



Use Cloud Storage Pools

StorageGRID

NetApp
July 19, 2022

Table of Contents

- Use Cloud Storage Pools 1
 - What a Cloud Storage Pool is 1
 - Lifecycle of a Cloud Storage Pool object 3
 - When to use Cloud Storage Pools 6
 - Considerations for Cloud Storage Pools 7
 - Comparing Cloud Storage Pools and CloudMirror replication 11
- Create a Cloud Storage Pool 12
- Edit a Cloud Storage Pool 23
- Remove a Cloud Storage Pool 24
- Troubleshoot Cloud Storage Pools 25

Use Cloud Storage Pools

What a Cloud Storage Pool is

A Cloud Storage Pool lets you use ILM to move object data outside of your StorageGRID system. For example, you might want to move infrequently accessed objects to lower-cost cloud storage, such as Amazon S3 Glacier, S3 Glacier Deep Archive, or the Archive access tier in Microsoft Azure Blob storage. Or, you might want to maintain a cloud backup of StorageGRID objects to enhance disaster recovery.

From an ILM perspective, a Cloud Storage Pool is similar to a storage pool. To store objects in either location, you select the pool when creating the placement instructions for an ILM rule. However, while storage pools consist of Storage Nodes or Archive Nodes within the StorageGRID system, a Cloud Storage Pool consists of an external bucket (S3) or container (Azure Blob storage).

The following table compares storage pools to Cloud Storage Pools and shows the high-level similarities and differences.

	Storage pool	Cloud Storage Pool
How is it created?	Using the ILM > Storage pools option in Grid Manager. You must set up storage grades before you can create the storage pool.	Using the ILM > Storage pools option in Grid Manager. You must set up the external bucket or container before you can create the Cloud Storage Pool.
How many pools can you create?	Unlimited.	Up to 10.

	Storage pool	Cloud Storage Pool
Where are objects stored?	On one or more Storage Nodes or Archive Nodes within StorageGRID.	<p>In an Amazon S3 bucket or Azure Blob storage container that is external to the StorageGRID system.</p> <p>If the Cloud Storage Pool is an Amazon S3 bucket:</p> <ul style="list-style-type: none"> You can optionally configure a bucket lifecycle to transition objects to low-cost, long-term storage, such as Amazon S3 Glacier or S3 Glacier Deep Archive. The external storage system must support the Glacier storage class and the S3 POST Object restore API. You can create Cloud Storage Pools for use with AWS Commercial Cloud Services (C2S), which supports the AWS Secret Region. <p>If the Cloud Storage Pool is an Azure Blob storage container, StorageGRID transitions the object to the Archive tier.</p> <p>Note: In general, do not configure Azure Blob Storage lifecycle management for the container used for a Cloud Storage Pool. POST Object restore operations on objects in the Cloud Storage Pool can be affected by the configured lifecycle.</p>
What controls object placement?	An ILM rule in the active ILM policy.	An ILM rule in the active ILM policy.
What data protection method is used?	Replication or erasure coding.	Replication.
How many copies of each object are allowed?	Multiple.	<p>One copy in the Cloud Storage Pool and, optionally, one or more copies in StorageGRID.</p> <p>Note: You cannot store an object in more than one Cloud Storage Pool at any given time.</p>
What are the advantages?	Objects are quickly accessible at any time.	Low-cost storage.

Lifecycle of a Cloud Storage Pool object

Before implementing Cloud Storage Pools, review the lifecycle of objects that are stored in each type of Cloud Storage Pool.

- [S3: Lifecycle of a Cloud Storage Pool object](#)
- [Azure: Lifecycle of a Cloud Storage Pool object](#)

S3: Lifecycle of a Cloud Storage Pool object


The figure shows the lifecycle stages of an object that is stored in an S3 Cloud Storage Pool.



In the figure and explanations, “Glacier” refers to both the Glacier storage class and the Glacier Deep Archive storage class, with one exception: the Glacier Deep Archive storage class does not support the Expedited restore tier. Only Bulk or Standard retrieval is supported.



The Google Cloud Platform (GCP) supports object retrieval from long-term storage without requiring a POST Restore operation.



1. Object stored in StorageGRID

To start the lifecycle, a client application stores an object in StorageGRID.

2. Object moved to S3 Cloud Storage Pool

- When the object is matched by an ILM rule that uses an S3 Cloud Storage Pool as its placement location, StorageGRID moves the object to the external S3 bucket specified by the Cloud Storage Pool.
- When the object has been moved to the S3 Cloud Storage Pool, the client application can retrieve it using an S3 GET Object request from StorageGRID, unless the object has been transitioned to Glacier storage.

3. Object transitioned to Glacier (non-retrievable state)

- Optionally, the object can be transitioned to Glacier storage. For example, the external S3 bucket might use lifecycle configuration to transition an object to Glacier storage immediately or after some number of days.



If you want to transition objects, you must create a lifecycle configuration for the external S3 bucket, and you must use a storage solution that implements the Glacier storage class and supports the S3 POST Object restore API.



Do not use Cloud Storage Pools for objects that have been ingested by Swift clients. Swift does not support POST Object restore requests, so StorageGRID will not be able to retrieve any Swift objects that have been transitioned to S3 Glacier storage. Issuing a Swift GET object request to retrieve these objects will fail (403 Forbidden).

- During the transition, the client application can use an S3 HEAD Object request to monitor the object's status.

4. Object restored from Glacier storage

If an object has been transitioned to Glacier storage, the client application can issue an S3 POST Object restore request to restore a retrievable copy to the S3 Cloud Storage Pool. The request specifies how many days the copy should be available in the Cloud Storage Pool and the data-access tier to use for the restore operation (Expedited, Standard, or Bulk). When the expiration date of the retrievable copy is reached, the copy is automatically returned to a non-retrievable state.




If one or more copies of the object also exist on Storage Nodes within StorageGRID, there is no need to restore the object from Glacier by issuing a POST Object restore request. Instead, the local copy can be retrieved directly, using a GET Object request.

5. Object retrieved

Once an object has been restored, the client application can issue a GET Object request to retrieve the restored object.

Azure: Lifecycle of a Cloud Storage Pool object

The figure shows the lifecycle stages of an object that is stored in an Azure Cloud Storage Pool.



1. Object stored in StorageGRID

To start the lifecycle, a client application stores an object in StorageGRID.

2. Object moved to Azure Cloud Storage Pool

When the object is matched by an ILM rule that uses an Azure Cloud Storage Pool as its placement location, StorageGRID moves the object to the external Azure Blob storage container specified by the Cloud Storage Pool



Do not use Cloud Storage Pools for objects that have been ingested by Swift clients. Swift does not support POST Object restore requests, so StorageGRID will not be able to retrieve any Swift objects that have been transitioned to the Azure Blob storage Archive tier. Issuing a Swift GET object request to retrieve these objects will fail (403 Forbidden).

3. Object transitioned to Archive tier (non-retrievable state)

Immediately after moving the object to the Azure Cloud Storage Pool, StorageGRID automatically transitions the object to the Azure Blob storage Archive tier.

4. Object restored from Archive tier

If an object has been transitioned to the Archive tier, the client application can issue an S3 POST Object

restore request to restore a retrievable copy to the Azure Cloud Storage Pool.

When StorageGRID receives the POST Object Restore, it temporarily transitions the object to the Azure Blob storage Cool tier. As soon as the expiration date in the POST Object restore request is reached, StorageGRID transitions the object back to the Archive tier.



If one or more copies of the object also exist on Storage Nodes within StorageGRID, there is no need to restore the object from the Archive access tier by issuing a POST Object restore request. Instead, the local copy can be retrieved directly, using a GET Object request.

5. Object retrieved

Once an object has been restored to the Azure Cloud Storage Pool, the client application can issue a GET Object request to retrieve the restored object.

Related information

[Use S3](#)

When to use Cloud Storage Pools

Cloud Storage Pools can provide significant benefits in several use cases.

Backing up StorageGRID data in an external location

You can use a Cloud Storage Pool to back up StorageGRID objects to an external location.

If the copies in StorageGRID are inaccessible, the object data in the Cloud Storage Pool can be used to serve client requests. However, you might need to issue S3 POST Object restore request to access the backup object copy in the Cloud Storage Pool.

The object data in a Cloud Storage Pool can also be used to recover data lost from StorageGRID because of a storage volume or Storage Node failure. If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID temporarily restores the object and creates a new copy on the recovered Storage Node.

To implement a backup solution:

1. Create a single Cloud Storage Pool.
2. Configure an ILM rule that simultaneously stores object copies on Storage Nodes (as replicated or erasure-coded copies) and a single object copy in the Cloud Storage Pool.
3. Add the rule to your ILM policy. Then, simulate and activate the policy.

Tiering data from StorageGRID to external location

You can use a Cloud Storage Pool to store objects outside of the StorageGRID system. For example, suppose you have a large number of objects that you need to retain, but you expect to access those objects rarely, if ever. You can use a Cloud Storage Pool to tier the objects to lower-cost storage and to free up space in StorageGRID.

To implement a tiering solution:

1. Create a single Cloud Storage Pool.

2. Configure an ILM rule that moves rarely used objects from Storage Nodes to the Cloud Storage Pool.
3. Add the rule to your ILM policy. Then, simulate and activate the policy.

Maintain multiple cloud endpoints

You can configure multiple Cloud Storage Pools if you want to tier or back up object data to more than one cloud. The filters in your ILM rules let you specify which objects are stored in each Cloud Storage Pool. For example, you might want to store objects from some tenants or buckets in Amazon S3 Glacier and objects from other tenant or buckets in Azure Blob storage. Or, you might want to move data between Amazon S3 Glacier and Azure Blob storage. When using multiple Cloud Storage Pools, keep in mind that an object can be stored in only one Cloud Storage Pool at a time.

To implement multiple cloud endpoints:

1. Create up to 10 Cloud Storage Pools.
2. Configure ILM rules to store the appropriate object data at the appropriate time in each Cloud Storage Pool. For example, store objects from bucket A in Cloud Storage Pool A, and store objects from bucket B in Cloud Storage Pool B. Or, store objects in Cloud Storage Pool A for some amount of time and then move them to Cloud Storage Pool B.
3. Add the rules to your ILM policy. Then, simulate and activate the policy.

Considerations for Cloud Storage Pools

If you plan to use a Cloud Storage Pool to move objects out of the StorageGRID system, you must review the considerations for configuring and using Cloud Storage Pools.

General considerations

- In general, cloud archival storage, such as Amazon S3 Glacier or Azure Blob storage, is an inexpensive place to store object data. However, the costs to retrieve data from cloud archival storage are relatively high. To achieve the lowest overall cost, you must consider when and how often you will access the objects in the Cloud Storage Pool. Using a Cloud Storage Pool is recommended only for content that you expect to access infrequently.
- Do not use Cloud Storage Pools for objects that have been ingested by Swift clients. Swift does not support POST Object restore requests, so StorageGRID will not be able to retrieve any Swift objects that have been transitioned to S3 Glacier storage or the Azure Blob storage Archive tier. Issuing a Swift GET object request to retrieve these objects will fail (403 Forbidden).
- Using Cloud Storage Pools with FabricPool is not supported because of the added latency to retrieve an object from the Cloud Storage Pool target.

Information required to create a Cloud Storage Pool

Before you can create a Cloud Storage Pool, you must create the external S3 bucket or the external Azure Blob storage container that you will use for the Cloud Storage Pool. Then, when you create the Cloud Storage Pool in StorageGRID, you must specify the following information:

- The provider type: Amazon S3 or Azure Blob storage.
- If you select Amazon S3, whether the Cloud Storage Pool is for use with the AWS Secret Region (**CAP (C2S Access Portal)**).

- The exact name of the bucket or container.
- The service endpoint needed to access the bucket or container.
- The authentication needed to access the bucket or container:
 - **S3**: Optionally, an access key ID and secret access key.
 - **C2S**: The complete URL for obtaining temporary credentials from the CAP server; a server CA certificate, a client certificate, a private key for the client certificate, and, if the private key is encrypted, the passphrase for decrypting it.
 - **Azure Blob storage**: An account name and account key. These credentials must have full permission for the container.
- Optionally, a custom CA certificate to verify TLS connections to the bucket or container.

Considerations for the ports used for Cloud Storage Pools

To ensure that the ILM rules can move objects to and from the specified Cloud Storage Pool, you must configure the network or networks that contain your system's Storage Nodes. You must ensure that the following ports can communicate with the Cloud Storage Pool.

By default, Cloud Storage Pools use the following ports:

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

You can specify a different port when you create or edit a Cloud Storage Pool.

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

Considerations for costs

Access to storage in the cloud using a Cloud Storage Pool requires network connectivity to the cloud. You must consider the cost of the network infrastructure you will use to access the cloud and provision it appropriately, based on the amount of data you expect to move between StorageGRID and the cloud using the Cloud Storage Pool.

When StorageGRID connects to the external Cloud Storage Pool endpoint, it issues various requests to monitor connectivity and to ensure it can perform the required operations. While some additional costs will be associated with these requests, the cost of monitoring a Cloud Storage Pool should only be a small fraction of the overall cost of storing objects in S3 or Azure.

More significant costs might be incurred if you need to move objects from an external Cloud Storage Pool endpoint back to StorageGRID. Objects might be moved back to StorageGRID in either of these cases:

- The only copy of the object is in a Cloud Storage Pool and you decide to store the object in StorageGRID instead. In this case, you simply reconfigure your ILM rules and policy. When ILM evaluation occurs, StorageGRID issues multiple requests to retrieve the object from the Cloud Storage Pool. StorageGRID then creates the specified number of replicated or erasure-coded copies locally. After the object is moved back to StorageGRID, the copy in the Cloud Storage Pool is deleted.
- Objects are lost because of Storage Node failure. If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID temporarily restores the object and creates a new copy on the recovered Storage Node.



When objects are moved back to StorageGRID from a Cloud Storage Pool, StorageGRID issues multiple requests to the Cloud Storage Pool endpoint for each object. Before moving large numbers of objects, contact technical support for help in estimating the time frame and associated costs.

S3: Permissions required for the Cloud Storage Pool bucket

The bucket policy for the external S3 bucket used for a Cloud Storage Pool must grant StorageGRID permission to move an object to the bucket, get an object's status, restore an object from Glacier storage when required, and more. Ideally, StorageGRID should have full-control access to the bucket (`s3:*`); however, if this is not possible, the bucket policy must grant the following S3 permissions to StorageGRID:

- `s3:AbortMultipartUpload`
- `s3:DeleteObject`
- `s3:GetObject`
- `s3:ListBucket`
- `s3:ListBucketMultipartUploads`
- `s3:ListMultipartUploadParts`
- `s3:PutObject`
- `s3:RestoreObject`

S3: Considerations for the external bucket's lifecycle

The movement of objects between StorageGRID and the external S3 bucket specified in the Cloud Storage Pool is controlled by ILM rules and the active ILM policy in StorageGRID. In contrast, the transition of objects from the external S3 bucket specified in the Cloud Storage Pool to Amazon S3 Glacier or S3 Glacier Deep Archive (or to a storage solution that implements the Glacier storage class) is controlled by that bucket's lifecycle configuration.

If you want to transition objects from the Cloud Storage Pool, you must create the appropriate lifecycle configuration on the external S3 bucket, and you must use a storage solution that implements the Glacier storage class and supports the S3 POST Object restore API.

For example, suppose you want all objects that are moved from StorageGRID to the Cloud Storage Pool to be transitioned to Amazon S3 Glacier storage immediately. You would create a lifecycle configuration on the external S3 bucket that specifies a single action (**Transition**) as follows:

```
<LifecycleConfiguration>
  <Rule>
    <ID>Transition Rule</ID>
    <Filter>
      <Prefix></Prefix>
    </Filter>
    <Status>Enabled</Status>
    <Transition>
      <Days>0</Days>
      <StorageClass>GLACIER</StorageClass>
    </Transition>
  </Rule>
</LifecycleConfiguration>
```

This rule would transition all bucket objects to Amazon S3 Glacier on the day they were created (that is, on the day they were moved from StorageGRID to the Cloud Storage Pool).



When configuring the external bucket's lifecycle, never use **Expiration** actions to define when objects expire. Expiration actions cause the external storage system to delete expired objects. If you later attempt to access an expired object from StorageGRID, the deleted object will not be found.

If you want to transition objects in the Cloud Storage Pool to S3 Glacier Deep Archive (instead of to Amazon S3 Glacier), specify `<StorageClass>DEEP_ARCHIVE</StorageClass>` in the bucket lifecycle. However, be aware that you cannot use the `Expedited` tier to restore objects from S3 Glacier Deep Archive.

Azure: Considerations for Access tier

When you configure an Azure storage account, you can set the default Access tier to Hot or Cool. When creating a storage account for use with a Cloud Storage Pool, you should use the Hot tier as the default tier. Even though StorageGRID immediately sets the tier to Archive when it moves objects to the Cloud Storage Pool, using a default setting of Hot ensures that you will not be charged an early deletion fee for objects removed from the Cool tier before the 30-day minimum.

Azure: Lifecycle management not supported

Do not use Azure Blob Storage lifecycle management for the container used with a Cloud Storage Pool. The lifecycle operations might interfere with Cloud Storage Pool operations.

Related information

- [Create a Cloud Storage Pool](#)
- [S3: Specify authentication details for a Cloud Storage Pool](#)
- [C2S S3: Specify authentication details for a Cloud Storage Pool](#)
- [Azure: Specify authentication details for a Cloud Storage Pool](#)

Comparing Cloud Storage Pools and CloudMirror replication

As you begin using Cloud Storage Pools, it might be helpful to understand the similarities and differences between Cloud Storage Pools and the StorageGRID CloudMirror replication service.

	Cloud Storage Pool	CloudMirror replication service
What is the primary purpose?	A Cloud Storage Pool acts as an archive target. The object copy in the Cloud Storage Pool can be the only copy of the object, or it can be an additional copy. That is, instead of keeping two copies on-premise, you can keep only one copy within StorageGRID and send a copy to the Cloud Storage Pool.	The CloudMirror replication service enables a tenant to automatically replicate objects from a bucket in StorageGRID (source) to an external S3 bucket (destination). CloudMirror replication creates an independent copy of an object in an independent S3 infrastructure.
How is it set up?	Cloud Storage Pools are defined in the same way as storage pools, using the Grid Manager or the Grid Management API. A Cloud Storage Pool can be selected as the placement location in an ILM rule. While a storage pool consists of a group of Storage Nodes, a Cloud Storage Pool is defined using a remote S3 or Azure endpoint (IP address, credentials, and so on).	A tenant user configures CloudMirror replication by defining a CloudMirror endpoint (IP address, credentials, and so on) using the Tenant Manager or the S3 API. After the CloudMirror endpoint is set up, any bucket owned by that tenant account can be configured to point to the CloudMirror endpoint.
Who is responsible for setting it up?	Typically, a grid administrator	Typically, a tenant user
What is the destination?	<ul style="list-style-type: none"> Any compatible S3 infrastructure (including Amazon S3) Azure Blob Archive tier 	<ul style="list-style-type: none"> Any compatible S3 infrastructure (including Amazon S3)
What causes objects to be moved to the destination?	One or more ILM rules in the active ILM policy. The ILM rules define which objects StorageGRID moves to the Cloud Storage Pool and when the objects are moved.	The act of ingesting a new object into a source bucket that has been configured with a CloudMirror endpoint. Objects that existed in the source bucket before the bucket was configured with the CloudMirror endpoint are not replicated, unless they are modified.

	Cloud Storage Pool	CloudMirror replication service
How are objects retrieved?	Applications must make requests to StorageGRID to retrieve objects that have been moved to a Cloud Storage Pool. If the only copy of an object has been transitioned to archival storage, StorageGRID manages the process of restoring the object so it can be retrieved.	Because the mirrored copy in the destination bucket is an independent copy, applications can retrieve the object by making requests either to StorageGRID or to the S3 destination. For example, suppose you use CloudMirror replication to mirror objects to a partner organization. The partner can use its own applications to read or update objects directly from the S3 destination. Using StorageGRID is not required.
Can you read from the destination directly?	No. Objects moved to a Cloud Storage Pool are managed by StorageGRID. Read requests must be directed to StorageGRID (and StorageGRID will be responsible for retrieval from Cloud Storage Pool).	Yes, because the mirrored copy is an independent copy.
What happens if an object is deleted from the source?	The object is also deleted in the Cloud Storage Pool.	The delete action is not replicated. A deleted object no longer exists in the StorageGRID bucket, but it continues to exist in the destination bucket. Similarly, objects in the destination bucket can be deleted without affecting the source.
How do you access objects after a disaster (StorageGRID system not operational)?	Failed StorageGRID nodes must be recovered. During this process, copies of replicated objects might be restored using the copies in the Cloud Storage Pool.	The object copies in the CloudMirror destination are independent of StorageGRID, so they can be accessed directly before the StorageGRID nodes are recovered.

Create a Cloud Storage Pool

When you create a Cloud Storage Pool, you specify the name and location of the external bucket or container that StorageGRID will use to store objects, the cloud provider type (Amazon S3 or Azure Blob Storage), and the information StorageGRID needs to access the external bucket or container.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.
- You have reviewed the guidelines for configuring Cloud Storage Pools.
- The external bucket or container referenced by the Cloud Storage Pool already exists.
- You have all of the authentication information needed to access the bucket or container.

About this task

A Cloud Storage Pool specifies a single external S3 bucket or Azure Blob storage container. StorageGRID validates the Cloud Storage Pool as soon as you save it, so you must ensure that the bucket or container specified in the Cloud Storage Pool exists and is reachable.

Steps

1. Select **ILM > Storage pools**.

The Storage Pools page appears. This page includes two sections: Storage Pools and Cloud Storage Pools.

Storage Pools

Storage Pools

A storage pool is a logical group of Storage Nodes or Archive Nodes and is used in ILM rules to determine where object data is stored.

Name	Used Space	Free Space	Total Capacity	ILM Usage
All Storage Nodes	1.10 MB	102.90 TB	102.90 TB	Used in 1 ILM rule

Displaying 1 storage pool.

Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

No Cloud Storage Pools found.

2. In the Cloud Storage Pools section of the page, select **Create**.

The Create Cloud Storage Pool dialog box appears.

Create Cloud Storage Pool

Display Name

Provider Type

Bucket or Container

Cancel Save

3. Enter the following information:

Field	Description
Display Name	A name that briefly describes the Cloud Storage Pool and its purpose. Use a name that will be easy to identify when you configure ILM rules.

Field	Description
Provider Type	<p>Which cloud provider you will use for this Cloud Storage Pool:</p> <ul style="list-style-type: none"> • Amazon S3: Select this option for an S3, C2S S3, or Google Cloud Platform (GCP) endpoint. • Azure Blob Storage <p>Note: When you select a Provider Type, the Service Endpoint, Authentication and Server Verification sections appear at the bottom on the page.</p>
Bucket or Container	<p>The name of the external S3 bucket or Azure container that was created for the Cloud Storage Pool. The name you specify here must exactly match the bucket or container's name or Cloud Storage Pool creation will fail. You cannot change this value after the Cloud Storage Pool is saved.</p>

- Complete the Service Endpoint, Authentication and Server Verification sections of the page, based on the selected provider type.
 - [S3: Specify authentication details for a Cloud Storage Pool](#)
 - [C2S S3: Specify authentication details for a Cloud Storage Pool](#)
 - [Azure: Specify authentication details for a Cloud Storage Pool](#)


S3: Specifying authentication details for a Cloud Storage Pool


When you create a Cloud Storage Pool for S3, you must select the type of authentication that is required for the Cloud Storage Pool endpoint. You can specify Anonymous or enter an Access Key ID and Secret Access Key.


What you'll need

- You have entered the basic information for the Cloud Storage Pool and specified **Amazon S3** as the provider type.


Create Cloud Storage Pool


Display Name  S3 Cloud Storage Pool


Provider Type  Amazon S3 ▼


Bucket or Container  my-s3-bucket

Service Endpoint


Protocol  HTTP HTTPS

Hostname  example.com or 0.0.0.0


Port (optional)  443

URL Style  Auto-Detect ▼

Authentication

Authentication Type  ▼

Server Verification

Certificate Validation  Use operating system CA certificate ▼

[Cancel](#) [Save](#)

- If you are using access key authentication, you know the Access Key ID and Secret Access Key for the external S3 bucket.

Steps

1. In the **Service Endpoint** section, provide the following information:

- a. Select which protocol to use when connecting to the Cloud Storage Pool.

The default protocol is HTTPS.

- b. Enter the server hostname or IP address of the Cloud Storage Pool.

For example:

`s3-aws-region.amazonaws.com`



Do not include the bucket name in this field. You include the bucket name in the **Bucket or Container** field.

- c. Optionally, specify the port that should be used when connecting to the Cloud Storage Pool.

Leave this field blank to use the default port: port 443 for HTTPS or port 80 for HTTP.

- d. Select the URL style for the Cloud Storage Pool bucket:

Option	Description
Virtual Hosted-Style	Use a virtual hosted-style URL to access the bucket. Virtual hosted-style URLs include the bucket name as part of the domain name, for example <code>https://bucket-name.s3.company.com/key-name</code> .
Path-Style	Use a path-style URL to access the bucket. Path-style URLs include the bucket name at the end, for example <code>https://s3.company.com/bucket-name/key-name</code> . Note: The path-style URL is being deprecated.
Auto-Detect	Attempt to automatically detect which URL style to use, based on the information provided. For example, if you specify an IP address, StorageGRID will use a path-style URL. Select this option only if you don't know which specific style to use.

2. In the **Authentication** section, select the type of authentication that is required for the Cloud Storage Pool endpoint.

Option	Description
Access Key	An Access Key ID and Secret Access Key are required to access the Cloud Storage Pool bucket.
Anonymous	Everyone has access to the Cloud Storage Pool bucket. An Access Key ID and Secret Access Key are not required.
CAP (C2S Access Portal)	Used for C2S S3 only. Go to C2S S3: Specifying authentication details for a Cloud Storage Pool .

3. If you selected Access Key, enter the following information:

Option	Description
Access Key ID	The Access Key ID for the account that owns the external bucket.
Secret Access Key	The associated Secret Access Key.

4. In the Server Verification section, select which method should be used to validate the certificate for TLS connections to the Cloud Storage Pool:

Option	Description
Use operating system CA certificate	Use the default Grid CA certificates installed on the operating system to secure connections.
Use custom CA certificate	Use a custom CA certificate. Select Select New , and upload the PEM-encoded CA certificate.
Do not verify certificate	The certificate used for the TLS connection is not verified.

5. Select **Save**.

When you save a Cloud Storage Pool, StorageGRID does the following:

- Validates that the bucket and the service endpoint exist and that they can be reached using the credentials that you specified.
- Writes a marker file to the bucket to identify the bucket as a Cloud Storage Pool. Never remove this file, which is named `x-ntap-sgws-cloud-pool-uuid`.

If Cloud Storage Pool validation fails, you receive an error message that explains why validation failed. For example, an error might be reported if there is a certificate error or if the bucket you specified does not already exist.

Error

422: Unprocessable Entity

Validation failed. Please check the values you entered for errors.

Cloud Pool test failed. Could not create or update Cloud Pool. Error from endpoint: NoSuchBucket: The specified bucket does not exist. status code: 404, request id: 4211567681, host id:

OK

See the instructions for [troubleshooting Cloud Storage Pools](#), resolve the issue, and then try saving the Cloud Storage Pool again.

C2S S3: Specify authentication details for a Cloud Storage Pool

To use the Commercial Cloud Services (C2S) S3 service as a Cloud Storage Pool, you must configure C2S Access Portal (CAP) as the authentication type, so that StorageGRID can request temporary credentials to access the S3 bucket in your C2S account.

What you'll need

- You have entered the basic information for an Amazon S3 Cloud Storage Pool, including the service endpoint.
- You know the complete URL that StorageGRID will use to obtain temporary credentials from the CAP server, including all the required and optional API parameters assigned to your C2S account.
- You have a server CA certificate issued by an appropriate Government Certificate Authority (CA). StorageGRID uses this certificate to verify the identity of the CAP server. The server CA certificate must use PEM encoding.
- You have a client certificate issued by an appropriate Government Certificate Authority (CA). StorageGRID uses this certificate to identify itself to the CAP server. The client certificate must use PEM encoding and must have been granted access to your C2S account.
- You have a PEM-encoded private key for the client certificate.
- If the private key for the client certificate is encrypted, you have the passphrase for decrypting it.

Steps


1. In the **Authentication** section, select **CAP (C2S Access Portal)** from the **Authentication Type** dropdown.

The CAP C2S authentication fields appear.

Create Cloud Storage Pool

Display Name  C2S Cloud Storage Pool

Provider Type  Amazon S3

Bucket or Container  my-c2s-bucket

Service Endpoint

Protocol  HTTP HTTPS

Hostname  s3-aws-region.amazonaws.com

Port (optional)  443

URL Style  Auto-Detect

Authentication

Authentication Type  CAP (C2S Access Portal)

Temporary Credentials URL  https://example.com/CAP/api/v1/cred


Server CA Certificate  [Select New](#)

Client Certificate  [Select New](#)

Client Private Key  [Select New](#)

Client Private Key Passphrase (optional) 

Server Verification

Certificate Validation  Use operating system CA certificate

Cancel

Save

2. Provide the following information:

- a. For **Temporary Credentials URL**, enter the complete URL that StorageGRID will use to obtain temporary credentials from the CAP server, including all the required and optional API parameters assigned to your C2S account.
- b. For **Server CA Certificate**, select **Select New**, and upload the PEM-encoded CA certificate that StorageGRID will use to verify the CAP server.
- c. For **Client Certificate**, select **Select New**, and upload the PEM-encoded certificate that StorageGRID will use to identify itself to the CAP server.
- d. For **Client Private Key**, select **Select New**, and upload the PEM-encoded private key for the client certificate.

If the private key is encrypted, the traditional format must be used. (PKCS #8 encrypted format is not supported.)

- e. If the client private key is encrypted, enter the passphrase for decrypting the client private key. Otherwise, leave the **Client Private Key Passphrase** field blank.

3. In the Server Verification section, provide the following information:

- a. For **Certificate Validation**, select **Use custom CA certificate**.
- b. Select **Select New**, and upload the PEM-encoded CA certificate.

4. Select **Save**.

When you save a Cloud Storage Pool, StorageGRID does the following:

- Validates that the bucket and the service endpoint exist and that they can be reached using the credentials that you specified.
- Writes a marker file to the bucket to identify the bucket as a Cloud Storage Pool. Never remove this file, which is named `x-ntap-sgws-cloud-pool-uuid`.

If Cloud Storage Pool validation fails, you receive an error message that explains why validation failed. For example, an error might be reported if there is a certificate error or if the bucket you specified does not already exist.

Error

422: Unprocessable Entity

Validation failed. Please check the values you entered for errors.

Cloud Pool test failed. Could not create or update Cloud Pool. Error from endpoint: NoSuchBucket: The specified bucket does not exist. status code: 404, request id: 4211567681, host id:

OK

See the instructions for [troubleshooting Cloud Storage Pools](#), resolve the issue, and then try saving the Cloud Storage Pool again.

Azure: Specify authentication details for a Cloud Storage Pool

When you create a Cloud Storage Pool for Azure Blob storage, you must specify an account name and account key for the external container that StorageGRID will use to store objects.

What you'll need

- You have entered the basic information for the Cloud Storage Pool and specified **Azure Blob Storage** as the provider type. **Shared Key** appears in the **Authentication Type** field.

Create Cloud Storage Pool

Display Name	<input type="text" value="Azure Cloud Storage Pool"/>
Provider Type	<input type="text" value="Azure Blob Storage"/>
Bucket or Container	<input type="text" value="my-azure-container"/>

Service Endpoint

URI	<input type="text" value="https://myaccount.blob.core.windows.net"/>
-----	--

Authentication

Authentication Type	Shared Key
Account Name	<input type="text"/>
Account Key	<input type="text"/>

Server Verification

Certificate Validation	<input type="text" value="Use operating system CA certificate"/>
------------------------	--

- You know the Uniform Resource Identifier (URI) used to access the Blob storage container used for the Cloud Storage Pool.
- You know the name of the storage account and the secret key. You can use the Azure portal to find these

values.

Steps

1. In the **Service Endpoint** section, enter the Uniform Resource Identifier (URI) used to access the Blob storage container used for the Cloud Storage Pool.

Specify the URI in one of the following formats:

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

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

Example URI for Azure Blob storage container:

`https://myaccount.blob.core.windows.net`

2. In the **Authentication** section, provide the following information:
 - a. For **Account Name**, enter the name of the Blob storage account that owns the external service container.
 - b. For **Account Key**, enter the secret key for the Blob storage account.



For Azure endpoints, you must use Shared Key authentication.

3. In the **Server Verification** section, select which method should be used to validate the certificate for TLS connections to the Cloud Storage Pool:

Option	Description
Use operating system CA certificate	Use the Grid CA certificates installed on the operating system to secure connections.
Use custom CA certificate	Use a custom CA certificate. Select Select New , and upload the PEM-encoded certificate.
Do not verify certificate	The certificate used for the TLS connection is not verified.

4. Select **Save**.

When you save a Cloud Storage Pool, StorageGRID does the following:

- Validates that the container and the URI exist and that they can be reached using the credentials that you specified.
- Writes a marker file to the container to identify it as a Cloud Storage Pool. Never remove this file, which is named `x-ntap-sgws-cloud-pool-uuid`.

If Cloud Storage Pool validation fails, you receive an error message that explains why validation failed. For example, an error might be reported if there is a certificate error or if the container you specified does not already exist.

See the instructions for [troubleshooting Cloud Storage Pools](#), resolve the issue, and then try saving the Cloud

Storage Pool again.

Edit a Cloud Storage Pool

You can edit a Cloud Storage Pool to change its name, service endpoint, or other details; however, you cannot change the S3 bucket or Azure container for a Cloud Storage Pool.

What you'll need

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

Steps

1. Select **ILM > Storage pools**.

The Storage Pools page appears. The Cloud Storage Pools table lists the existing Cloud Storage Pools.

Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

	Pool Name	URI	Pool Type	Container	Used in ILM Rule	Last Error
<input checked="" type="radio"/>	azure-endpoint	https://storagegrid.blob.core.windows.net	azure	azure-3	✓	
<input type="radio"/>	s3-endpoint	https://s3.amazonaws.com	s3	s3-1	✓	

Displaying 2 pools.

2. Select the radio button for the Cloud Storage Pool you want to edit.
3. Select **Edit**.
4. As required, change the display name, service endpoint, authentication credentials, or certificate validation method.



You cannot change the provider type or the S3 bucket or Azure container for a Cloud Storage Pool.

If you previously uploaded a server or client certificate, you can select **View Current** to review the certificate that is currently in use.

5. Select **Save**.

When you save a Cloud Storage Pool, StorageGRID validates that the bucket or container and the service endpoint exist, and that they can be reached using the credentials that you specified.

If Cloud Storage Pool validation fails, an error message is displayed. For example, an error might be reported if there is a certificate error.

See the instructions for [troubleshooting Cloud Storage Pools](#), resolve the issue, and then try saving the Cloud Storage Pool again.

Remove a Cloud Storage Pool

You can remove a Cloud Storage Pool that is not used in an ILM rule and that does not contain object data.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.
- You have confirmed that the S3 bucket or Azure container does not contain any objects. An error occurs if you attempt to remove a Cloud Storage Pool if it contains objects. See [Troubleshoot Cloud Storage Pools](#).



When you create a Cloud Storage Pool, StorageGRID writes a marker file to the bucket or container to identify it as a Cloud Storage Pool. Do not remove this file, which is named `x-ntap-sgws-cloud-pool-uuid`.

- You have already removed any ILM rules that might have used the pool.

Steps

1. Select **ILM > Storage pools**.

The Storage Pools page appears.

2. Select the radio button for a Cloud Storage Pool that is not currently used in an ILM rule.

You cannot remove a Cloud Storage Pool if it is used in an ILM rule. The **Remove** button is disabled.

Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

	Pool Name	URI	Pool Type	Container	Used in ILM Rule	Last Error
<input checked="" type="radio"/>	azure-endpoint	https://storagegrid.blob.core.windows.net	azure	azure-3	✓	
<input type="radio"/>	s3-endpoint	https://s3.amazonaws.com	s3	s3-1	✓	

Displaying 2 pools.

3. Select **Remove**.

A confirmation warning is displayed.

Warning

Remove Cloud Storage Pool

Are you sure you want to remove this Cloud Storage Pool: My Cloud Storage Pool?

Cancel

OK

4. Select **OK**.

The Cloud Storage Pool is removed.

Troubleshoot Cloud Storage Pools

If you encounter errors when creating, editing, or deleting a Cloud Storage Pool, use these troubleshooting steps to help resolve the issue.

Determine if an error has occurred

StorageGRID performs a simple health check on every Cloud Storage Pool once a minute to ensure that the Cloud Storage Pool can be accessed and that it is functioning correctly. If the health check detects an issue, a message is shown in the Last Error column of the Cloud Storage Pools table on the Storage Pools page.

The table shows the most recent error detected for each Cloud Storage Pool and indicates how long ago the error occurred.

Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

Pool Name	URI	Pool Type	Container	Used in ILM Rule	Last Error
<input checked="" type="radio"/> S3	10.96.106.142:18082	s3	s3	✓	Endpoint failure: DC2-S1-106-147: Could not create or update Cloud Storage Pool. Error from endpoint: RequestError: send request failed caused by: Get https://10.96.106.142:18082/s3-targetbucket/x-ntap-sgws-cloud-pool-uuid: net/http: request canceled while waiting for connection (Client.Timeout exceeded while awaiting headers) 8 minutes ago
<input type="radio"/> Azure	http://pboerkoe@10.96.100.254:10000/d evstoreaccount1	azure	azure	✓	

Displaying 2 pools.

In addition, a **Cloud Storage Pool connectivity error** alert is triggered if the health check detects that one or more new Cloud Storage Pool errors have occurred within the past 5 minutes. If you receive an email notification for this alert, go to the Storage Pool page (select **ILM > Storage pools**), review the error messages in the Last Error column, and refer to the troubleshooting guidelines below.

Check if an error has been resolved

After resolving any underlying issues, you can determine if the error has been resolved. From the Cloud Storage Pool page, select the radio button for the endpoint, and select **Clear Error**. A confirmation message indicates that StorageGRID has cleared the error for the Cloud Storage Pool.

Error successfully cleared. This error might reappear if the underlying problem is not resolved.



If the underlying problem has been resolved, the error message is no longer displayed. However, if the underlying problem has not been fixed (or if a different error is encountered), the error message will be shown in the Last Error column within a few minutes.

Error: This Cloud Storage Pool contains unexpected content

You might encounter this error when you try to create, edit, or delete a Cloud Storage Pool. This error occurs if the bucket or container includes the `x-ntap-sgws-cloud-pool-uuid` marker file, but that file does not have the expected UUID.

Typically, you will only see this error if you are creating a new Cloud Storage Pool and another instance of StorageGRID is already using the same Cloud Storage Pool.

Try these steps to correct the issue:

- Check to make sure that no one in your organization is also using this Cloud Storage Pool.
- Delete the `x-ntap-sgws-cloud-pool-uuid` file and try configuring the Cloud Storage Pool again.

Error: Could not create or update Cloud Storage Pool. Error from endpoint

You might encounter this error when you try to create or edit a Cloud Storage Pool. This error indicates that some kind of connectivity or configuration issue is preventing StorageGRID from writing to the Cloud Storage Pool.

To correct the issue, review the error message from the endpoint.

- If the error message contains `Get url: EOF`, check that the service endpoint used for the Cloud Storage Pool does not use the HTTP protocol for a container or bucket that requires HTTPS.
- If the error message contains `Get url: net/http: request canceled while waiting for connection`, verify that the network configuration allows Storage Nodes to access the service endpoint used for the Cloud Storage Pool.
- For all other endpoint error messages, try one or more of the following:
 - Create an external container or bucket with the same name you entered for the Cloud Storage Pool, and try to save the new Cloud Storage Pool again.
 - Correct the container or bucket name you specified for the Cloud Storage Pool, and try to save the new Cloud Storage Pool again.

Error: Failed to parse CA certificate

You might encounter this error when you try to create or edit a Cloud Storage Pool. The error occurs if StorageGRID could not parse the certificate you entered when configuring the Cloud Storage Pool.

To correct the issue, check the CA certificate you provided for issues.

Error: A Cloud Storage Pool with this ID was not found

You might encounter this error when you try to edit or delete a Cloud Storage Pool. This error occurs if the endpoint returns a 404 response, which can mean either of the following:

- The credentials used for the Cloud Storage Pool do not have read permission for the bucket.
- The bucket used for the Cloud Storage Pool does not include the `x-ntap-sgws-cloud-pool-uuid` marker file.

Try one or more of these steps to correct the issue:

- Check that the user associated with the configured Access Key has the requisite permissions.
- Edit the Cloud Storage Pool with credentials that have the requisite permissions.
- If the permissions are correct, contact support.

Error: Could not check the content of the Cloud Storage Pool. Error from endpoint

You might encounter this error when you try to delete a Cloud Storage Pool. This error indicates that some kind of connectivity or configuration issue is preventing StorageGRID from reading the contents of the Cloud Storage Pool bucket.

To correct the issue, review the error message from the endpoint.

Error: Objects have already been placed in this bucket

You might encounter this error when you try to delete a Cloud Storage Pool. You cannot delete a Cloud Storage Pool if it contains data that was moved there by ILM, data that was in the bucket before you configured the Cloud Storage Pool, or data that was put in the bucket by some other source after the Cloud Storage Pool was created.

Try one or more of these steps to correct the issue:

- Follow the instructions for moving objects back to StorageGRID in “Lifecycle of a Cloud Storage Pool object.”
- If you are certain the remaining objects were not placed in the Cloud Storage Pool by ILM, manually delete the objects from the bucket.



Never manually delete objects from a Cloud Storage Pool that might have been placed there by ILM. If you later attempt to access a manually deleted object from StorageGRID, the deleted object will not be found.

Error: Proxy encountered an external error while trying to reach the Cloud Storage Pool

You might encounter this error if you have configured a non-transparent Storage proxy between Storage Nodes and the external S3 endpoint used for the Cloud Storage Pool. This error occurs if the external proxy server cannot reach the Cloud Storage Pool endpoint. For example, the DNS server might not be able to resolve the hostname or there might be an external networking issue.

Try one or more of these steps to correct the issue:

- Check the settings for the Cloud Storage Pool (**ILM > Storage pools**).
- Check the networking configuration of the Storage proxy server.

Related information

[Lifecycle of a Cloud Storage Pool object](#)

Copyright Information

Copyright © 2022 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.