domain name or the IP address of an Admin Node</li> <li>◦ <i>port</i> is the tenant-only port</li> <li>◦ <i>20-digit-account-id</i> is the tenant's unique account ID</li> </ul> </li> </ol>

## Sign in later

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



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

## Configure the tenant

Follow the instructions in [Use a tenant account](#) to manage tenant groups and users, S3 access keys, buckets, platform services, and account clone and cross-grid replication.

## Edit tenant account

You can edit a tenant account to change the display name, storage quota, or tenant permissions.



If a tenant has the **Use grid federation connection** permission, you can edit tenant details from either grid in the connection. However, any changes you make on one grid in the connection will not be copied to the other grid. If you want to keep the tenant details exactly in sync between grids, make the same edits on both grids. See [Manage the permitted tenants for grid federation connection](#).

### Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the Root access or Tenant accounts permission.

### Steps

1. Select **TENANTS**.

# Tenants

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

Create

Export to CSV

Actions

Search tenants by name or ID

Displaying 5 results

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

## 2. Locate the tenant account you want to edit.

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

## 3. Select the tenant. You can do either of the following:

- Select the checkbox for the tenant, and select **Actions > Edit**.
- Select the tenant name to display the details page, and select **Edit**.

## 4. Optionally, change the values for these fields:

- **Name**
- **Description**
- **Storage quota**

## 5. Select **Continue**.

## 6. Select or clear the permissions for the tenant account.

- If you disable **Platform services** for a tenant who is already using them, the services that they have configured for their S3 buckets will stop working. No error message is sent to the tenant. For example, if the tenant has configured CloudMirror replication for an S3 bucket, they can still store objects in the bucket, but copies of those objects will no longer be made in the external S3 bucket that they have configured as an endpoint. See [Manage platform services for S3 tenant accounts](#).
- Change the setting of **Uses own identity source** to determine whether the tenant account will use its own identity source or the identity source that was configured for the Grid Manager.

If **Uses own identity source** is:

- Disabled and selected, the tenant has already enabled its own identity source. A tenant must disable its identity source before it can use the identity source that was configured for the Grid Manager.
- Disabled and not selected, SSO is enabled for the StorageGRID system. The tenant must use the identity source that was configured for the Grid Manager.
- Select or clear the **Allow S3 Select** permission as needed. See [Manage S3 Select for tenant accounts](#).

- To remove the **Use grid federation connection** permission, follow the instructions for [removing a tenant's permission to use grid federation](#).

## Change password for tenant's local root user

You might need to change the password for a tenant's local root user if the root user is locked out of the account.

### Before you begin

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

### About this task

If single sign-on (SSO) is enabled for your StorageGRID system, the local root user can't sign in to the tenant account. To perform root user tasks, users must belong to a federated group that has the Root access permission for the tenant.

### Steps

1. Select **TENANTS**.

## Tenants

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

[Create](#)
[Export to CSV](#)
[Actions](#)

Displaying 5 results

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

2. Select the tenant account. You can do either of the following:
  - Select the checkbox for the tenant, and select **Actions** > **Change root password**.
  - Select the tenant's name to display the details page, and select **Actions** > **Change root password**.
3. Enter the new password for the tenant account.
4. Select **Save**.

## Delete tenant account

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

to the system.

### Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.
- You have removed all buckets (S3), containers (Swift), and objects associated with the tenant account.
- If the tenant is permitted to use a grid federation connection, you have reviewed the considerations for [deleting a tenant with the Use grid federation connection permission](#).

### Steps

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

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

3. To delete multiple tenants, select the checkboxes and select **Actions > Delete**.
4. To delete a single tenant, do either of the following:
  - Select the checkbox, and select **Actions > Delete**.
  - Select the tenant name to display the details page, and then select **Actions > Delete**.
5. Select **Yes**.

## Manage platform services

### Manage platform services for tenants: Overview

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

#### What are platform services?

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

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

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

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



CloudMirror replication has some important similarities and differences with the cross-grid replication feature. To learn more, see [Compare cross-grid replication and CloudMirror replication](#).



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

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

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



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

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

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



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

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

## Recommendations for using platform services

Before using platform services, be aware of the following recommendations:

- If an S3 bucket in the StorageGRID system has both versioning and CloudMirror replication enabled, you should also enable S3 bucket versioning for the destination endpoint. This allows CloudMirror replication to generate similar object versions on the endpoint.
- You should not use more than 100 active tenants with S3 requests requiring CloudMirror replication, notifications, and search integration. Having more than 100 active tenants can result in slower S3 client performance.
- Requests to an endpoint that can't be completed will be queued to a maximum of 500,000 requests. This limit is equally shared among active tenants. New tenants are allowed to temporarily exceed this 500,000 limit so that newly created tenants aren't unfairly penalized.

## Related information

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

## Network and ports for platform services

If you allow an S3 tenant to use platform services, you must configure networking for the grid to ensure that platform services messages can be delivered to their destinations.

You can enable platform services for an S3 tenant account when you create or update the tenant account. If platform services are enabled, the tenant can create endpoints that serve as a destination for CloudMirror

replication, event notifications, or search integration messages from its S3 buckets. These platform services messages are sent from Storage Nodes that run the ADC service to the destination endpoints.

For example, tenants might configure the following types of destination endpoints:

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

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

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

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

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



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

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


#### Related information

- [Use a tenant account](#)


## Per-site delivery of platform services messages

All platform services operations are performed on a per-site basis.

That is, if a tenant uses a client to perform an S3 API Create operation on an object by connecting to a Gateway Node at Data Center Site 1, the notification about that action is triggered and sent from Data Center Site 1.



If the client subsequently performs an S3 API Delete operation on that same object from Data Center Site 2, the notification about the delete action is triggered and sent from Data Center Site 2.



Make sure that the networking at each site is configured such that platform services messages can be delivered to their destinations.

## Troubleshoot platform services

The endpoints used in platform services are created and maintained by tenant users in the Tenant Manager; however, if a tenant has issues configuring or using platform services, you might be able to use the Grid Manager to help resolve the issue.

### Issues with new endpoints

Before a tenant can use platform services, they must create one or more endpoints using the Tenant Manager. Each endpoint represents an external destination for one platform service, such as a StorageGRID S3 bucket, an Amazon Web Services bucket, a Simple Notification Service topic, or an Elasticsearch cluster hosted locally or on AWS. Each endpoint includes both the location of the external resource and the credentials needed to access that resource.

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



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




Endpoint creation will fail if platform services aren't enabled for the tenant account.

## Issues with existing endpoints

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



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

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

## Platform services endpoints

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















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

5 endpoints

Create endpoint

Delete endpoint

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



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

## Issues related to proxy servers

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

settings of your proxy server to ensure that platform service-related messages aren't blocked.

## Determine if an error has occurred

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

## Client operations fail

Some platform services issues might cause client operations on the S3 bucket to fail. For example, S3 client operations will fail if the internal Replicated State Machine (RSM) service stops, or if there are too many platform services messages queued for delivery.

To check the status of services:

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **site > Storage Node > SSM > Services**.

## Recoverable and unrecoverable endpoint errors

After endpoints have been created, platform service request errors can occur for various reasons. Some errors are recoverable with user intervention. For example, recoverable errors might occur for the following reasons:

- The user's credentials have been deleted or have expired.
- The destination bucket does not exist.
- The notification can't be delivered.

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

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

If StorageGRID encounters an unrecoverable endpoint error, the Total Events (SMTT) legacy alarm is triggered in the Grid Manager. To view the Total Events legacy alarm:

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **site > node > SSM > Events**.
3. View Last Event at the top of the table.

Event messages are also listed in `/var/local/log/bycast-err.log`.

4. Follow the guidance provided in the SMTT alarm contents to correct the issue.
5. Select the **Configuration** tab to reset event counts.
6. Notify the tenant of the objects whose platform services messages have not been delivered.
7. Instruct the tenant to re-trigger the failed replication or notification by updating the object's metadata or tags.

The tenant can resubmit the existing values to avoid making unwanted changes.

## Platform services messages can't be delivered

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

If platform services messages can't be delivered because of an unrecoverable error, the Total Events (SMTT) legacy alarm is triggered in the Grid Manager.

## Slower performance for platform service requests

StorageGRID software might throttle incoming S3 requests for a bucket if the rate at which the requests are being sent exceeds the rate at which the destination endpoint can receive the requests. Throttling only occurs when there is a backlog of requests waiting to be sent to the destination endpoint.



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

CloudMirror requests are more likely to be affected by the performance of the destination endpoint because these requests typically involve more data transfer than search integration or event notification requests.

## Platform service requests fail

To view the request failure rate for platform services:

1. Select **NODES**.
2. Select **site > Platform Services**.
3. View the Request error rate chart.

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

## Platform services unavailable alert

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

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

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



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

### Additional troubleshooting guidance for platform services endpoints

For additional information see [Use a tenant account > Troubleshoot platform services endpoints](#).

#### Related information

- [Troubleshoot StorageGRID system](#)

## Manage S3 Select for tenant accounts

You can allow certain S3 tenants to use S3 Select to issue `SelectObjectContent` requests on individual objects.

S3 Select provides an efficient way to search through large amounts of data without having to deploy a database and associated resources to enable searches. It also reduces the cost and latency of retrieving data.

### What is S3 Select?

S3 Select allows S3 clients to use `SelectObjectContent` requests to filter and retrieve only the data needed from an object. The StorageGRID implementation of S3 Select includes a subset of S3 Select commands and features.

### Considerations and requirements for using S3 Select

#### Grid administration requirements

The grid administrator must grant tenants S3 Select ability. Select **Allow S3 Select** when [creating a tenant](#) or [editing a tenant](#).

#### Object format requirements

The object you want to query must be in one of the following formats:

- **CSV.** Can be used as is or compressed into GZIP or BZIP2 archives.
- **Parquet.** Additional requirements for Parquet objects:
  - S3 Select supports only columnar compression using GZIP or Snappy. S3 Select doesn't support whole-object compression for Parquet objects.
  - S3 Select doesn't support Parquet output. You must specify the output format as CSV or JSON.
  - The maximum uncompressed row group size is 512 MB.
  - You must use the data types specified in the object's schema.
  - You can't use INTERVAL, JSON, LIST, TIME, or UUID logical types.

#### Endpoint requirements

The `SelectObjectContent` request must be sent to a [StorageGRID load balancer endpoint](#).

The Admin and Gateway Nodes used by the endpoint must be one of the following:

- An SG100 or SG1000 appliance node
- A VMware-based software node
- A bare metal node running a kernel with cgroup v2 enabled

## General considerations

Queries can't be sent directly to Storage Nodes.



SelectObjectContent requests can decrease load-balancer performance for all S3 clients and all tenants. Enable this feature only when required and only for trusted tenants.

See the [instructions for using S3 Select](#).

To view [Grafana charts](#) for S3 Select operations over time, select **SUPPORT > Tools > Metrics** in the Grid Manager.

## Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

## Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.