



# Cloud Sync Documentation

## Cloud Sync

NetApp  
May 08, 2024

This PDF was generated from <https://docs.netapp.com/us-en/cloudsync/index.html> on May 08, 2024.  
Always check [docs.netapp.com](https://docs.netapp.com) for the latest.

# Table of Contents

Cloud Sync Documentation	1
Discover what's new	1
Get started	1
Automate with APIs	1
Learn about Cloud Sync	1
Get help and connect with peers	1
Release notes	2
What's new in Cloud Sync	2
Known limitations	21
Concepts	23
Cloud Sync overview	23
How Cloud Sync works	23
How Cloud Sync licenses work	24
Accounts	25
Data privacy	26
Get started	28
Quick start for Cloud Sync	28
Networking overview	28
Preparing the source and target	30
Endpoints that are required for Cloud Sync	41
Install the data broker	42
Creating a sync relationship	53
Paying for sync relationships after your free trial ends	56
Tutorials	58
Copying ACLs	58
Syncing NFS data using data-in-flight encryption	60
Setting up the data broker to use an external HashiCorp Vault	65
Managing sync relationships	69
Performing an immediate data sync	69
Accelerating sync performance	69
Changing the settings for a sync relationship	70
Creating and viewing reports about paths	72
Deleting relationships	74
Manage data brokers	75
Data broker groups	75
Add a new data broker	75
View a data broker's configuration	77
Remove a data broker from a group	77
Edit a group's name	78
Address issues with a data broker	79
Defining a unified configuration for a data broker group	79
Associating users to an account	81
Uninstalling the data broker	83

Cloud Sync APIs	84
Getting started	84
API reference	85
Using list APIs	85
Cloud Sync technical FAQ	88
Getting started	88
Supported sources and targets	89
Networking	90
Data synchronization	90
Security	91
Permissions	91
Performance	92
Deleting things	93
Troubleshooting	93
Data broker deep dive	93
How to get help and find more information	94
Self-support resources	94
Chatting with NetApp cloud experts	94
Activating NetApp support	95
Contacting NetApp support	97
Sending AutoSupport messages to NetApp	99
Legal notices	100
Copyright	100
Trademarks	100
Patents	100
Privacy policy	100
Open source	100

# Cloud Sync Documentation

Cloud Sync offers a simple, secure, and automated way to migrate your data from any source destination to any target destination, in the cloud or on your premises.

## Discover what's new

[What's new in Cloud Sync](#)

## Get started

- [View supported sync relationships](#)
- [Quick start](#)

## Automate with APIs

- [Get started with APIs](#)
- [API reference](#)

## Learn about Cloud Sync

- [What it is](#)
- [How it works](#)
- [Licensing](#)

## Get help and connect with peers

- [FAQ](#)
- [NetApp Community: Cloud Data Services](#)

# Release notes

## What's new in Cloud Sync

NetApp periodically updates Cloud Sync to bring you new features, enhancements, and bug fixes.

### 12 Mar 2021

The standalone Cloud Sync service has been retired. You should now access Cloud Sync directly from [Cloud Manager](#) where all of the same features and functionality are available.

After logging in to Cloud Manager, you can switch to the **Sync** tab to view and manage your relationships, just like before.

[Go to the Cloud Sync section in the Cloud Manager documentation.](#)

### 9 Mar 2021

- Cloud Sync now supports sync relationships between ONTAP S3 Storage and SMB servers:
  - ONTAP S3 Storage to an SMB server
  - An SMB server to ONTAP S3 Storage

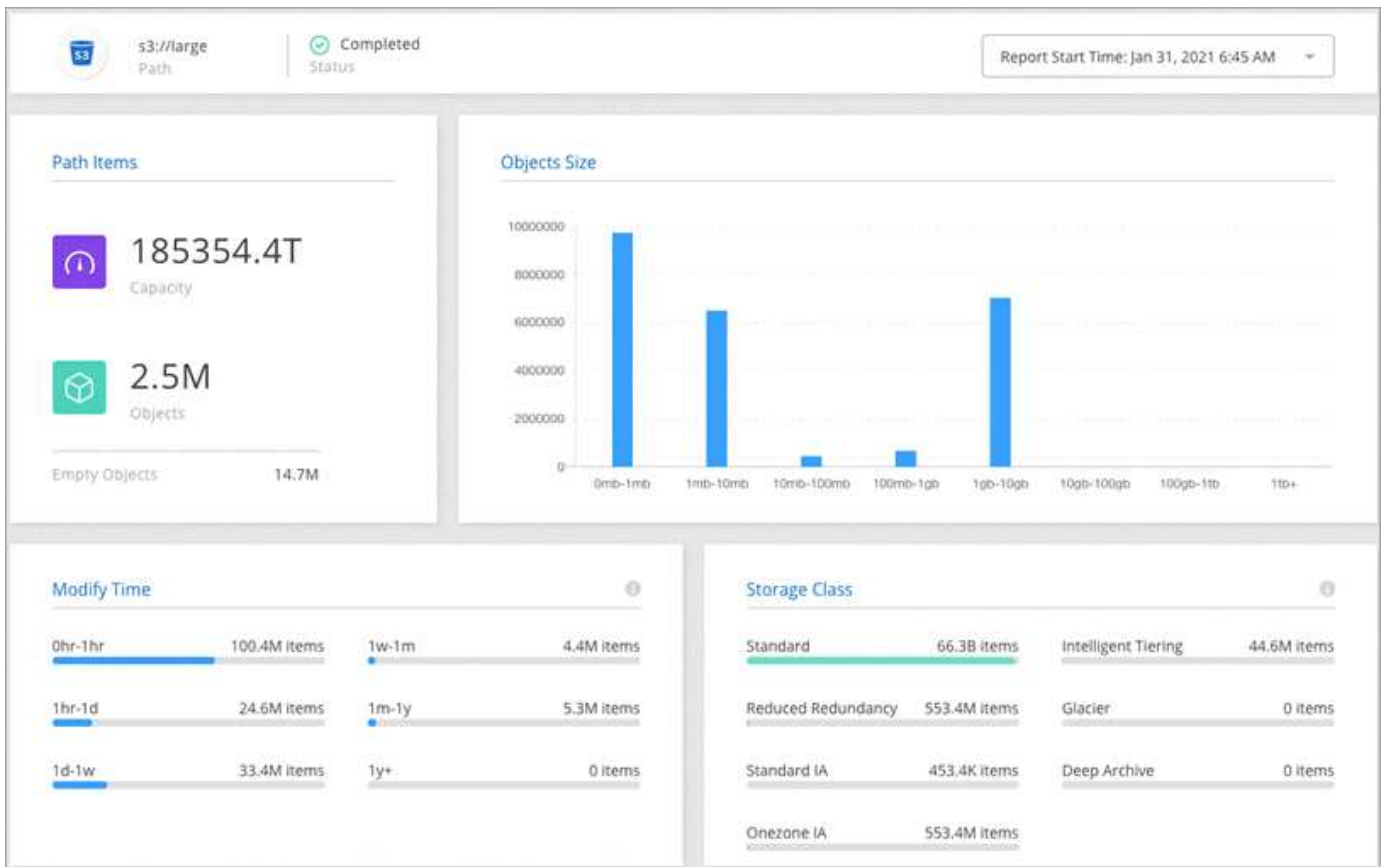
[View supported sync relationships.](#)
- Cloud Sync now enables you to unify a data broker group's configuration directly from the user interface.

We don't recommend changing the configuration on your own. You should consult with NetApp to understand when to change the configuration and how to change it.

[Learn more about defining a unified configuration.](#)

### 10 Feb 2021

In the last release, we introduced a new Reports feature that provides information that you can use with the help of NetApp personnel to tune a data broker's configuration and improve performance. These reports are now supported with object storage. [Learn more about these reports.](#)



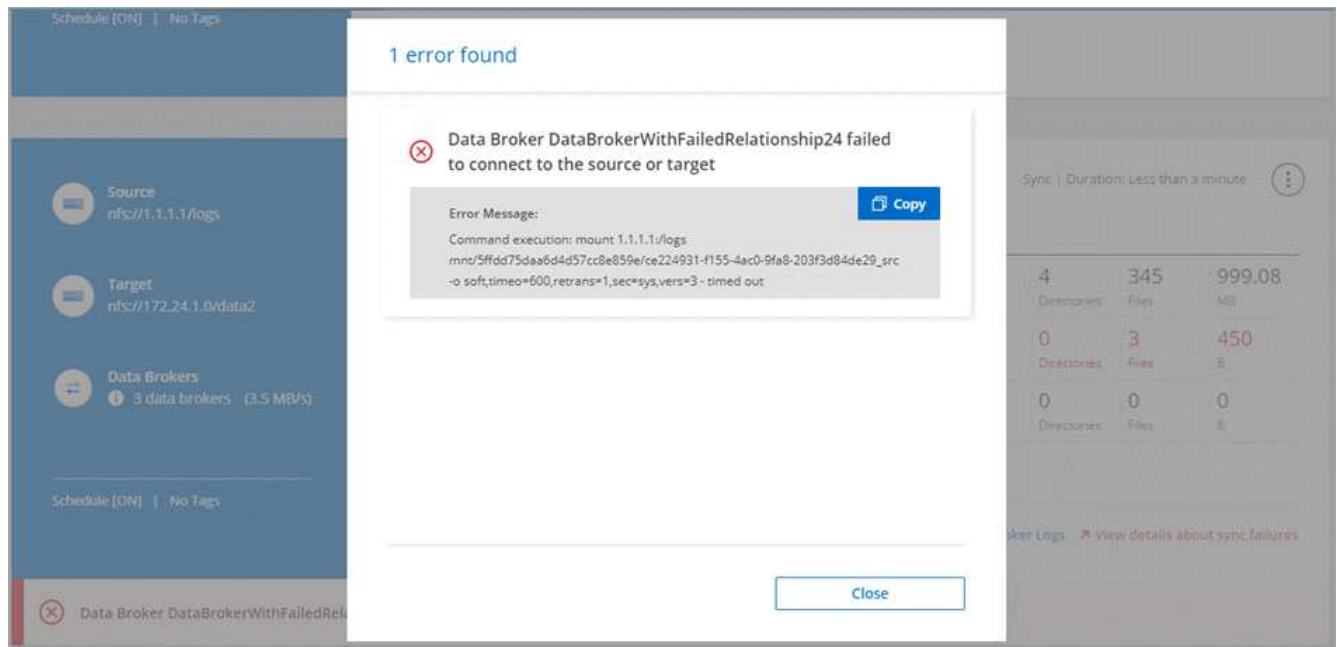
## 13 Jan 2021

- A new reporting feature provides information that you can use with the help of NetApp personnel to tune a data broker's configuration and improve performance.

Each report provides in-depth details about a path in a sync relationship: how many directories and files there are, the distribution of file size, how deep and wide the directories are, and more. [Learn more about these reports.](#)

- Cloud Sync now supports sync relationships from SFTP to S3 by using the API.
- Data broker connection errors now display on the Dashboard.

These errors can help you identify issues that prevent the data broker from connecting to the source or target in a sync relationship. The most typical issues are related to connectivity or permissions. Here's an example:



## 7 Dec 2020

- You can now manage data broker groups.

Grouping data brokers together can help improve the performance of sync relationships. Manage groups by adding a new data broker to a group, viewing information about data brokers, and more.

[Learn how to manage data brokers.](#)

- Cloud Sync now supports ONTAP S3 Storage to ONTAP S3 Storage sync relationships.

[View supported sync relationships.](#)

- You can now choose whether to copy access control lists (ACLs) between a source NFS server and a target NFS server when using NFS version 4 or later.

[Learn how to copy ACLs between NFS servers.](#)

## 1 Nov 2020

- When you create a data broker, Cloud Sync now enables you to specify a proxy server for the data broker.

### Proxy Configuration *(optional)*

Host

Example: 172.16.254.1

Port

Example: 8080

Define credentials for this proxy

User Name

Password

- Improved performance when copying large NFS files in a low latency environment.
- A new uninstall script for the data broker removes packages and directories that were created when the data broker was installed. [Learn more](#).

## 6 Sept 2020

This update includes performance improvements and usability enhancements to the data broker creation wizard.

## 2 Aug 2020

The deployment of the AWS data broker has been simplified. You no longer need to use a CloudFormation template to install the data broker in AWS. Cloud Sync now gives you the option to enter an AWS access key and then fill out a quick deployment wizard.

You still have the option to use a CloudFormation template, if you'd rather not provide an access key.

[Learn more about installing the data broker in AWS](#).

## 6 July 2020

- Cloud Sync now provides recommendations when it identifies ways for you to accelerate or optimize a sync relationship.



For example, Cloud Sync might recommend how to optimize a sync relationship for Azure NetApp Files. Or it might recommend that you accelerate a sync relationship if many large directories are being synced or if there are too many relationships per data broker.

You'll find these recommendations on the Sync Relationships dashboard. For example, here's a recommendation for an Azure NetApp Files relationship:

The screenshot displays the Cloud Sync Sync Relationships dashboard. On the left, a blue sidebar contains icons for Source, Target, and On-Prem Data Broker, along with a 'Schedule [ON]' button. The main area shows a 'Synced Successfully' status with a green checkmark. Below this is a 'Scan' table with the following data:

Scan	Directories	Files	MB
Succeeded	683	2.7K	457.42
Failed	0		


An information tooltip is visible over the 'On-Prem Data Broker' section, containing the text: 'Verify that you're using the Premium or Ultra service level with Azure NetApp Files. You might experience failures and performance issues if the disk service level is Standard. Consult a solutions architect if you need help determining the right service level. The volume size and volume tier determines the throughput that you can get.'

At the bottom right, there is an 'Optimization recommendation' link with an information icon.

- The deployment of the Google Cloud data broker has been simplified. You no longer need to manually enter commands to install the data broker in GCP. Cloud Sync now provides a Google login form and a quick deployment wizard.

[Learn more about installing the data broker in Google Cloud.](#)

- When an Amazon S3 bucket is the target in a sync relationship, you can now enable S3 bucket encryption by selecting an AWS KMS key, entering the ARN of a KMS key, or by selecting AES-256 encryption.



- SMB version 3.11 is now supported.



- Cloud Sync now shows the storage class used with an S3 sync relationship, even when it's the Standard storage class.

The screenshot shows the NetApp Cloud Sync configuration page. On the left, the 'Source' is an NFS volume at `nfs://.../mnt/kingkong...` and the 'Target' is an AWS S3 bucket at `s3://.../testush`. A dark blue box highlights the message: "Selected Storage Class for target bucket: 'Standard'". Below this, it says "Schedule [ON] | No Tags". On the right, a green checkmark icon indicates "Synced Successfully". Below this is a "Scan" table showing the results of the sync operation.

Scan			
Succeeded	684 Directories	2.7K Files	457.42 MB
Failed	0 Directories		
Marked for Copy	683 Directories	2.7K Files	457.42 MB
Marked for Delete	0 Directories	0 Files	0 B

## 4 June 2020

When you create a sync relationship, you can now refresh the list of directories or shares for the source or target.

The screenshot shows the "Select Directory" page in the NetApp Cloud Sync interface. The breadcrumb navigation at the top reads: "NFS Server > Data Broker > Directories > Target Cloud Volumes Service > Target Directories > Settings > Review". The main heading is "Select Directory". On the right, there is a link "Refresh Directories" with a circular arrow icon. A red arrow points to this link. Below the heading, there is a search bar and a list of directories. One directory is listed: `/MyExportName` with a blue checkmark icon to its left.

## 5 May 2020

- You can now sync NFS data from Azure NetApp Files to Azure NetApp Files using data-in-flight encryption. This makes it easy to securely transfer data across subnets or regions.

[Learn more about data-in-flight encryption.](#)

- ONTAP S3 is now supported in a sync relationship with StorageGRID.

ONTAP 9.7 supports the Amazon Simple Storage Service (Amazon S3) as a public preview. [Learn more about ONTAP support for Amazon S3.](#)

[Review Cloud Sync requirements for ONTAP S3 Storage.](#)

- Two new settings are available for sync relationships:

- Delete files on source

When you enable this setting, Cloud Sync deletes files from the source location after it copies the files to the target location. [Learn more about this setting.](#)

- Object tagging

When AWS S3 is the target in a sync relationship, Cloud Sync tags S3 objects with metadata that's relevant to the sync operation. You can disable tagging of S3 objects, if it's not desired in your environment. There's no impact to Cloud Sync if you disable tagging—Cloud Sync just stores the sync metadata in a different way.

## 5 Apr 2020

When you create a sync relationship, Cloud Sync might be unable to retrieve the shares for an SMB server. If this happens, you can now click **Add Share Manually** and enter the name of an SMB share.

The screenshot shows the Cloud Sync interface with a breadcrumb trail: SMB Server > Data Broker > Shares > Target SMB Server > Target Shares > Settings > Review. In the 'Target Shares' section, there is a button labeled '+ Add Share Manually'. A modal dialog box titled 'Add A Share Manually' is open. It contains the text: 'Enter the name of an SMB to manually add it. This step might be necessary if Cloud Sync cannot retrieve shares from the SMB server.' Below this text is a text input field labeled 'Share name' with the value '/MyShareName'. At the bottom of the dialog are two buttons: 'Add' (in blue) and 'Cancel' (in grey).

## 24 Mar 2020

Last month, we released a feature that enables Cloud Sync to copy access control lists (ACLs) between source SMB shares and target SMB shares. Prior to this update, the only data broker supported with this feature was the On-Prem Data Broker option with a CentOS 7.0 host.

Starting today, this feature works with *any* type of data broker: the AWS, Azure, Google Cloud Platform, or on-prem data broker. The on-prem data broker can run [any supported operating system](#).

[Learn more about copying ACLs between SMB servers.](#)

## 8 Mar 2020

- Credentials are now locally encrypted on the data broker machine.

The credentials that you provide while using the Cloud Sync service are stored directly on the data broker machine. These credentials are now encrypted using HashiCorp Vault.

- You can now download data broker logs at any time. Just click the download button that's available from each sync relationship.

The screenshot shows the AWS Cloud Sync console interface. On the left, a blue sidebar lists the components: Source (nfs://172.31.91.49/disk1/data/...), Target (s3://vadim-service-2test-empt...), and AWS Data Broker (vadimBroker1). Below this, it shows 'Schedule [ON]' and 'No Tags'. The main area has a green checkmark and 'Synced Successfully'. It displays 'Sync | Duration: 2 minutes | 5 days ago'. There are two tables: 'Scan' and 'Copy'. The 'Scan' table shows 42 Succeeded Directories, 181 Files, and 1.77 GB. The 'Copy' table shows 41 Succeeded Directories, 181 Files, and 1.77 GB. Both tables show 0 Failed and 0 Deleted items. At the bottom right, there is a red arrow pointing to a 'Download Data Broker Logs' button.

Scan			
Succeeded	42	181	1.77
	Directories	Files	GB
Failed	0		
	Directories		
Marked for Copy	41	181	1.77
	Directories	Files	GB
Marked for Delete	0	0	0
	Directories	Files	B

Copy			
Succeeded	41	181	1.77
	Directories	Files	GB
Failed	0	0	0
	Directories	Files	B
Deleted	0	0	0
	Directories	Files	B

## 23 Feb 2020

- Cloud Sync now supports multi-tenancy through Cloud Central accounts. Accounts enable multiple users to manage the same sync relationships in an account.

[Learn more about accounts.](#)

- We enhanced how Cloud Sync copies access control lists (ACLs) between SMB servers.

You no longer need to set up a Windows server to copy ACLs. You simply need to check a box when you create a relationship or after you create a relationship. Note that you'll need a data broker running on CentOS 7.0.

[Learn more about copying ACLs between SMB servers and review the requirements.](#)

- New sync relationships are supported:
  - Azure Blob Storage to Google Cloud Storage
  - Google Cloud Storage to Azure Blob Storage

[View supported sync relationships](#)

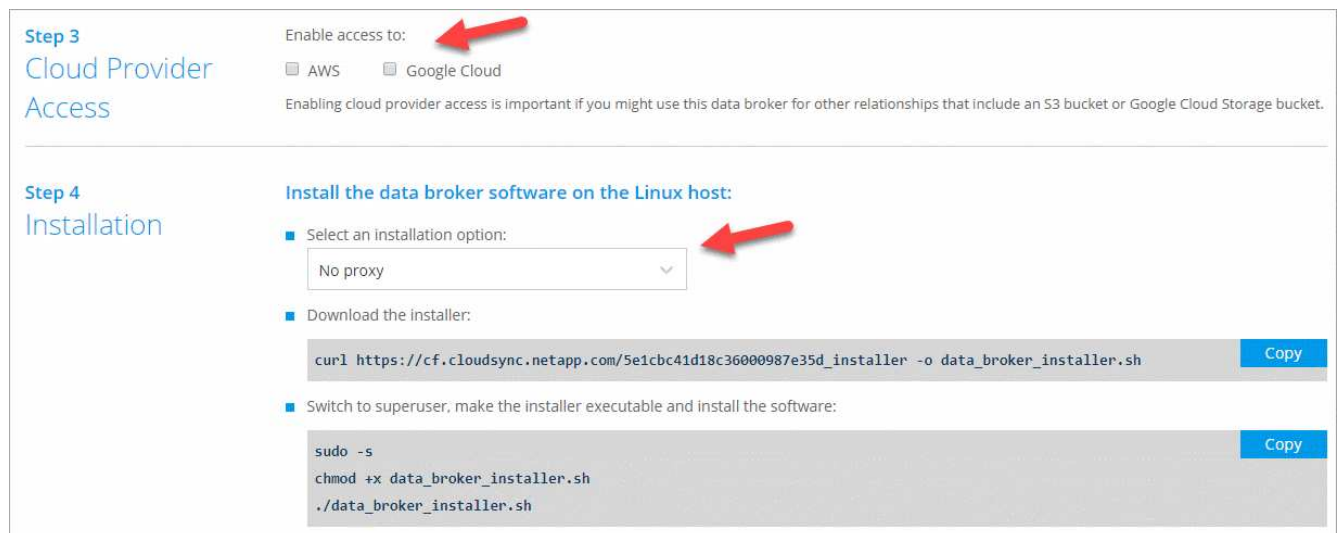
## 14 Jan 2020

- NFS 4.1 is now supported with Azure NetApp Files.



- We improved the sync performance of folders that contain greater than 10,000 files.
- A data broker that uses a proxy server can now be part of a sync relationship that includes an Azure Blob container.
- The on-prem data broker is now supported with additional Linux distributions:
  - CentOS 7.7 and 8.0
  - Red Hat Enterprise Linux 7.7 and 8.0
  - Ubuntu Server 16.04 LTS and 18.04 LTS
  - SUSE Linux Enterprise Server 12.4 SP5 and 15 SP1
- We simplified the installation steps for the on-premises data broker.

You now choose whether the data broker should have AWS access, GCP access, or both, and the type of proxy configuration that you have. Cloud Sync then shows you the right commands to use.



## 9 Dec 2019

- We improved how to schedule data syncs.

The settings for each sync relationship enables you to schedule recurring syncs. The following image shows the Schedule setting for a new sync relationship:



Schedule

☒ Start sync now
 ☐ Future sync
 ☐ One time copy

The first sync will start after the wizard is completed.

The next sync will automatically start at 14:00

and will repeat every 1 Days

Retries

Retry 3 times before skipping file

While this image shows the Schedule setting for an existing sync relationship:

General

Schedule

☒ Future sync
 ☐ Turn sync schedule off

The sync can be scheduled to start within the next 24 hours.

The next sync will automatically start at 05:30

and will repeat every 10 Minutes



You can schedule a relationship to sync data as often as every 1 minute.

- Additional statistics are available for each sync relationship on the **Sync Relationships** page. These stats enable you to see more details about the most recent data sync.

Here's an example:

Source

nfs://172.31.91.49/disk1/data/...

Target

nfs://172.31.91.49/disk2/targe...

Data Brokers

2 data brokers

Schedule [OFF]

No Tags

Sync Completed

Sync | Duration: 2 minutes | 15 days ago

Scan

Succeeded	8.2K Directories	18.3K Files	4.66 GB
Failed	1.8K Directories		
Marked for Copy	6.9K Directories	18.3K Files	4.66 GB
Marked for Delete	0 Directories	0 Files	0 B

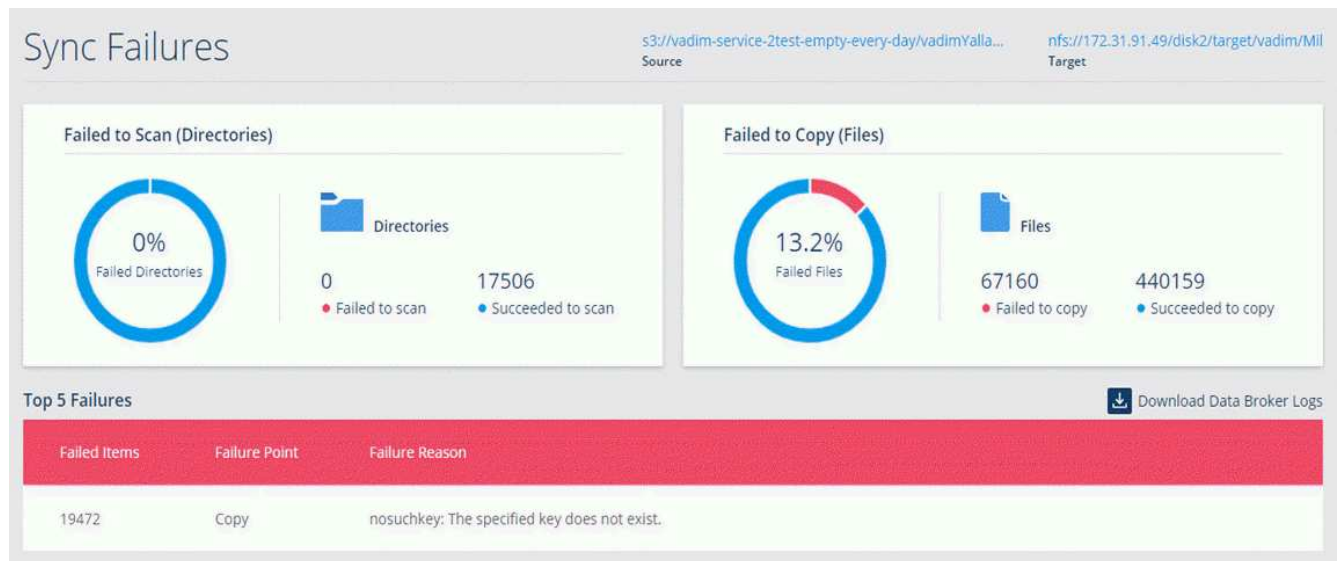
Copy

Succeeded	6.9K Directories	18.2K Files	4.56 GB
Failed	1 Directories	104 Files	95.37 MB
Deleted	0 Directories	0 Files	0 B

View details about sync failures


- Cloud Sync now shows you details about failures that occurred in the most recent data sync. If errors occur, you can use the failure reason to correct the issue.

Click **View details about sync failures** (as shown in the image above) and you'll see a page similar to the following:



Click **Download Data Broker Logs** to download logs that can help with troubleshooting.

- The sync status for an in-progress sync now shows an estimated time remaining that is more precise.



## 17 Nov 2019

- You can now contact NetApp technical support to get help with Cloud Sync. You will need to register your support serial number to activate support before you can contact NetApp technical support.

[Learn how to activate support and contact NetApp technical support.](#)

- The deployment of the Azure data broker has been simplified. You no longer need to manually enter commands to install the data broker in Azure. Cloud Sync now provides a Microsoft Azure login form and a quick deployment wizard.

[Learn more about installing the data broker in Azure.](#)

- When a sync is in progress, Cloud Sync now displays the remaining estimated time to finish syncing the data.
- Two new regions are now supported in AWS: Middle East (Bahrain) and Asia Pacific (Hong Kong).



## 4 Nov 2019

File metadata is now copied between sync relationships that include an S3 bucket and StorageGRID.



Cloud Sync doesn't sync any metadata that includes special characters.

## 8 Sept 2019

- You can now set up sync relationships to sync data between:
  - An AWS S3 bucket and an Azure Blob container
  - An AWS S3 bucket and a Google Cloud Storage bucket
- Azure NetApp Files (SMB) is now supported as the source or target in a sync relationship.

[View the list of supported sync relationships.](#)

- A new sync relationship setting enables you to sync files based on the last modified date. Choose all files regardless of their last modified date, files modified after a specific date, before a specific date, or between a time range.

## 15 July 2019

- You can now subscribe to Cloud Sync from Azure where you can pay as-you-go with hourly rates, or pay up front for a year.
  - [Learn more about paying for sync relationships](#)
  - [Learn how to subscribe from Azure](#)
- You can now deploy a data broker in AWS using your own IAM role, rather than the IAM role that Cloud Sync creates for you. You might use this option if your organization has strict security policies.

[Review details about using your own IAM role.](#)

- Cloud Sync now supports NFSv4 ACLs. When syncing data, Cloud Sync copies ACLs between NFS servers that use NFS versions 4.0, 4.1, or 4.2.
- When you create a sync relationship to or from Google Cloud Storage, Cloud Sync no longer prompts you to provide a project ID, client email, and private key for a Cloud Storage service account. GCP access must now be provided through the data broker.

Sync relationships that include GCP storage require a GCP data broker or an on-prem data broker that has GCP access:

- When you create a GCP data broker, Cloud Sync now prompts you for a service account that has "Storage Admin" permissions, along with the previously required permissions.

[Learn how to deploy the data broker in GCP.](#)

- When you deploy the data broker on an existing Linux host and GCP storage is the source or target in the relationship, Cloud Sync now prompts you to prepare the Linux host for GCP access.

[Learn how to install the data broker on an existing Linux host.](#)

## 20 June 2019

- New sync relationships are supported:
  - Azure NetApp Files to Azure Blob Storage
  - Azure Blob Storage to Azure NetApp Files

[View supported sync relationships](#)

- Additional S3 storage classes are now supported when AWS S3 is the target in a sync relationship:
  - Glacier
  - Glacier Deep Archive


[Learn about S3 storage classes](#)

- New settings enable you to define the number of retries and file types for a sync relationship.
- Sync relationship settings were moved to a new page when setting up a relationship and when editing a relationship.

Here's the Settings page when creating a new relationship:

Sync Relationship Settings		
You can modify settings that define how source files and folders are synced and maintained in the target location		
Recently Modified Files	Exclude files that are modified up to <b>30 Seconds</b> before a scheduled sync	▼
Delete Files On Target	Delete files from the target location if they were deleted from the source	▼
Retries	Retry 3 times before skipping file	▼
File Types	Include All: Files, Directories, Symbolic Links	▼

Here's where to access the Settings option for an existing relationship:



- We improved the speed of the user interface.
- A few bugs were fixed.

## 16 May 2019

You can now accelerate the performance of a sync relationship by adding an additional data broker to the relationship.

[Learn how to accelerate sync performance.](#)

## 21 Mar 2019

- You can now sync data between NFS servers using data-in-flight encryption.

[Learn more about data-in-flight encryption.](#)

- Two new sync relationships are supported:
  - Azure NetApp Files to Azure NetApp Files
  - AWS EFS to Azure NetApp Files

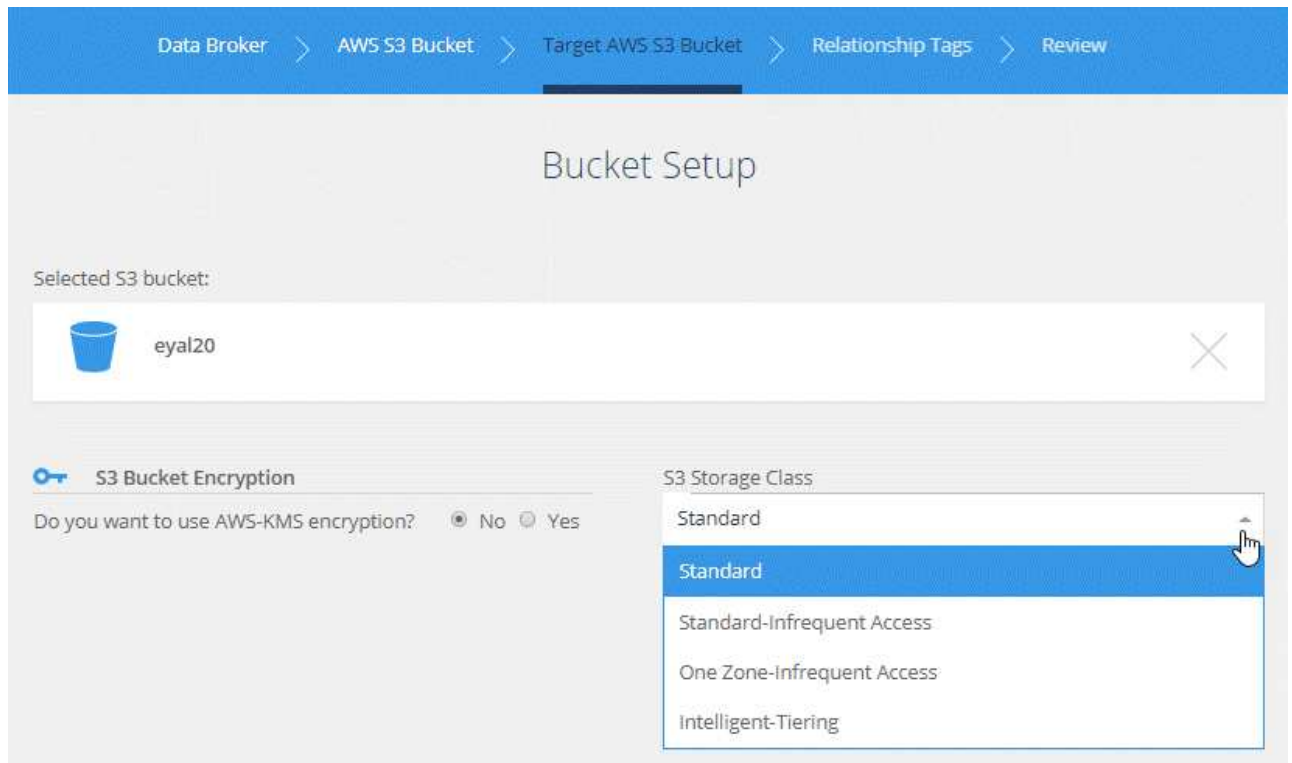
## 25 Feb 2019

Two new sync relationships are supported:

- StorageGRID to StorageGRID
- IBM Cloud Object Storage to IBM Cloud Object Storage

## 2 Dec 2018

- S3 to S3 sync relationships are now supported.
- When an S3 bucket is the target in a sync relationship, you can now choose an [S3 storage classes](#):
  - Standard (this is the default class)
  - Intelligent-Tiering
  - Standard-Infrequent Access
  - One Zone-Infrequent Access



## 8 Nov 2018

- Google Cloud Storage is now supported as the source or target in a sync relationship with an NFS server.

When setting up a relationship, you need to provide a project ID, client email, and private key for a Cloud Storage service account.

[Google Cloud Documentation: Creating and Managing Service Account Keys](#)

- You can now sync data between two Azure Blob containers.

When you set up the sync relationship, you need to provide a storage account connection string that includes a shared access signature (SAS).

[View requirements for Azure Blob storage](#)

- When creating a sync relationship between an Azure Blob container and an NFS or SMB server, you must now provide the storage account connection string.

[View requirements for Azure Blob storage](#)

## 11 Oct 2018

Cloud Sync now provides additional stats about the last data sync for each sync relationship:

- How many directories and files were scanned
- How many directories failed to scan
- How many directories and files were marked for copy and marked for deletion

For example, the following image shows that 684 directories were scanned and no files were marked for copy

or for deletion:

Cloud Sync

Sync Relationships

Cost

Timeline

Help

Hili at Netapp

Sync Relationships

Create New Sync Relationship

3 Syncs

Source

Target

AWS Data Broker

Schedule [ON]

No Tags

No Custom Settings

Synced Successfully

Sync | Duration: a minute | 3 days ago

Scan		
Succeeded		Failed
684 Directories	1.4K Files	457.42 MB
		0 Directories
Marked For Copy		
0 Directories	0 Files	0 B
Marked For Deletion		
0 Directories	0 Files	0 B

Copy			
Succeeded		Deleted	Failed
0 Files	0 B	0 Files	0 B
			0 Files
			0 B

FREE TRIAL [11 Days left]

API API Documentation

12 Sept 2018

Cloud Sync now supports deploying a data broker in Google Cloud Platform.

Just follow the prompts in Cloud Sync to deploy a virtual machine in Google Cloud Platform that runs the data broker software.

Add a NetApp Data Broker

aws

AWS Data Broker

Azure Data Broker

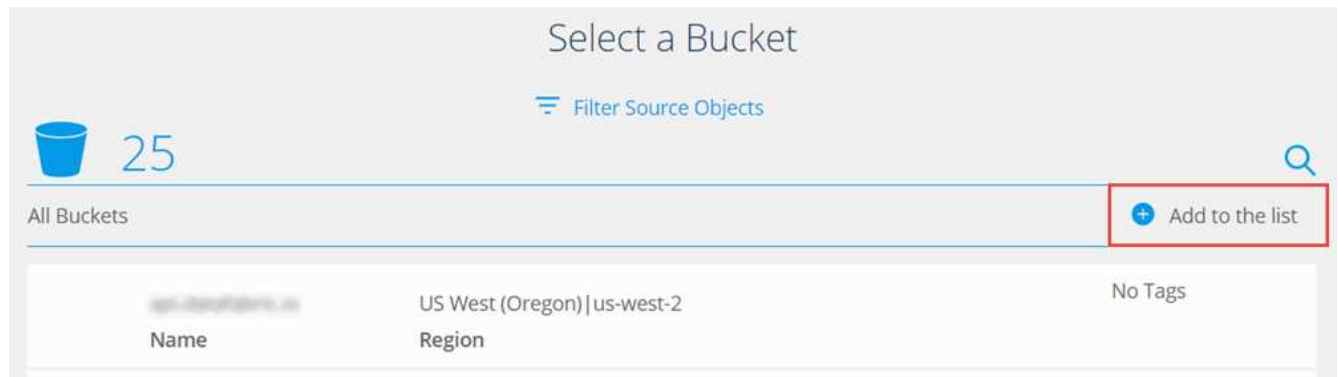
Google Cloud Data Broker

On-Prem Data Broker

Cancel

21 Aug 2018

- IBM Cloud Object Storage is now supported as the source or target in a sync relationship with an NFS or SMB server.
- When creating a new sync relationship, you can choose an S3 bucket that is not associated with your AWS account.



## 14 Aug 2018

Cloud Sync can now preserve access control lists (ACLs) between a source SMB/CIFS share and a target SMB/CIFS share when creating a new sync relationship.

[Learn more about copying ACLs between SMB/CIFS shares.](#)

## 17 July 2018

- You can now change the sync schedule for a relationship to as frequently as 5 minutes. The default is 24 hours.
- StorageGRID Webscale is now supported as the source in a sync relationship. The target can be an NFS or SMB server.

## 10 July 2018

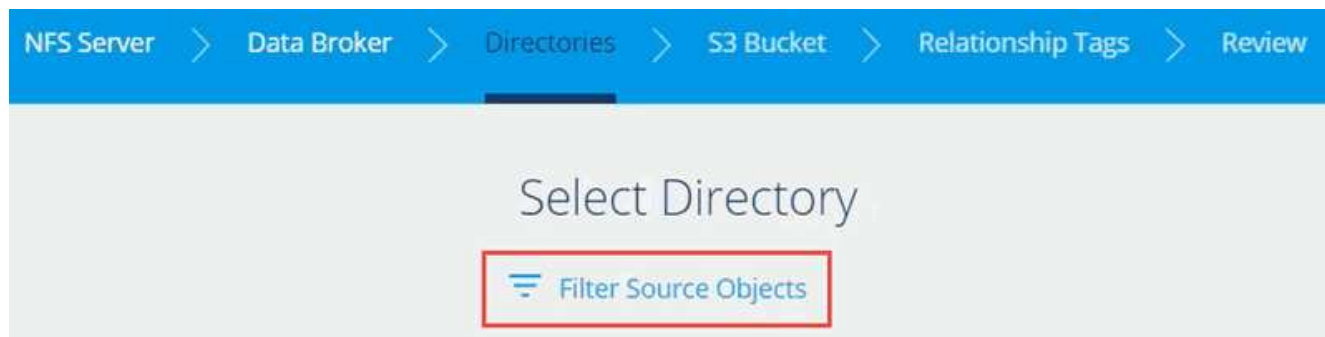
- An Azure Blob container is now supported as the source or target in a sync relationship with an NFS or SMB server.

When setting up the sync relationship, you simply need to enter the Azure storage account name and the access key for the storage account. Then you can select the Blob container.

- You must now select an NFS version or SMB version when setting up a new sync relationship for an NFS or SMB server.
  - For NFS, you can select version 3, 4.1, or 4.2.
  - For SMB, you can select version 1.0, 2.0, 2.1, or 3.0.
- You can now filter source objects when setting up a new sync relationship.

Filtering source objects enables you to define how source files and folders are synced and maintained in the target location.

You can access the option when selecting a directory:



The following options are available when filtering source objects:

General	
Schedule	ON   Every 1 Day
Retries	Retry 3 times before skipping file

Files and Directories	
Recently Modified Files	Exclude files that are modified up to 30 Seconds before a scheduled sync
Delete Files On Source	Never delete files from the source location
Delete Files On Target	Never delete files from the target location
Object Tagging	Allow Cloud Sync to tag S3 objects
File Types	Include All: Files, Directories, Symbolic Links
Exclude File Extensions	None
File Size	All
Date Modified	All

[Reset to defaults](#)

- Incremental updates from an S3 bucket to an NFS or SMB server are no longer event-driven—they are based on a sync schedule.

## April 2018

Cloud Sync now supports the NetApp Cloud Volumes Service as an NFS or SMB server in a sync relationship.

## February 2018

- EFS to S3 and S3 to EFS sync relationships are now supported.
- Bugs were fixed.

## January 2018

- You can now abort an in-progress sync.

This does not break the sync relationship. Cloud Sync syncs data at the next scheduled time.

- You can now view and select objects from S3 buckets that belong to other AWS accounts if they are shared with your account.
- You can now use Cloud Sync with S3 buckets that are protected with AWS KMS encryption.
- A few bugs were fixed. Most notably, you no longer have to enter AWS credentials when using the on-premises data broker.

## December 2017

- Cloud Sync now supports installing the data broker in Microsoft Azure, which enables you to sync data in and out of Azure.
- A few bugs were fixed.

## November 2017

- Cloud Sync is now integrated with NetApp Cloud Central, which enables centralized user authentication.
- EFS to NFS and NFS to EFS sync relationships are now supported.
- SMB to SMB sync relationships are now supported.
- A few bugs were fixed.

## October 2017

- NFS to NFS sync relationships are now supported.
- You can now specify whether files modified prior to the scheduled sync should be excluded.

For example, you can exclude files modified 30 seconds before the scheduled sync. This setting helps avoid copying partial changes to files that frequently change.

- Cloud Sync now displays the number of failed transfers.
- A few bugs were fixed.

## Known limitations

These known limitations identify platforms or features that are not supported.

### Unsupported regions

- Cloud Sync is not supported in China.



- In addition to China, the Cloud Sync data broker is not supported in the following regions:
  - AWS GovCloud (US)
  - Azure US Gov
  - Azure US DoD

## **SMB sync behavior due to case-insensitivity limitation**

The SMB protocol is case-insensitive, which means uppercase and lowercase letters are treated as being the same. This behavior can result in overwritten files and directory copy errors, if a sync relationship includes an SMB server and data already exists on the target.

For example, let's say that there's a file named "a" on the source and a file named "A" on the target. When Cloud Sync copies the file named "a" to the target, file "A" is overwritten by file "a" from the source.

In the case of directories, let's say that there's a directory named "b" on the source and a directory named "B" on the target. When Cloud Sync tries to copy the directory named "b" to the target, Cloud Sync receives an error that says the directory already exists. As a result, Cloud Sync always fails to copy the directory named "b."

The best way to avoid this limitation is to ensure that you sync data to an empty directory.

## **SMB limitation for hidden directories and files**


An SMB limitation affects hidden directories and files when syncing data between SMB servers. If any of the directories or files on the source SMB server were hidden through Windows, the hidden attribute isn't copied to the target SMB server.

# Concepts

## Cloud Sync overview

The NetApp Cloud Sync service offers a simple, secure, and automated way to migrate your data to any target, in the cloud or on your premises. Whether it's a file-based NAS dataset (NFS or SMB), Amazon Simple Storage Service (S3) object format, a NetApp StorageGRID® appliance, or any other cloud provider object store, Cloud Sync can convert and move it for you.

Watch the following video for an overview of Cloud Sync:




For more details about the value that Cloud Sync provides, [read the datasheet](#).

## How Cloud Sync works

Cloud Sync is a software-as-a-service (SaaS) platform that consists of a data broker, a cloud-based portal, and a source and target.

The following image shows the relationship between Cloud Sync components:



The NetApp data broker software syncs data from a source to a target (this is called a *sync relationship*). You can run the data broker in AWS, Azure, Google Cloud Platform, or on your premises. The data broker needs an outbound internet connection over port 443 so it can communicate with the Cloud Sync service and contact a few other services and repositories. [View the list of endpoints](#).

After the initial copy, the service syncs any changed data based on the schedule that you set.

## How Cloud Sync licenses work

There are two ways to pay for sync relationships after your 14-day free trial ends. The first option is to subscribe from AWS or Azure, which enables you to pay-as-you-go or to pay annually. The second option is to purchase licenses directly from NetApp.

### AWS or Azure subscriptions

Subscribing to the Cloud Sync service from AWS or Azure enables you to pay at an hourly rate, or to pay annually. You can subscribe through either AWS or Azure, depending on where you want to be billed.

### Pay-as-you-go subscriptions

With a pay-as-you-go subscription, the Cloud Sync service charges hourly based on the number of sync relationships that you create. For pricing details, go to the [Cloud Sync service page](#).

### Annual subscriptions

An annual subscription provides a license for 20 sync relationships that you pay for up front.

If you go above 20 sync relationships and you've subscribed through Azure, you pay for the additional relationships by the hour.

- [View pay-as-you-go pricing in AWS](#)
- [View annual pricing in AWS](#)
- [View pricing in Azure](#)
- [Learn how to subscribe through AWS or Azure](#)

## Licenses from NetApp

Another way to pay for sync relationships up front is by purchasing licenses directly from NetApp. Each license enables you to create up to 20 sync relationships.

You can use these licenses with an AWS or Azure subscription. For example, if you have 25 sync relationships, you can pay for the first 20 sync relationships using a license and then pay-as-you-go from AWS or Azure with the remaining 5 sync relationships.

[Learn how to purchase licenses and add them to Cloud Sync.](#)

### License terms

Customers who purchase a Bring Your Own License (BYOL) to the Cloud Sync service should be aware of limitations associated with the license entitlement.

- Customers are entitled to leverage the BYOL license for a term not to exceed one year from the date of delivery.
- Customers are entitled to leverage the BYOL license to establish and not to exceed a total of 20 individual connections between a source and a target (each a “sync relationship”).
- A customer’s entitlement expires at the conclusion of the one-year license term, irrespective as to whether Customer has reached the 20 sync relationship limitation.
- In the event the Customer chooses to renew its license, unused sync relationships associated from the previous license grant DO NOT roll over to the license renewal.

## Accounts

Each Cloud Central user is associated with one or more Cloud Central accounts. An account enables multi-tenancy: multiple users can manage the sync relationships in a single account.

For example, two users might be associated with the same Cloud Central account. Both of those users can see the same sync relationships and data brokers that are created in that account.

But if those two users are associated with *separate* Cloud Central accounts, then the users would only see the sync relationships and data brokers in the account that they are associated with.

If a user is associated with multiple Cloud Central accounts, you can change to a different account at any time from the User Settings menu in Cloud Sync.

NetApp User Settings interface. The header is dark blue with a headset icon and a user icon. Below the header, there's a 'User Settings' section with a user icon and a 'Logout' button. The main content area has fields for 'Name' (with 'Ben' entered), 'Email', and 'Company' (with 'NetApp' entered). At the bottom, there's a red-bordered box containing 'MyAccount' and 'Account' on the left, and a blue 'Switch Account' link on the right. Below this box, a message says 'To edit user info or password go to NetApp Cloud Central'.

If you want to associate a user to a specific account, you can use Cloud Central APIs, Cloud Manager's user interface, or contact us for help using the in-product chat.

[NetApp Cloud Central Services API](#)

## Data privacy

NetApp doesn't have access to any credentials that you provide while using the Cloud Sync service. The credentials are stored directly on the data broker machine, which resides in your network.

Depending on the configuration that you choose, Cloud Sync might prompt you for credentials when you create a new relationship. For example, when setting up a relationship that includes an SMB server, or when deploying the data broker in AWS.

These credentials are always saved directly to the data broker itself. The data broker resides on a machine in your network, whether it's on premises or in your cloud account. The credentials are never made available to

NetApp.

The credentials are locally encrypted on the data broker machine using HashiCorp Vault.

# Get started

## Quick start for Cloud Sync

Getting started with the Cloud Sync service includes a few steps.



### Prepare your source and target

Verify that your source and target are supported and setup. The most important requirement is to verify connectivity between the data broker and the source and target locations. [Learn more](#).



### Prepare a location for the NetApp data broker

The NetApp data broker software syncs data from a source to a target (this is called a *sync relationship*). You can run the data broker in AWS, Azure, Google Cloud Platform, or on your premises. The data broker needs an outbound internet connection over port 443 so it can communicate with the Cloud Sync service and contact a few other services and repositories. [View the list of endpoints](#).

Cloud Sync guides you through the installation process when you create a sync relationship, at which point you can deploy the data broker in the cloud or download an install script for your own Linux host.

- [Review AWS installation](#)
- [Review Azure installation](#)
- [Review GCP installation](#)
- [Review Linux host installation](#)



### Create your first sync relationship

Start your free trial from [NetApp Cloud Central](#), drag and drop your selections for the source and target, and follow the prompts to complete the setup. [Learn more](#).



### Pay for your sync relationships after your free trial ends

Subscribe from AWS or Azure to pay-as-you-go or to pay annually. Or purchase licenses directly from NetApp. Just go to the License Settings page in Cloud Sync to set it up. [Learn more](#).

## Networking overview

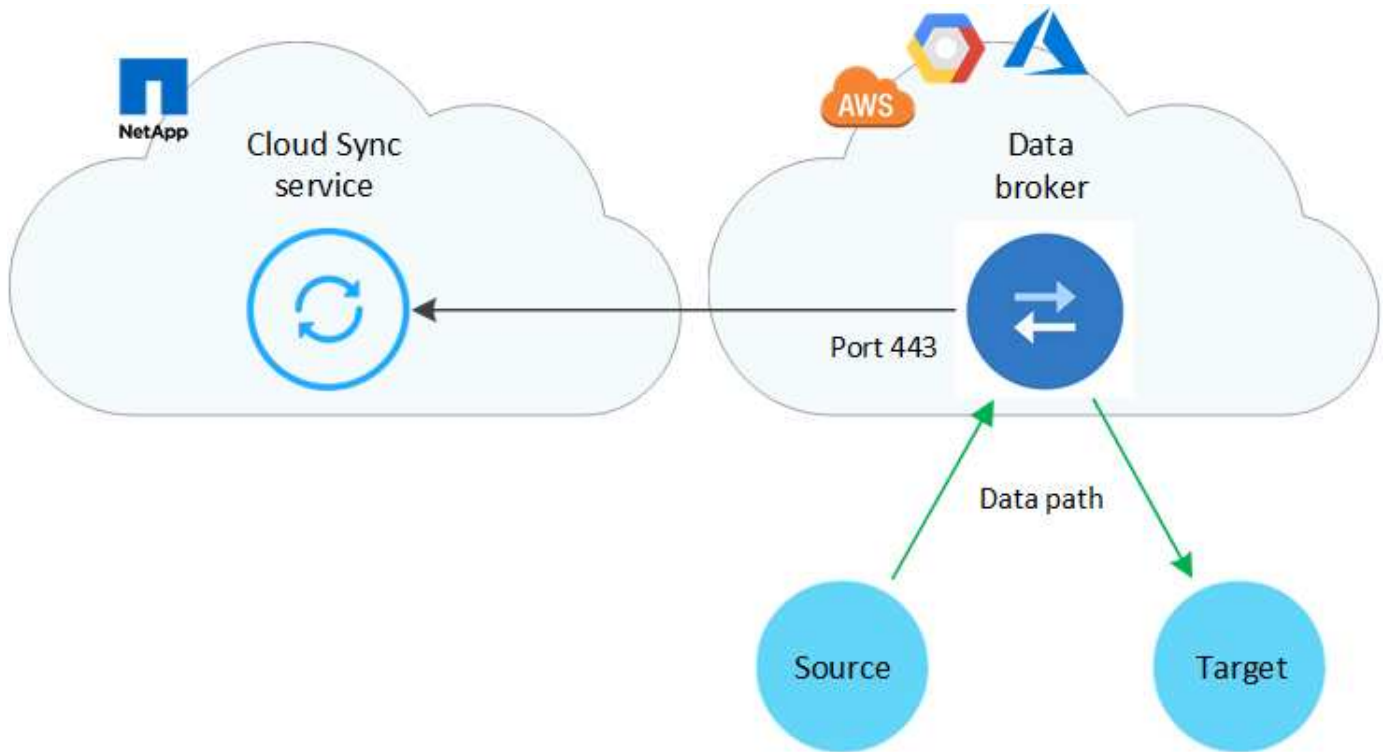
Networking for Cloud Sync includes connectivity between the data broker and the source and target locations, and an outbound internet connection from the data broker over port 443.

## Data broker in the cloud

The following image shows the data broker running in the cloud, in either AWS, GCP, or Azure. The source and target can be in any location, as long as there's a connection to the data broker. For example, you might have a VPN connection from your data center to your cloud provider.




When Cloud Sync deploys the data broker in AWS, Azure, or GCP, it creates a security group that enables the required outbound communication.



## Data broker on your premises

The following image shows the data broker running on-prem, in a data center. Again, the source and target can be in any location, as long as there's a connection to the data broker.





#### Related link

[Endpoints that the data broker contacts](#)

## Preparing the source and target

Prepare to sync data by verifying that your source and target are supported and setup.

### Supported sync relationships

Cloud Sync enables you to sync data from a source to a target (this is called a *sync relationship*). You should understand the supported relationships before you get started.

Source location	Supported target locations
AWS EFS	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (NFS)</li> <li>• Cloud Volumes ONTAP (NFS)</li> <li>• Cloud Volumes Service (NFS)</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• StorageGRID</li> </ul>
AWS S3	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files</li> <li>• Cloud Volumes ONTAP</li> <li>• Cloud Volumes Service</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>

Source location	Supported target locations
Azure Blob	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files</li> <li>• Cloud Volumes ONTAP</li> <li>• Cloud Volumes Service</li> <li>• Google Cloud Storage</li> <li>• IBM Cloud Object Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>
Azure NetApp Files (NFS)	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (NFS)</li> <li>• Cloud Volumes ONTAP (NFS)</li> <li>• Cloud Volumes Service (NFS)</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• StorageGRID</li> </ul>
Azure NetApp Files (SMB)	<ul style="list-style-type: none"> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (SMB)</li> <li>• Cloud Volumes ONTAP (SMB)</li> <li>• Cloud Volumes Service (SMB)</li> <li>• Google Cloud Storage</li> <li>• IBM Cloud Object Storage</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>

Source location	Supported target locations
Cloud Volumes ONTAP (NFS)	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (NFS)</li> <li>• Cloud Volumes ONTAP (NFS)</li> <li>• Cloud Volumes Service (NFS)</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• StorageGRID</li> </ul>
Cloud Volumes ONTAP (SMB)	<ul style="list-style-type: none"> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (SMB)</li> <li>• Cloud Volumes ONTAP (SMB)</li> <li>• Cloud Volumes Service (SMB)</li> <li>• Google Cloud Storage</li> <li>• IBM Cloud Object Storage</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>
Cloud Volumes Service (NFS)	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (NFS)</li> <li>• Cloud Volumes ONTAP (NFS)</li> <li>• Cloud Volumes Service (NFS)</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• StorageGRID</li> </ul>

Source location	Supported target locations
Cloud Volumes Service (SMB)	<ul style="list-style-type: none"> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (SMB)</li> <li>• Cloud Volumes ONTAP (SMB)</li> <li>• Cloud Volumes Service (SMB)</li> <li>• Google Cloud Storage</li> <li>• IBM Cloud Object Storage</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>
Google Cloud Storage	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files</li> <li>• Cloud Volumes ONTAP</li> <li>• Cloud Volumes Service</li> <li>• Google Cloud Storage</li> <li>• IBM Cloud Object Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>
IBM Cloud Object Storage	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files</li> <li>• Cloud Volumes ONTAP</li> <li>• Cloud Volumes Service</li> <li>• Google Cloud Storage</li> <li>• IBM Cloud Object Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>

Source location	Supported target locations
NFS server	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (NFS)</li> <li>• Cloud Volumes ONTAP (NFS)</li> <li>• Cloud Volumes Service (NFS)</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• StorageGRID</li> </ul>
On-prem ONTAP cluster (NFS)	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (NFS)</li> <li>• Cloud Volumes ONTAP (NFS)</li> <li>• Cloud Volumes Service (NFS)</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• StorageGRID</li> </ul>
On-prem ONTAP cluster (SMB)	<ul style="list-style-type: none"> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (SMB)</li> <li>• Cloud Volumes ONTAP (SMB)</li> <li>• Cloud Volumes Service (SMB)</li> <li>• Google Cloud Storage</li> <li>• IBM Cloud Object Storage</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>

Source location	Supported target locations
ONTAP S3 Storage	<ul style="list-style-type: none"> <li>• SMB server</li> <li>• StorageGRID</li> <li>• ONTAP S3 Storage</li> </ul>
SFTP <sup>1</sup>	S3
SMB server	<ul style="list-style-type: none"> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (SMB)</li> <li>• Cloud Volumes ONTAP (NFS)</li> <li>• Cloud Volumes Service (NFS)</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• On-premises ONTAP cluster</li> <li>• ONTAP S3 Storage</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>
StorageGRID	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files</li> <li>• Cloud Volumes ONTAP</li> <li>• Cloud Volumes Service</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• ONTAP S3 Storage</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>

Notes:

1. Cloud Sync supports sync relationships from SFTP to S3 by using the API only.
2. You can choose a specific Azure Blob storage tier when a Blob container is the target:
  - Hot storage
  - Cool storage

3. You can choose a specific S3 storage class when AWS S3 is the target:
  - Standard (this is the default class)
  - Intelligent-Tiering
  - Standard-Infrequent Access
  - One Zone-Infrequent Access
  - Glacier
  - Glacier Deep Archive

## Networking requirements

- The source and target must have a network connection to the data broker.

For example, if an NFS server is in your data center and the data broker is in AWS, then you need a network connection (VPN or Direct Connect) from your network to the VPC.

- NetApp recommends configuring the source, target, and data broker to use a Network Time Protocol (NTP) service. The time difference between the three components should not exceed 5 minutes.

## Azure NetApp Files requirement

Use the Premium or Ultra service level when you sync data to or from Azure NetApp Files. You might experience failures and performance issues if the disk service level is Standard.



Consult a solutions architect if you need help determining the right service level. The volume size and volume tier determines the throughput that you can get.

[Learn more about Azure NetApp Files service levels and throughput.](#)

## NFS server requirements

- The NFS server can be a NetApp system or a non-NetApp system. For example, Cloud Volumes ONTAP or an AFF cluster.
- The file server must allow the data broker host to access the exports.
- NFS versions 3, 4.0, 4.1, and 4.2 are supported.

The desired version must be enabled on the server.

- If you want to sync NFS data from an ONTAP system, ensure that access to the NFS export list for an SVM is enabled (`vserver nfs modify -vserver svm_name -showmount enabled`).



The default setting for showmount is *enabled* starting with ONTAP 9.2.

## SMB server requirements

- The SMB server can be a NetApp system or a non-NetApp system. For example, Cloud Volumes ONTAP or an AFF cluster.
- The file server must allow the data broker host to access the exports.



- SMB versions 1.0, 2.0, 2.1, 3.0 and 3.11 are supported.
- Grant the "Administrators" group with "Full Control" permissions to the source and target folders.

If you don't grant this permission, then the data broker might not have sufficient permissions to get the ACLs on a file or directory. If this occurs, you'll receive the following error: "getxattr error 95"

## AWS S3 bucket requirements

Make sure that your AWS S3 bucket meets the following requirements.

### Supported data broker locations for AWS S3

Sync relationships that include S3 storage require a data broker deployed in AWS or on your premises. In either case, Cloud Sync prompts you to associate the data broker with an AWS account during installation.

- [Learn how to deploy the AWS data broker](#)
- [Learn how to install the data broker on a Linux host](#)

### Supported AWS regions

All regions are supported except for the China and GovCloud (US) regions.

### Permissions required for S3 buckets in other AWS accounts

When setting up a sync relationship, you can specify an S3 bucket that resides in an AWS account that isn't associated with the data broker.

[The permissions included in this JSON file](#) must be applied to that S3 bucket so the data broker can access it. These permissions enable the data broker to copy data to and from the bucket and to list the objects in the bucket.


Note the following about the permissions included in the JSON file:

1. *<BucketName>* is the name of the bucket that resides in the AWS account that isn't associated with the data broker.
2. *<RoleARN>* should be replaced with one of the following:
  - If the data broker was manually installed on a Linux host, *RoleARN* should be the ARN of the AWS user for which you provided AWS credentials when deploying the data broker.
  - If the data broker was deployed in AWS using the CloudFormation template, *RoleARN* should be the ARN of the IAM role created by the template.

You can find the Role ARN by going to the EC2 console, selecting the data broker instance, and clicking the IAM role from the Description tab. You should then see the Summary page in the IAM console that contains the Role ARN.

## Summary

Delete role

**Role ARN** `arn:aws:iam::542991749851:role/tanyaBroker0304-DataBrokerIamRole-1VMHXXMW3AQ05` 

Role description [Edit](#)

## Azure Blob storage requirements

Make sure that your Azure Blob storage meets the following requirements.

### Supported data broker locations for Azure Blob

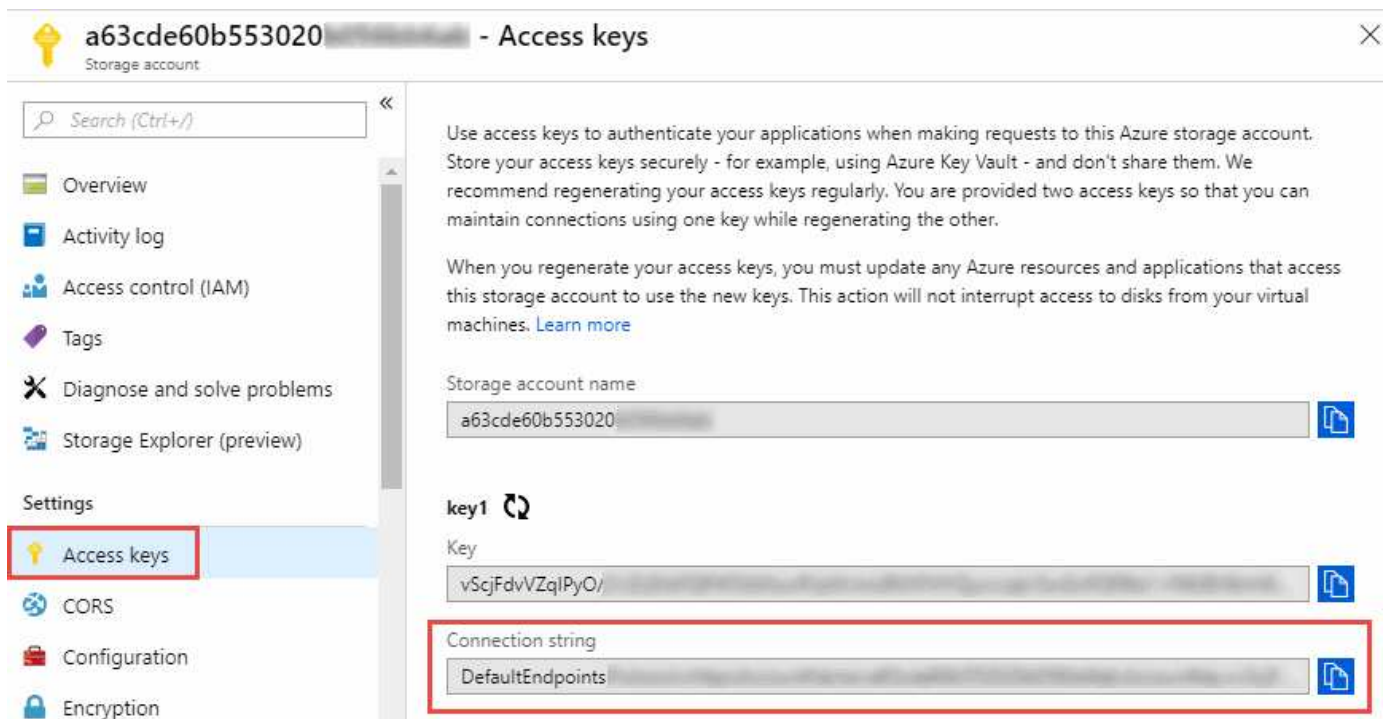
The data broker can reside in any location when a sync relationship includes Azure Blob storage.

### Supported Azure regions

All regions are supported except for the China, US Gov, and US DoD regions.

### Connection string required for relationships that include Azure Blob and NFS/SMB

When creating a sync relationship between an Azure Blob container and an NFS or SMB server, you need to provide Cloud Sync with the storage account connection string:



The screenshot shows the 'Access keys' page for an Azure storage account. The left sidebar contains navigation links: Overview, Activity log, Access control (IAM), Tags, Diagnose and solve problems, Storage Explorer (preview), Settings, Access keys (highlighted with a red box), CORS, Configuration, and Encryption. The main content area has a title bar 'a63cde60b553020 - Access keys' and a close button. Below the title bar is a search bar and a list of navigation links. The main content area contains instructions on using access keys, a 'Storage account name' field with the value 'a63cde60b553020', and a 'key1' section with a 'Key' field containing 'vScjFdvVZqIPyO/'. Below the key field is a 'Connection string' field with the value 'DefaultEndpoints', which is highlighted with a red box.

If you want to sync data between two Azure Blob containers, then the connection string must include a [shared access signature](#) (SAS). You also have the option to use a SAS when syncing between a Blob container and an NFS or SMB server.

The SAS must allow access to the Blob service and all resource types (Service, Container, and Object). The SAS must also include the following permissions:

- For the source Blob container: Read and List
- For the target Blob container: Read, Write, List, Add, and Create

Search (Ctrl+/)

- Overview
- Activity log
- Access control (IAM)
- Tags
- Diagnose and solve problems
- Storage Explorer (preview)
- Settings
  - Access keys
  - CORS
  - Configuration
  - Encryption
  - Shared access signature
  - Firewalls and virtual networks
  - Advanced Threat Protection (pr...
  - Properties
  - Locks

Allowed services ⓘ

☒ Blob
☐ File
☐ Queue
☐ Table

Allowed resource types ⓘ

☒ Service
☒ Container
☒ Object

Allowed permissions ⓘ

☒ Read
☒ Write
☒ Delete
☒ List
☒ Add
☒ Create
☐ Update
☐ Process

Start and expiry date/time ⓘ

Start

2018-10-23
10:07:32 AM

End

2019-10-23
6:07:32 PM

(UTC-04:00) --- Current Time Zone ---

Allowed IP addresses ⓘ

for example, 168.1.5.65 or 168.1.5.65-168.1.5.70

Allowed protocols ⓘ

☒ HTTPS only
☐ HTTPS and HTTP

Signing key ⓘ

key1

Generate SAS and connection string

## Google Cloud Storage bucket requirements

Make sure that your Google Cloud Storage bucket meets the following requirements.

### Supported data broker locations for Google Cloud Storage

Sync relationships that include Google Cloud Storage require a data broker deployed in GCP or on your premises. Cloud Sync guides you through the data broker installation process when you create a sync relationship.

- [Learn how to deploy the GCP data broker](#)
- [Learn how to install the data broker on a Linux host](#)

### Supported GCP regions

All regions are supported.

## ONTAP requirements

If the sync relationship includes Cloud Volumes ONTAP or an on-prem ONTAP cluster and you selected NFSv4 or later, then you'll need to enable NFSv4 ACLs on the ONTAP system. This is required to copy the ACLs.

## Permissions for a SnapMirror destination

If the source for a sync relationship is a SnapMirror destination (which is read-only), "read/list" permissions are sufficient to sync data from the source to a target.

## ONTAP S3 Storage requirements

When you set up a sync relationship that includes [ONTAP S3 Storage](#), you'll need to provide the following:

- The IP address of the LIF that's connected to ONTAP S3
- The access key and secret key that ONTAP is configured to use

## Endpoints that are required for Cloud Sync

The NetApp data broker requires outbound internet access over port 443 to communicate with the Cloud Sync service and to contact a few other services and repositories. Your local web browser also requires access to endpoints for certain actions. If you need to limit outbound connectivity, refer to the following list of endpoints when configuring your firewall for outbound traffic.

### Data broker endpoints

The data broker contacts the following endpoints:

Endpoints	Purpose
olcentgbl.trafficmanager.net:443	To contact a repository for updating CentOS packages for the data broker host. This endpoint is contacted only if you manually install the data broker on a CentOS host.
rpm.nodesource.com:443 registry.npmjs.org:443 nodejs.org:443	To contact repositories for updating Node.js, npm, and other 3rd party packages used in development.
tgz.pm2.io:443	To access a repository for updating PM2, which is a 3rd party package used to monitor Cloud Sync.
sqs.us-east-1.amazonaws.com:443 kinesis.us-east-1.amazonaws.com:443	To contact the AWS services that Cloud Sync uses for operations (queuing files, registering actions, and delivering updates to the data broker).
s3.region.amazonaws.com:443  For example: s3.us-east-2.amazonaws.com:443 <a href="#">See AWS documentation for a list of S3 endpoints</a>	To contact Amazon S3 when a sync relationship includes an S3 bucket.
cf.cloudsync.netapp.com:443 repo.cloudsync.netapp.com:443	To contact the Cloud Sync service.
support.netapp.com:443	To contact NetApp support when using a BYOL license for sync relationships.

Endpoints	Purpose
fedoraproject.org:443	To install 7z on the data broker virtual machine during installation and updates. 7z is needed to send AutoSupport messages to NetApp technical support.
sts.amazonaws.com:443	To verify AWS credentials when the data broker is deployed in AWS or when it's deployed on your premises and AWS credentials are provided. The data broker contacts this endpoint during deployment, when it's updated, and when it's restarted.

## Web browser endpoints

Your web browser needs access to the following endpoint to download logs for troubleshooting purposes:

logs.cloudsync.netapp.com:443

# Install the data broker

## Installing the data broker in AWS

When you create a new data broker, choose the AWS Data Broker option to deploy the data broker software on a new EC2 instance in a VPC. Cloud Sync guides you through the installation process, but the requirements and steps are repeated on this page to help you prepare for installation.

You also have the option to install the data broker on an existing Linux host in the cloud or on your premises. [Learn more.](#)

## Supported AWS regions

All regions are supported except for the China and GovCloud (US) regions.

## Networking requirements

- The data broker needs an outbound internet connection so it can poll the Cloud Sync service for tasks over port 443.

When Cloud Sync deploys the data broker in AWS, it creates a security group that enables the required outbound communication. Note that you can configure the data broker to use a proxy server during the installation process.

If you need to limit outbound connectivity, see [the list of endpoints that the data broker contacts](#).

- NetApp recommends configuring the source, target, and data broker to use a Network Time Protocol (NTP) service. The time difference between the three components should not exceed 5 minutes.

## Permissions required to deploy the data broker in AWS

The AWS user account that you use to deploy the data broker must have the permissions included in [this NetApp-provided policy](#).

## Requirements to use your own IAM role with the AWS data broker

When Cloud Sync deploys the data broker, it creates an IAM role for the data broker instance. You can deploy the data broker using your own IAM role, if you prefer. You might use this option if your organization has strict security policies.

The IAM role must meet the following requirements:

- The EC2 service must be allowed to assume the IAM role as a trusted entity.
- [The permissions defined in this JSON file](#) must be attached to the IAM role so the data broker can function properly.


Follow the steps below to specify the IAM role when deploying the data broker.

## Installing the data broker

You can install a data broker in AWS when you create a sync relationship.

### Steps

1. Click **Create New Sync Relationship**.
2. On the **Define Sync Relationship** page, choose a source and target and click **Continue**.  
  
Complete the steps until you reach the **Data Broker** page.
3. On the **Data Broker** page, click **Create Data Broker** and then select **Amazon Web Services**.

If you already have a data broker, you'll need to click the  icon first.

4. Enter a name for the data broker and click **Continue**.
5. Enter an AWS access key so Cloud Sync can create the data broker in AWS on your behalf.

The keys aren't saved or used for any other purposes.

If you'd rather not provide access keys, click the link at the bottom of the page to use a CloudFormation template instead. When you use this option, you don't need to provide credentials because you are logging in directly to AWS.

The following video shows how to launch the data broker instance using a CloudFormation template:

► [https://docs.netapp.com/us-en/cloudsync//media/video\\_cloud\\_sync.mp4](https://docs.netapp.com/us-en/cloudsync//media/video_cloud_sync.mp4) (video)

6. If you entered an AWS access key, select a location for the instance, select a key pair, choose whether to enable a public IP address, and then select an existing IAM role, or leave the field blank so Cloud Sync creates the role for you.

If you choose your own IAM role, [you'll need to provide the required permissions](#).

### Basic Settings

**Location**

Region  

US West | Oregon ▼

VPC  

vpc-3c46c059 - 10.60.21.0/25 ▼

Subnet  

10.60.21.0/25 ▼

**Connectivity**

Key Pair  


newKey ▼

**Enable Public IP?**  
☒ Enable ☐ Disable

IAM Role (optional) ?

7. Specify a proxy configuration, if a proxy is required for internet access in the VPC.
8. After the data broker is available, click **Continue** in Cloud Sync.

The following image shows a successfully deployed instance in AWS:

Select a NetApp Data Broker			
1 NetApp Data Brokers			
 name			Active
US West (Oregon) Region	10.60.21.0/25   vpc-3c46c059 VPC	10.60.21.5 Private IP	5f5002eecf378e000a560988 Broker ID
us-west-2c Availability Zone	10.60.21.0/25   subnet-e7f526be Subnet	i-0fc5c97e2f5f22c20 Instance ID	

9. Complete the pages in the wizard to create the new sync relationship.

### Result

You have deployed a data broker in AWS and created a new sync relationship. You can use this data broker with additional sync relationships.

## Installing the data broker in Azure

When you create a new data broker, choose the Azure Data Broker option to deploy the data broker software on a new virtual machine in a VNet. Cloud Sync guides you through the installation process, but the requirements and steps are repeated on this page to help you prepare for installation.

You also have the option to install the data broker on an existing Linux host in the cloud or on your premises. [Learn more.](#)

## Supported Azure regions

All regions are supported except for the China, US Gov, and US DoD regions.

## Networking requirements

- The data broker needs an outbound internet connection so it can poll the Cloud Sync service for tasks over port 443.

When Cloud Sync deploys the data broker in Azure, it creates a security group that enables the required outbound communication.

If you need to limit outbound connectivity, see [the list of endpoints that the data broker contacts](#).

- NetApp recommends configuring the source, target, and data broker to use a Network Time Protocol (NTP) service. The time difference between the three components should not exceed 5 minutes.

## Authentication method

When you deploy the data broker, you'll need to choose an authentication method: a password or an SSH public-private key pair.

For help with creating a key pair, refer to [Azure Documentation: Create and use an SSH public-private key pair for Linux VMs in Azure](#).

## Installing the data broker


You can install a data broker in Azure when you create a sync relationship.


### Steps

1. Click **Create New Sync Relationship**.
2. On the **Define Sync Relationship** page, choose a source and target and click **Continue**.

Complete the pages until you reach the **Data Broker** page.

3. On the **Data Broker** page, click **Create Data Broker** and then select **Microsoft Azure**.

If you already have a data broker, you'll need to click the  icon first.



4. Enter a name for the data broker and click **Continue**.



5. If you're prompted, log in to your Microsoft account. If you're not prompted, click **Log in to Azure**.

The form is owned and hosted by Microsoft. Your credentials are not provided to NetApp.

6. Choose a location for the data broker and enter basic details about the virtual machine.

The screenshot shows the 'Location' and 'Virtual Machine' configuration sections of the Azure portal. The 'Location' section includes dropdowns for Subscription (OCCM Dev), Azure Region (West US 2), VNet (Vnet1), and Subnet (Subnet1). The 'Virtual Machine' section includes text inputs for VM Name (netappdatabroker), User Name (databroker), and Enter Password (masked). It also features radio buttons for Authentication Method (Password selected) and Resource Group (Generate a new group selected).

7. Specify a proxy configuration, if a proxy is required for internet access in the VNet.

8. Click **Continue** and keep the page open until the deployment is complete.

The process can take up to 7 minutes.

9. In Cloud Sync, click **Continue** once the data broker is available.

10. Complete the pages in the wizard to create the new sync relationship.

## Result

You have deployed a data broker in Azure and created a new sync relationship. You can use this data broker with additional sync relationships.

## Getting a message about needing admin consent?

If Microsoft notifies you that admin approval is required because Cloud Sync needs permission to access resources in your organization on your behalf, then you have two options:

1. Ask your AD admin to provide you with the following permission:

In Azure, go to **Admin Centers > Azure AD > Users and Groups > User Settings** and enable **Users can consent to apps accessing company data on their behalf**.

2. Ask your AD admin to consent on your behalf to **CloudSync-AzureDataBrokerCreator** using the following URL (this is the admin consent endpoint):

```
https://login.microsoftonline.com/{FILL HERE YOUR TENANT ID}/v2.0/adminconsent?client_id=8ee4ca3a-bafa-4831-97cc-5a38923cab85&redirect_uri=https://cloudsync.netapp.com&scope=https://management.azure.com/user_impersonationhttps://graph.microsoft.com/User.Read
```

As shown in the URL, our app URL is `https://cloudsync.netapp.com` and the application client ID is `8ee4ca3a-bafa-4831-97cc-5a38923cab85`.

## Installing the data broker in Google Cloud Platform

When you create a new data broker, choose the GCP Data Broker option to deploy the data broker software on a new virtual machine instance in a VPC. Cloud Sync guides you through the installation process, but the requirements and steps are repeated on this page to help you prepare for installation.

You also have the option to install the data broker on an existing Linux host in the cloud or on your premises. [Learn more](#).

### Supported GCP regions

All regions are supported.

### Networking requirements

- The data broker needs an outbound internet connection so it can poll the Cloud Sync service for tasks over port 443.

When Cloud Sync deploys the data broker in GCP, it creates a security group that enables the required outbound communication.

If you need to limit outbound connectivity, see [the list of endpoints that the data broker contacts](#).

- NetApp recommends configuring the source, target, and data broker to use a Network Time Protocol (NTP) service. The time difference between the three components should not exceed 5 minutes.

### Permissions required to deploy the data broker in GCP

Ensure that the GCP user who deploys the data broker has the following permissions:

- `compute.networks.list`
- `compute.regions.list`
- `deploymentmanager.deployments.create`
- `deploymentmanager.deployments.delete`
- `deploymentmanager.operations.get`
- `iam.serviceAccounts.list`

### Permissions required for the service account

When you deploy the data broker, you need to select a service account that has the following permissions:

- `logging.logEntries.create`
- `resourcemanager.projects.get`
- `storage.buckets.get`
- `storage.buckets.list`
- `storage.objects.*`

### Installing the data broker


You can install a data broker in GCP when you create a sync relationship.


#### Steps

1. Click **Create New Sync Relationship**.
2. On the **Define Sync Relationship** page, choose a source and target and click **Continue**.

Complete the steps until you reach the **Data Broker** page.

3. On the **Data Broker** page, click **Create Data Broker** and then select **Google Cloud Platform**.

If you already have a data broker, you'll need to click the  icon first.



4. Enter a name for the data broker and click **Continue**.
5. If you're prompted, log in with your Google account.

The form is owned and hosted by Google. Your credentials are not provided to NetApp.

6. Select a project and service account and then choose a location for the data broker.

### Basic Settings

<b>Project</b>	<b>Location</b>
Project OCCM-Dev	Region us-west1
Service Account test	Zone us-west1-a
Select a Service Account that includes <a href="#">these permissions</a>	VPC default
	Subnet default

7. Specify a proxy configuration, if a proxy is required for internet access in the VPC.

If a proxy is required for internet access, then the proxy must be in Google Cloud and use the same service account as the data broker.

8. Once the data broker is available, click **Continue** in Cloud Sync.

The instance takes approximately 5 to 10 minutes to deploy. You can monitor the progress from the Cloud Sync service, which automatically refreshes when the instance is available.

9. Complete the pages in the wizard to create the new sync relationship.

## Result

You've deployed a data broker in GCP and created a new sync relationship. You can use this data broker with additional sync relationships.

## Installing the data broker on a Linux host

When you create a new data broker, choose the On-Prem Data Broker option to install the data broker software on an on-premises Linux host, or on an existing Linux host in the cloud. Cloud Sync guides you through the installation process, but the requirements and steps are repeated on this page to help you prepare for installation.

### Linux host requirements

- **Operating system:**
  - CentOS 7.0, 7.7, and 8.0
  - Red Hat Enterprise Linux 7.7 and 8.0

- Ubuntu Server 20.04 LTS
- SUSE Linux Enterprise Server 15 SP1

The command `yum update all` must be run on the host before you install the data broker.

A Red Hat Enterprise Linux system must be registered with Red Hat Subscription Management. If it is not registered, the system cannot access repositories to update required 3rd party software during installation.

- **RAM:** 16 GB
- **CPU:** 4 cores
- **Free disk space:** 10 GB
- **SELinux:** We recommend that you disable [SELinux](#) on the host.

SELinux enforces a policy that blocks data broker software updates and can block the data broker from contacting endpoints required for normal operation.

- **OpenSSL:** OpenSSL must be installed on the Linux host.

## Networking requirements

- The Linux host must have a connection to the source and target.
- The file server must allow the Linux host to access the exports.
- Port 443 must be open on the Linux host for outbound traffic to AWS (the data broker constantly communicates with the Amazon SQS service).
- NetApp recommends configuring the source, target, and data broker to use a Network Time Protocol (NTP) service. The time difference between the three components should not exceed 5 minutes.

## Enabling access to AWS

If you plan to use the data broker with a sync relationship that includes an S3 bucket, then you should prepare the Linux host for AWS access. When you install the data broker, you'll need to provide AWS keys for an AWS user that has programmatic access and specific permissions.

### Steps

1. Create an IAM policy using [this NetApp-provided policy](#). [View AWS instructions](#).
2. Create an IAM user that has programmatic access. [View AWS instructions](#).

Be sure to copy the AWS keys because you need to specify them when you install the data broker software.

## Enabling access to Google Cloud

If you plan to use the data broker with a sync relationship that includes a Google Cloud Storage bucket, then you should prepare the Linux host for GCP access. When you install the data broker, you'll need to provide a key for a service account that has specific permissions.

### Steps

1. Create a GCP service account that has Storage Admin permissions, if you don't already have one.

2. Create a service account key saved in JSON format. [View GCP instructions](#).

The file should contain at least the following properties: "project\_id", "private\_key", and "client\_email"



When you create a key, the file gets generated and downloaded to your machine.

3. Save the JSON file to the Linux host.

## Enabling access to Microsoft Azure

Access to Azure is defined per relationship by providing a storage account and a connection string in the Sync Relationship wizard.

## Installing the data broker


You can install a data broker on a Linux host when you create a sync relationship.

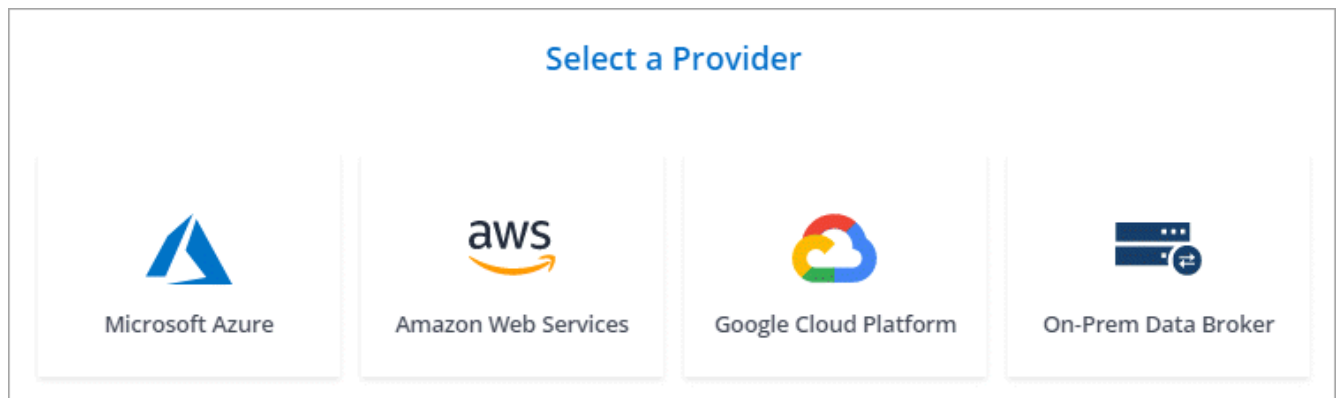
### Steps

1. Click **Create New Sync Relationship**.
2. On the **Define Sync Relationship** page, choose a source and target and click **Continue**.

Complete the steps until you reach the **Data Broker** page.

3. On the **Data Broker** page, click **Create Data Broker** and then select **On-Prem Data Broker**.

If you already have a data broker, you'll need to click the  icon first.



Even though the option is labeled **On-Prem Data Broker**, it applies to a Linux host on your premises or in the cloud.

4. Enter a name for the data broker and click **Continue**.

The instructions page loads shortly. You'll need to follow these instructions—they include a unique link to download the installer.

5. On the instructions page:
  - a. Select whether to enable access to **AWS**, **Google Cloud**, or both.
  - b. Select an installation option: **No proxy**, **Use proxy server**, or **Use proxy server with authentication**.

- c. Use the commands to download and install the data broker.

The following steps provide details about each possible installation option. Follow the instructions page to get the exact command based on your installation option.

- d. Download the installer:

- No proxy:

```
curl <URI> -o data_broker_installer.sh
```

- Use proxy server:

```
curl <URI> -o data_broker_installer.sh -x <proxy_host>:<proxy_port>
```

- Use proxy server with authentication:

```
curl <URI> -o data_broker_installer.sh -x  
<proxy_username>:<proxy_password>@<proxy_host>:<proxy_port>
```

### URI

Cloud Sync displays the URI of the installation file on the instructions page, which loads when you follow the prompts to deploy the On-Prem Data Broker. That URI isn't repeated here because the link is generated dynamically and can be used only once. [Follow these steps to obtain the URI from Cloud Sync.](#)

- e. Switch to superuser, make the installer executable and install the software:



Each command listed below includes parameters for AWS access and GCP access. Follow the instructions page to get the exact command based on your installation option.

- No proxy configuration:

```
sudo -s  
chmod +x data_broker_installer.sh  
./data_broker_installer.sh -a <aws_access_key> -s <aws_secret_key> -g  
<absolute_path_to_the_json_file>
```

- Proxy configuration:

```
sudo -s  
chmod +x data_broker_installer.sh  
./data_broker_installer.sh -a <aws_access_key> -s <aws_secret_key> -g  
<absolute_path_to_the_json_file> -h <proxy_host> -p <proxy_port>
```

- Proxy configuration with authentication:

```
sudo -s  
chmod +x data_broker_installer.sh  
./data_broker_installer.sh -a <aws_access_key> -s <aws_secret_key> -g  
<absolute_path_to_the_json_file> -h <proxy_host> -p <proxy_port> -u  
<proxy_username> -w <proxy_password>
```

### AWS keys

These are the keys for the user that you should have prepared [following these steps](#). The AWS keys are stored on the data broker, which runs in your on-premises or cloud network. NetApp doesn't use the keys outside of the data broker.

### JSON file

This is the JSON file that contains a service account key that you should have prepared [following these steps](#).

6. Once the data broker is available, click **Continue** in Cloud Sync.
7. Complete the pages in the wizard to create the new sync relationship.

## Creating a sync relationship

When you create a sync relationship, the Cloud Sync service copies files from the source to the target. After the initial copy, the service syncs any changed data every 24 hours.

The steps below provide an example that shows how to set up a sync relationship from an NFS server to an S3 bucket.

### Steps

1. Go to [NetApp Cloud Central](#).
2. Sign up or log in and then start a free trial of Cloud Sync.
3. After you log in, review details about using the service after the free trial ends, and then click **OK**.
4. On the **Select Source & Target** page, choose a source and target.

The following steps provide an example of how to create a sync relationship from an NFS server to an S3 bucket.



5. Review the details about how the service works and then click **Continue**.
6. On the **NFS Server** page, enter the IP address or fully qualified domain name of the NFS server that you want to sync to AWS.
7. On the **Data Broker** page, follow the prompts to create a data broker virtual machine in AWS, Azure, or Google Cloud Platform, or to install the data broker software on an existing Linux host.

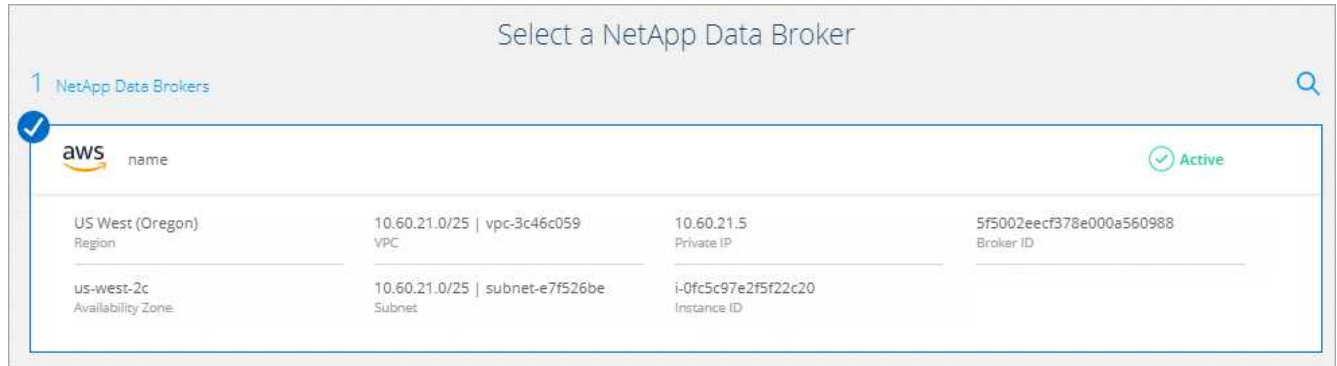
For more details, refer to the following pages:



- [Installing the data broker in AWS](#)
- [Installing the data broker in Azure](#)
- [Installing the data broker in GCP](#)
- [Installing the data broker on a Linux host](#)

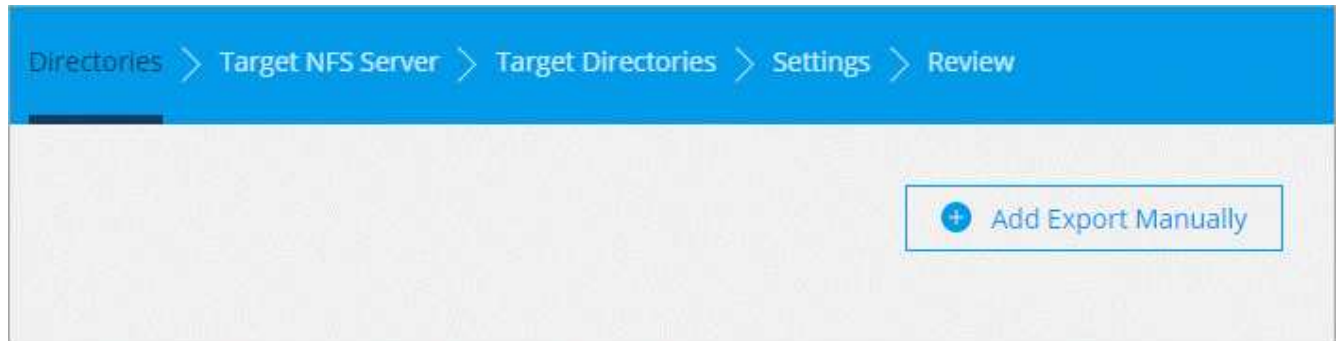
8. After you install the data broker, click **Continue**.

The following image shows a successfully deployed data broker in AWS:



9. On the **Directories** page, select a top-level directory or subdirectory.

If Cloud Sync is unable to retrieve the exports, click **Add Export Manually** and enter the name of an NFS export.



If you want to sync more than one directory on the NFS server, then you must create additional sync relationships after you are done.

10. On the **AWS S3 Bucket** page, select a bucket:

- Drill down to select an existing folder within the bucket or to select a new folder that you create inside the bucket.
- Click **Add to the list** to select an S3 bucket that is not associated with your AWS account. [Specific permissions must be applied to the S3 bucket.](#)

11. On the **Bucket Setup** page, set up the bucket:

- Choose whether to enable S3 bucket encryption and then select an AWS KMS key, enter the ARN of a KMS key, or select AES-256 encryption.
- Select an S3 storage class. [View the supported storage classes.](#)

12. On the **Settings** page, define how source files and folders are synced and maintained in the target location:

### Schedule

Choose a recurring schedule for future syncs or turn off the sync schedule. You can schedule a relationship to sync data as often as every 1 minute.

### Retries

Define the number of times that Cloud Sync should retry to sync a file before skipping it.

### Recently Modified Files

Choose to exclude files that were recently modified prior to the scheduled sync.

### Delete Files on Source

Choose to delete files from the source location after Cloud Sync copies the files to the target location. This option includes the risk of data loss because the source files are deleted after they're copied.

If you enable this option, you also need to change a parameter in the `local.json` file on the data broker. Open the file and change the parameter named `workers.transferrer.delete-on-source` to **true**.

### Delete Files on Target

Choose to delete files from the target location, if they were deleted from the source. The default is to never deletes files from the target location.

### Object tagging

When AWS S3 is the target in a sync relationship, Cloud Sync tags S3 objects with metadata that's relevant to the sync operation. You can disable tagging of S3 objects, if it's not desired in your environment. There's no impact to Cloud Sync if you disable tagging—Cloud Sync just stores the sync metadata in a different way.

## File Types

Define the file types to include in each sync: files, directories, and symbolic links.

## Exclude File Extensions

Specify file extensions to exclude from the sync by typing the file extension and pressing **Enter**. For example, type *log* or *.log* to exclude \*.log files. A separator isn't required for multiple extensions. The following video provides a short demo:

► [https://docs.netapp.com/us-en/cloudsync//media/video\\_file\\_extensions.mp4](https://docs.netapp.com/us-en/cloudsync//media/video_file_extensions.mp4) (video)

## File Size

Choose to sync all files regardless of their size or just files that are in a specific size range.

## Date Modified

Choose all files regardless of their last modified date, files modified after a specific date, before a specific date, or between a time range.

13. On the **Relationship Tags** page, enter up to 9 relationship tags and then click **Continue**.

The Cloud Sync service assigns the tags to each object that it syncs to the S3 bucket.

14. Review the details of the sync relationship and then click **Create Relationship**.
15. After the Cloud Sync service successfully creates the relationship, click **View in Dashboard** to view details about the data sync relationship.

## Result

Cloud Sync starts syncing data between the source and target.

# Paying for sync relationships after your free trial ends

There are two ways to pay for sync relationships after your 14-day free trial ends. The first option is to subscribe from AWS or Azure to pay-as-you-go or to pay annually. The second option is to purchase licenses directly from NetApp.

You can use licenses from NetApp with an AWS or Azure subscription. For example, if you have 25 sync relationships, you can pay for the first 20 sync relationships using a license and then pay-as-you-go from AWS or Azure with the remaining 5 sync relationships.

[Learn more about how licenses work.](#)

## What if I don't immediately pay after my free trial ends?

You won't be able to create any additional relationships. Existing relationships are not deleted, but you cannot make any changes to them until you subscribe or enter a license.

## Subscribing from AWS

AWS enables you to pay-as-you-go or to pay annually.

## Steps to pay-as-you-go

1. Go to the [License Settings](#) page.
2. Select **AWS**
3. Click **Subscribe** and then click **Continue**.
4. Subscribe from the AWS Marketplace, and then log back in to the Cloud Sync service to complete the registration.

The following video shows the process:

► [https://docs.netapp.com/us-en/cloudsync//media/video\\_cloud\\_sync\\_registering.mp4](https://docs.netapp.com/us-en/cloudsync//media/video_cloud_sync_registering.mp4) (video)

### Steps to pay annually

1. [Go to the AWS Marketplace page](#).
2. Click **Continue to Subscribe**.
3. Select your contract options and click **Create contract**.

## Subscribing from Azure

Azure enables you to pay-as-you-go or to pay annually.

### What you'll need

An Azure user account that has Contributor or Owner permissions in the relevant subscription.

### Steps

1. Go to the [License Settings](#) page.
2. Select **Azure**.
3. Click **Subscribe** and then click **Continue**.
4. In the Azure portal, click **Create**, select your options, and click **Subscribe**.

Select **Monthly** to pay by the hour, or **Yearly** to pay for a year up front.

5. When deployment is complete, click the name of the SaaS resource in the notification pop-up.
6. Click **Configure Account** to return to Cloud Sync.

The following video shows the process:

► [https://docs.netapp.com/us-en/cloudsync//media/video\\_cloud\\_sync\\_registering\\_azure.mp4](https://docs.netapp.com/us-en/cloudsync//media/video_cloud_sync_registering_azure.mp4) (video)

## Purchasing licenses from NetApp and adding them to Cloud Sync

To pay for your sync relationships up front, you must purchase one or more licenses and add them to the Cloud Sync service.

### Steps

1. Purchase a license by [contacting NetApp](#).
2. Go to the [License Settings](#) page and add the license.

# Tutorials

## Copying ACLs

Cloud Sync can copy access control lists (ACLs) between a source SMB share and a target SMB share, or between a source NFS server and target NFS server. If needed, you can manually preserve ACLs for SMB shares yourself by using robocopy.

### Choices

- [Set up Cloud Sync to automatically copy ACLs](#)
- [Manually copy ACLs between SMB shares](#)

## Setting up Cloud Sync to copy ACLs

Copy ACLs between SMB servers or between NFS servers by enabling a setting when you create a relationship or after you create a relationship.

### What you'll need

- A new sync relationship or an existing sync relationship.

For SMB shares, note that this feature is available for new sync relationships created after the 23 Feb 2020 release. If you'd like to use this feature with existing relationships created prior to that date, then you'll need to recreate the relationship.

- Any type of data broker.

This feature works with *any* type of data broker: the AWS, Azure, Google Cloud Platform, or on-prem data broker. The on-prem data broker can run [any supported operating system](#).

- For NFS, you'll need to use version 4 or later.

Copying ACLs isn't supported with NFS version 3.


### Steps for a new relationship

1. From Cloud Sync, click **Create New Sync Relationship**.
2. Drag and drop **SMB Server** to the source and target or **NFS Server** to the source and target and click **Continue**.
3. On the **SMB Server** or **NFS Server** page:
  - a. Enter a new server or select an existing server and click **Continue**.
  - b. Select **Copy Access Control Lists to the target** and click **Continue**.

SMB Server > Data Broker > Shares > Target SMB Server > Target Shares > Settings > Review

### Select an SMB Source

SMB Version : 2.1 ▼

 Selected SMB Server:  
10.20.30.152

Define SMB Credentials:

User Name	Password	Domain (Optional)
user1	*****	

ACL - Access Control List

☒ Copy Access Control Lists to the target

**Notice:** Copying ACLs can affect sync performance.  
You can change this setting after you create the relationship.

4. Follow the remaining prompts to create the sync relationship.

#### Steps for an existing relationship

1. Hover over the sync relationship and click the action menu.
2. Click **Settings**.
3. Select **Copy Access Control Lists to the target**.
4. Click **Save Settings**.

#### Result

When syncing data, Cloud Sync preserves the ACLs between the source and target servers.

## Manually copying ACLs between SMB shares

You can manually preserve ACLs between SMB shares by using the Windows robocopy command.

#### Steps

1. Identify a Windows host that has full access to both SMB shares.
2. If either of the endpoints require authentication, use the **net use** command to connect to the endpoints from the Windows host.

You must perform this step before you use robocopy.

3. From Cloud Sync, create a new relationship between the source and target SMB shares or sync an existing relationship.
4. After the data sync is complete, run the following command from the Windows host to sync the ACLs and

ownership:

```
robocopy /E /COPY:SOU /secfix [source] [target] /w:0 /r:0 /XD ~snapshots  
/UNILog:"[logfilepath]
```

Both *source* and *target* should be specified using the UNC format. For example: \\<server>\<share>\<path>


## Syncing NFS data using data-in-flight encryption

If your business has strict security policies, you can sync NFS data using data-in-flight encryption. This feature is supported from an NFS server to another NFS server and from Azure NetApp Files to Azure NetApp Files.

For example, you might want to sync data between two NFS servers that are in different networks. Or you might need to securely transfer data on Azure NetApp Files across subnets or regions.

### How data-in-flight encryption works

Data-in-flight encryption encrypts NFS data when it's sent over the network between two data brokers. The following image shows a relationship between two NFS servers and two data brokers:



One data broker functions as the *initiator*. When it's time to sync data, it sends a connection request to the other data broker, which is the *listener*. That data broker listens for requests on port 443. You can use a different port, if needed, but be sure to check that the port is not in use by another service.

For example, if you sync data from an on-premises NFS server to a cloud-based NFS server, you can choose which data broker listens for the connection requests and which sends them.

Here's how in-flight encryption works:

1. After you create the sync relationship, the initiator starts an encrypted connection with the other data broker.
2. The source data broker encrypts data from the source using TLS 1.3.
3. It then sends the data over the network to the target data broker.
4. The target data broker decrypts the data before sending it to the target.
5. After the initial copy, the service syncs any changed data every 24 hours. If there is data to sync, the process starts with the initiator opening an encrypted connection with the other data broker.

If you prefer to sync data more frequently, [you can change the schedule after you create the relationship](#).

## Supported NFS versions

- For NFS servers, data-in-flight encryption is supported with NFS versions 3, 4.0, 4.1, and 4.2.
- For Azure NetApp Files, data-in-flight encryption is supported with NFS versions 3 and 4.1.

## Proxy server limitation

If you create an encrypted sync relationship, the encrypted data is sent over HTTPS and isn't routable through a proxy server.

## What you'll need to get started

Be sure to have the following:

- Two NFS servers that meet [source and target requirements](#) or Azure NetApp Files in two subnets or regions.
- The IP addresses or fully qualified domain names of the servers.
- Network locations for two data brokers.

You can select an existing data broker but it must function as the initiator. The listener data broker must be a *new* data broker.

If you have not yet deployed a data broker, review the data broker requirements. Because you have strict security policies, be sure to review the networking requirements, which includes outbound traffic from port 443 and the [internet endpoints](#) that the data broker contacts.

- [Review AWS installation](#)
- [Review Azure installation](#)
- [Review GCP installation](#)
- [Review Linux host installation](#)

## Syncing NFS data using data-in-flight encryption

Create a new sync relationship between two NFS servers or between Azure NetApp Files, enable the in-flight encryption option, and follow the prompts.

### Steps

1. Log in to [NetApp Cloud Central](#) and select Cloud Sync.
2. Click **Create New Sync Relationship**.
3. Drag and drop **NFS Server** to the source and target locations or **Azure NetApp Files** to the source and target locations and select **Yes** to enable data-in-flight encryption.

The following image shows what you'd select to sync data between two NFS servers:





The following image shows what you'd select to sync data between Azure NetApp Files:

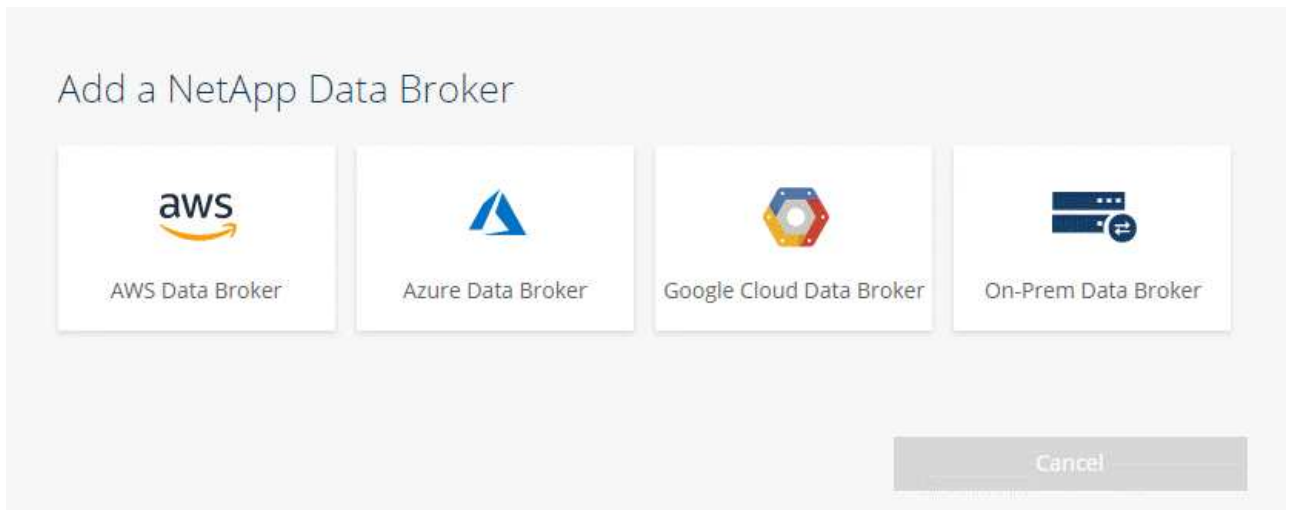


4. Follow the prompts to create the relationship:

- a. **NFS Server/Azure NetApp Files:** Choose the NFS version and then specify a new NFS source or select an existing server.
- b. **Define Data Broker Functionality:** Define which data broker *listens* for connection requests on a port and which one *initiates* the connection. Make your choice based on your networking requirements.
- c. **Data Broker:** Follow the prompts to add a new source data broker or select an existing data broker.

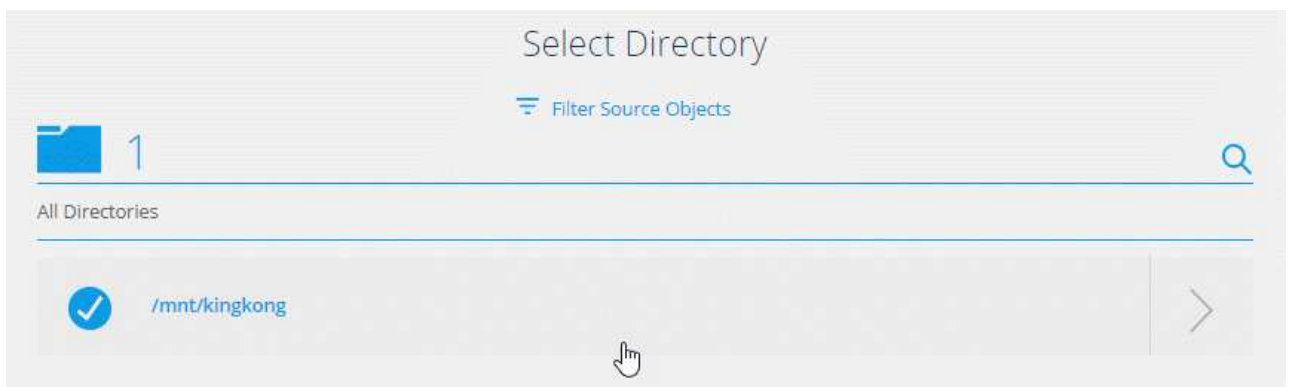
If the source data broker acts as the listener, then it must be a new data broker.

If you need a new data broker, Cloud Sync prompts you with the installation instructions. You can deploy the data broker in the cloud or download an installation script for your own Linux host.



- d. **Directories:** Choose the directories that you want to sync by selecting all directories, or by drilling down and selecting a subdirectory.

Click **Filter Source Objects** to modify settings that define how source files and folders are synced and maintained in the target location.





- e. **Target NFS Server/Target Azure NetApp Files:** Choose the NFS version and then enter a new NFS target or select an existing server.
- f. **Target Data Broker:** Follow the prompts to add a new source data broker or select an existing data broker.


If the target data broker acts as the listener, then it must be a new data broker.


Here's an example of the prompt when the target data broker functions as the listener. Notice the option to specify the port.

## Add a NetApp Data Broker

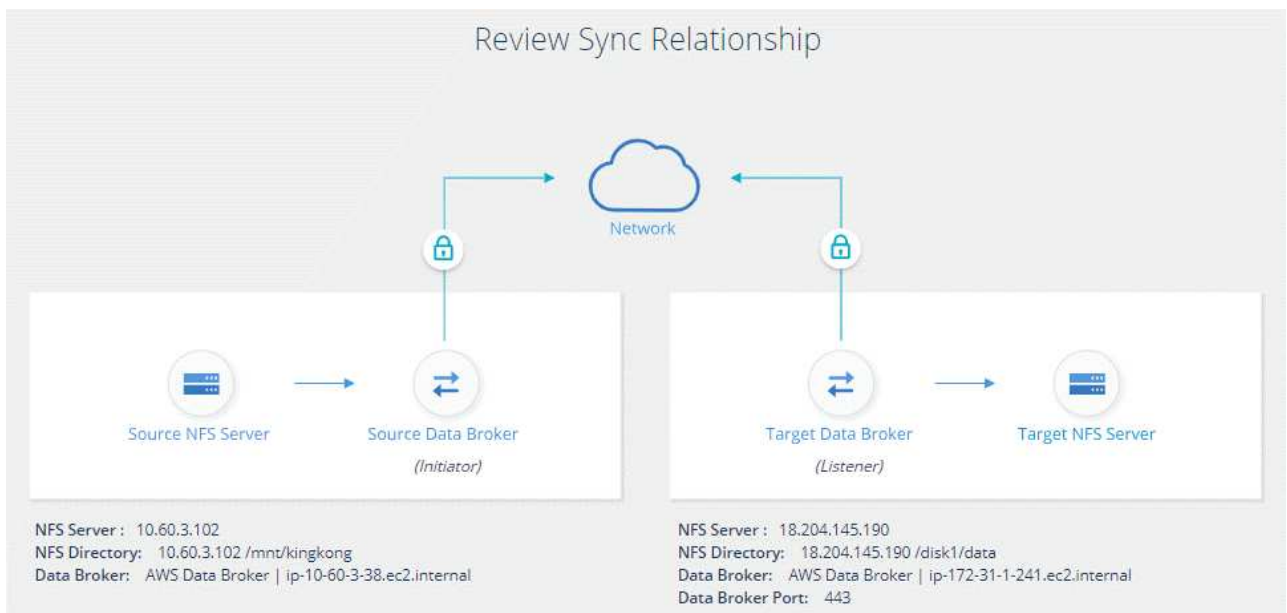
  
AWS Data Broker

  
Azure Data Broker

  
GCP Data Broker

  
On-Prem Data Broker

- g. **Target Directories:** Select a top-level directory, or drill down to select an existing subdirectory or to create a new folder inside an export.
- h. **Settings:** Define how source files and folders are synced and maintained in the target location.
- i. **Review:** Review the details of the sync relationship and then click **Create Relationship**.



## Result

Cloud Sync starts creating the new sync relationship. When it's done, click **View in Dashboard** to view details about the new relationship.

## Setting up the data broker to use an external HashiCorp Vault

When you create a sync relationship that requires Amazon S3, Azure, or Google Cloud credentials, you need to specify those credentials through the Cloud Sync user interface or API. An alternative is to set up the data broker to access the credentials (or *secrets*) directly from an external HashiCorp Vault.

This feature is supported through the Cloud Sync API with sync relationships that require Amazon S3, Azure, or Google Cloud credentials.



### Prepare the vault

Prepare the vault to supply credentials to the data broker by setting up the URLs. The URLs to the secrets in the vault must end with *Creds*.



### Prepare the data broker

Prepare the data broker to fetch credentials from the external vault by modifying the local config file for the data broker.



### Create a sync relationship using the API

Now that everything is set up, you can send an API call to create a sync relationship that uses your vault to get the secrets.

## Preparing the vault

You'll need to provide Cloud Sync with the URL to the secrets in your vault. Prepare the vault by setting up those URLs. You need to set up URLs to the credentials for each source and target in the sync relationships that you plan to create.

The URL must be set up as follows:

```
/<path>/<requestid>/<endpoint-protocol>Creds
```

### Path

The prefix path to the secret. This can be any value that's unique to you.

### Request ID

A request ID that you need to generate. You'll need to provide the ID in one of the headers in the API POST request when you create the sync relationship.

### Endpoint protocol

One of the following protocols, as defined [in the post relationship v2 documentation](#): S3, AZURE, or GCP (each must be in uppercase).

## Creds

The URL must end with *Creds*.

## Examples

The following examples show URLs to secrets.

### Example for the full URL and path for source credentials

`http://example.vault.com:8200/my-path/all-secrets/hb312vdsr2/S3Creds`

As you can see in the example, the prefix path is */my-path/all-secrets/*, the request ID is *hb312vdsr2* and the source endpoint is S3.

### Example for the full URL and path for target credentials

`http://example.vault.com:8200/my-path/all-secrets/n32hcbnejk2/AZURECreds`

The prefix path is */my-path/all-secrets/*, the request ID is *n32hcbnejk2*, and the target endpoint is Azure.

## Preparing the data broker

Prepare the data broker to fetch credentials from the external vault by modifying the local config file for the data broker.

### Steps

1. SSH to the data broker.
2. Edit the `local.json` file that resides in `/opt/netapp/databroker/config`.
3. Set `enable` to **true** and set the config parameter fields under *external-integrations.hashicorp* as follows:

#### **enabled**

- Valid values: `true/false`
- Type: Boolean
- Default value: `false`
- True: The data broker gets secrets from your own external HashiCorp Vault
- False: The data broker stores credentials in its local vault

#### **url**

- Type: string
- Value: The URL to your external vault

#### **path**

- Type: string
- Value: Prefix path to the secret with your credentials

#### **Reject-unauthorized**

- Determines if you want the data broker to reject unauthorized external vault
- Type: Boolean
- Default: `false`

**auth-method**

- Your authentication method to the external vault
- Type: string
- Valid values: “aws-iam” / “role-app”

**role-name**

- Type: string
- Your role name (in case you use aws-iam)

**Secretid & rootid**

- Type: string (in case you use app-role)

**Namespace**

- Type: string
- Your namespace (X-Vault-Namespace header if needed)

**Example**

```
{
  "external-integrations": {
    "hashicorp": {
      "enabled": true,
      "url": "https://example.vault.com:8200",
      "path": "my-path/all-secrets",
      "reject-unauthorized": false,
      "auth-method": "aws-role",
      "aws-role": {
        "role-name": "my-role"
      }
    }
  }
}
```

**Creating a new sync relationship using secrets from the vault**

Now that everything is set up, you can send an API call to create a sync relationship that uses your vault to get the secrets.

Post the relationship using the Cloud Sync REST API.

Headers:

Authorization: Bearer <user-token>

Content-Type: application/json

x-account-id: <accountid>

x-netapp-external-request-id-src: request ID as part of path for source credentials

x-netapp-external-request-id-trg: request ID as part of path for target credentials

Body: post relationship v2 body

- To obtain a user token and your Cloud Central account ID, [refer to this page in the documentation](#).
- To build a body for your post relationship, [refer to the relationships-v2 API call](#).

### Example

Example for the POST request:

```
url: https://api.cloudsync.netapp.com/api/relationships-v2
headers:
"x-account-id": "CS-SasdW"
"x-netapp-external-request-id-src": "hb312vdasr2"
"Content-Type": "application/json"
"Authorization": "Bearer eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6Ik..."
Body:
{
  "dataBrokerId": "5e6e111d578dtyuul555sa60",
  "source": {
    "protocol": "s3",
    "s3": {
      "provider": "sgws",
      "host": "1.1.1.1",
      "port": "443",
      "bucket": "my-source"
    }
  },
  "target": {
    "protocol": "s3",
    "s3": {
      "bucket": "my-target-bucket"
    }
  }
}
```

# Managing sync relationships

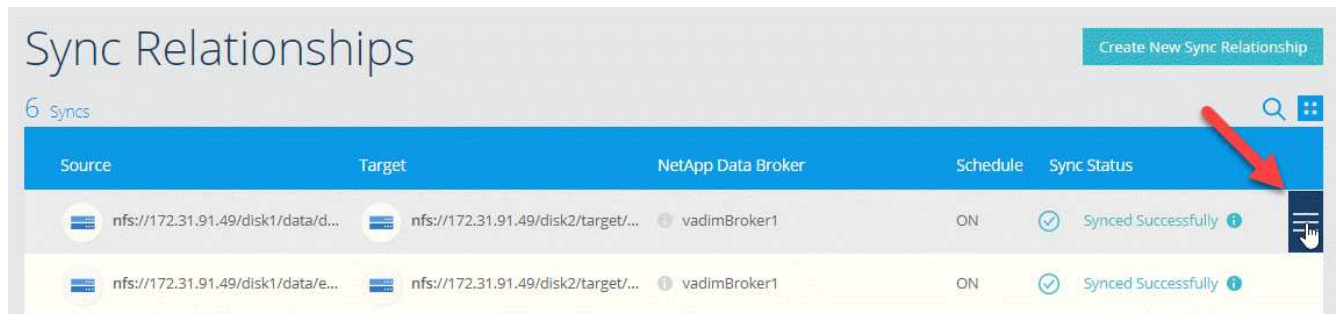
You can manage sync relationships at any time by immediately syncing data, changing schedules, and more.

## Performing an immediate data sync


Rather than wait for the next scheduled sync, you can press a button to immediately sync data between the source and target.

### Steps

1. Hover over the sync relationship and click the action menu.



2. Click **Sync Now** and then click **Sync** to confirm.



### Result

Cloud Sync starts the data sync process for the relationship.

## Accelerating sync performance

Accelerate the performance of a sync relationship by adding an additional data broker to the relationship. The additional data broker must be a *new* data broker.

### How this works

If the existing data brokers in the relationship are used in other sync relationships, then Cloud Sync automatically adds the new data broker to those relationships, as well.

For example, let's say you have three relationships:

- Relationship 1 uses data broker A
- Relationship 2 uses data broker B




- Relationship 3 uses data broker A

You want to accelerate the performance of relationship 1 so you add a new data broker to that relationship (data broker C). Because data broker A is also used in relationship 3, the new data broker is automatically added to relationship 3, as well.

### Steps

1. Go to the **Sync Relationships** dashboard.
2. Ensure that at least one of the existing data brokers in the relationship are online.
3. Hover over the sync relationship and click the action menu.
4. Click **Accelerate**.



5. Follow the prompts to create a new data broker.

### Result

Cloud Sync adds the new data broker to the sync relationships. The performance of the next data sync should be accelerated.

## Changing the settings for a sync relationship

Modify settings that define how source files and folders are synced and maintained in the target location.

1. Hover over the sync relationship and click the action menu.
2. Click **Settings**.
3. Modify any of the settings.

General

Schedule	ON   Every 1 Day	▼
Retries	Retry 3 times before skipping file	▼

Files and Directories

Recently Modified Files	Exclude files that are modified up to 30 Seconds before a scheduled sync	▼
Delete Files On Source	Never delete files from the source location	▼
Delete Files On Target	Never delete files from the target location	▼
Object Tagging	Allow Cloud Sync to tag S3 objects	▼
File Types	Include All: Files, Directories, Symbolic Links	▼
Exclude File Extensions	None	▼
File Size	All	▼
Date Modified	All	▼

Reset to defaults

Here's a brief description of each setting:

### Schedule

Choose a recurring schedule for future syncs or turn off the sync schedule. You can schedule a relationship to sync data as often as every 1 minute.

### Retries

Define the number of times that Cloud Sync should retry to sync a file before skipping it.

### Recently Modified Files

Choose to exclude files that were recently modified prior to the scheduled sync.

### Delete Files on Source

Choose to delete files from the source location after Cloud Sync copies the files to the target location. This option includes the risk of data loss because the source files are deleted after they're copied.

If you enable this option, you also need to change a parameter in the local.json file on the data broker. Open the file and change the parameter named *workers.transferrer.delete-on-source* to **true**.

### Delete Files on Target

Choose to delete files from the target location, if they were deleted from the source. The default is to never delete files from the target location.

### Object tagging

When AWS S3 is the target in a sync relationship, Cloud Sync tags S3 objects with metadata that's relevant to the sync operation. You can disable tagging of S3 objects, if it's not desired in your environment. There's no impact to Cloud Sync if you disable tagging—Cloud Sync just stores the sync metadata in a different way.

### File Types

Define the file types to include in each sync: files, directories, and symbolic links.

### Exclude File Extensions

Specify file extensions to exclude from the sync by typing the file extension and pressing **Enter**. For example, type *log* or *.log* to exclude \*.log files. A separator isn't required for multiple extensions. The following video provides a short demo:

► [https://docs.netapp.com/us-en/cloudsync//media/video\\_file\\_extensions.mp4](https://docs.netapp.com/us-en/cloudsync//media/video_file_extensions.mp4) (video)

### File Size

Choose to sync all files regardless of their size or just files that are in a specific size range.

### Date Modified

Choose all files regardless of their last modified date, files modified after a specific date, before a specific date, or between a time range.

### Copy Access Control Lists to the target

Choose to copy access control lists (ACLs) between source SMB shares and target SMB shares. Note that this option is only available for sync relationships created after the 23 Feb 2020 release.

4. Click **Save Settings**.

### Result

Cloud Sync modifies the sync relationship with the new settings.

## Creating and viewing reports about paths

Create and view reports to get information that you can use with the help of NetApp personnel to tune a data broker's configuration and improve performance.

Each report provides in-depth details about a path in a sync relationship. For example, the report for a file system shows how many directories and files there are, the distribution of file size, how deep and wide the directories are, and more.

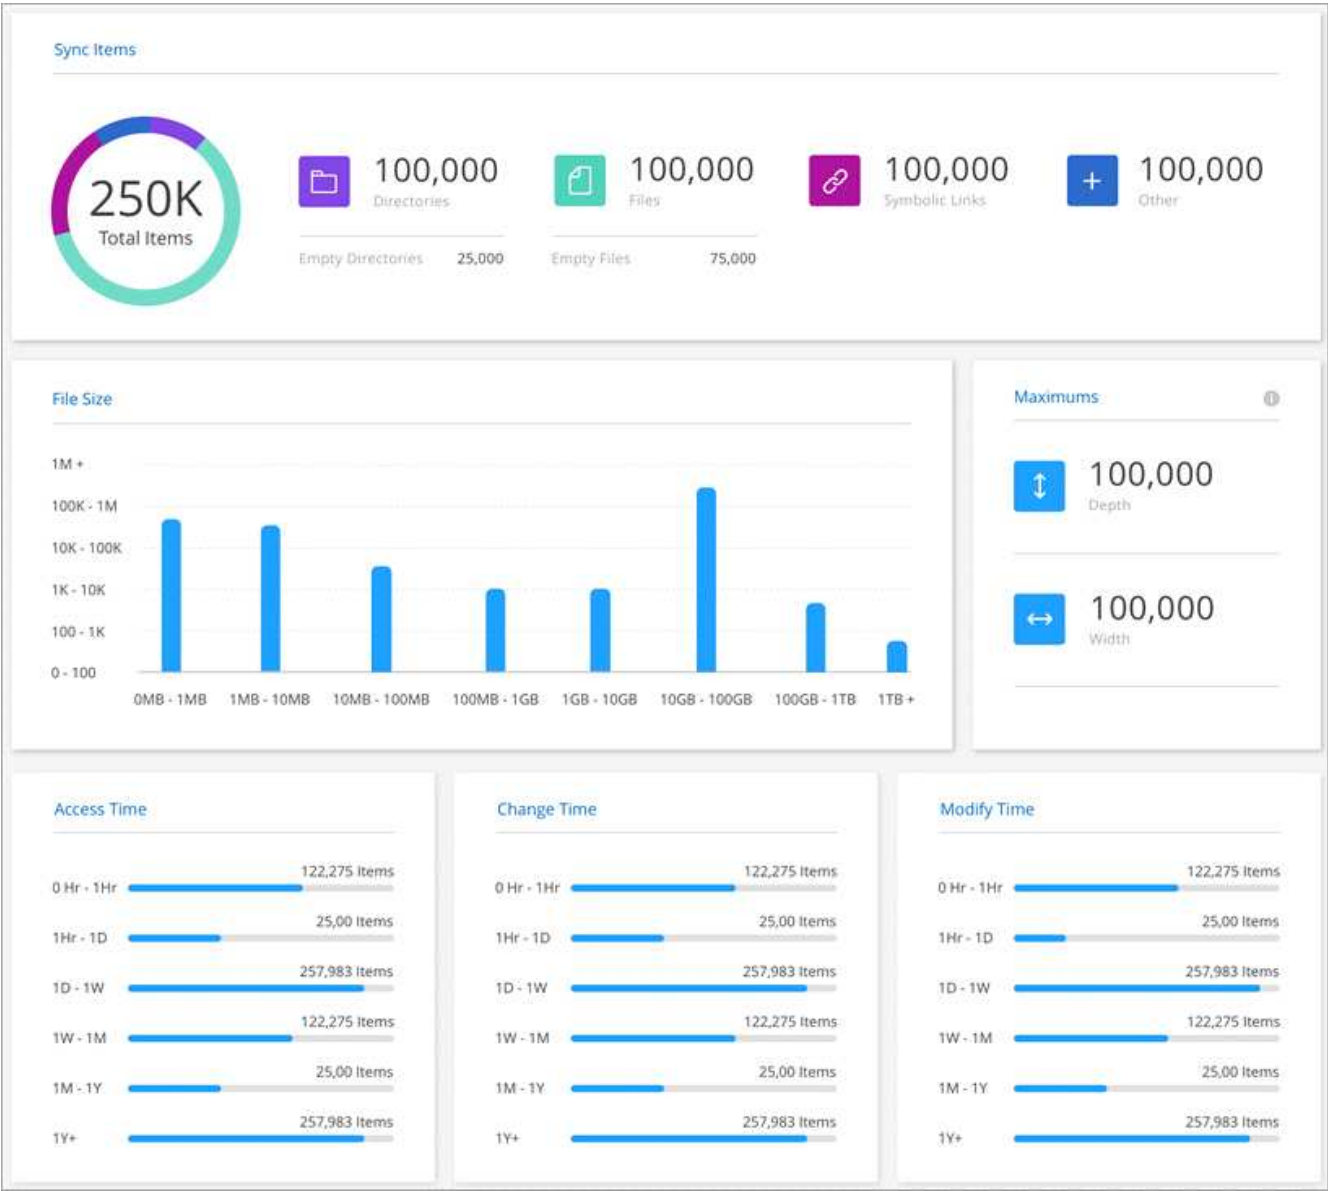
### Steps

1. Click **Reports**.


The paths (source or target) in each of your sync relationships display in a table.

2. In the **Reports** column, click **Create New** for a path.
3. When the report is ready, click **View**.

Here’s a sample report for a file system path.



And here’s a sample report for object storage.



## Deleting relationships

You can delete a sync relationship, if you no longer need to sync data between the source and target. This action does not delete the data broker instance and it does not delete data from the target.

### Steps

1. Hover over the sync relationship and click the action menu.
2. Click **Delete** and then click **Delete** again to confirm.

### Result

Cloud Sync deletes the sync relationship.

# Manage data brokers

A data broker syncs data from a source location to a target location. A data broker is required for each sync relationship that you create. Manage data brokers by adding a new data broker to a group, by viewing information about data brokers, and more.

## Data broker groups

Grouping data brokers together can help improve the performance of sync relationships.

### Determining the number of data brokers

In many cases, a single data broker can meet the performance requirements for a sync relationship. If it doesn't, you can accelerate sync performance by adding additional data brokers to the group. But you should first check other factors that can impact sync performance. [Learn more about how to determine when multiple data brokers are required.](#)

### Groups can manage several relationships

A data broker group can manage one or more sync relationships at a time.

For example, let's say you have three relationships:

- Relationship 1 uses data broker A
- Relationship 2 uses data broker B
- Relationship 3 uses data broker A

You want to accelerate the performance of relationship 1 so you add a new data broker (data broker C) to the group. Because data broker A is also used to manage relationship 3, having two data brokers in the group also accelerates the performance of this relationship.

### Supported types of data broker in a group

A data broker group can consist of one or more AWS, Azure, or GCP data brokers and one or more on-prem data brokers. A group can't consist of a mixture of AWS, Azure, and GCP data brokers. For example, a group can have two AWS data brokers and one on-prem data broker, but not one AWS data broker and one Azure data broker.

### New data brokers only

You can only add new data brokers to a group. You can't add existing data brokers to a group.

## Add a new data broker

There are several ways to create a new data broker:

- When creating a new sync relationship

[Learn how to create a new data broker when creating a sync relationship.](#)

- From the **Manage Data Brokers** page by clicking **Add New Data Broker** which creates the data broker in a new group
- From the **Manage Data Brokers** page by creating a new data broker in an existing group

### Things you should know

- You can't add data brokers to a group that manages an encrypted sync relationship.
- If you want to create a data broker in an existing group, the data broker must be an on-prem data broker or the same type of data broker.

For example, if a group includes an AWS data broker, then you can create an AWS data broker or on-prem data broker in that group. You can't create an Azure data broker or GCP data broker because they aren't the same data broker type.

### Steps to create a new data broker in a new group


1. Click **Sync > Manage Data Brokers**.
2. Click **Add New Data Broker**.
3. Follow the prompts to create the data broker.

For help, refer to the following pages:

- [Installing the data broker in AWS](#)
- [Installing the data broker in Azure](#)
- [Installing the data broker in GCP](#)
- [Installing the data broker on a Linux host](#)

### Steps to create a new data broker in an existing group

1. Click **Manage Data Brokers**.
2. Click the action menu and select **Edit Group Name**.



3. Follow the prompts to create the data broker.

For help, refer to the following pages:

- [Installing the data broker in AWS](#)
- [Installing the data broker in Azure](#)
- [Installing the data broker in GCP](#)
- [Installing the data broker on a Linux host](#)

# View a data broker's configuration

You might want to view details about a data broker to identify things like its host name, IP address, available CPU and RAM, and more.



Cloud Sync provides the following details about a data broker:


- Basic information: Instance ID, host name, etc.
- Network: Region, network, subnet, private IP, etc.
- Software: Linux distribution, data broker version, etc.
- Hardware: CPU and RAM
- Configuration: Details about the data broker's two kinds of main processes—scanner and transferrer





The scanner scans the source and target and decides what should be copied. The transferrer does the actual copying. NetApp personnel might use these configuration details to suggest actions that can optimize performance.

## Steps

1. Click **Manage Data Brokers**.
2. Click  to expand the list of data brokers in a group.
3. Click  to view details about a data broker.



 tanyagcp0212


 

2  
Data Brokers

968.5 B/s  
Transfer Rate


1  
Relationships



 1 Active  1 Unknown  
Data Brokers Status

 tanyagcp0212

GCP

Transfer Rate: 968.5 B/s

 Active

Information	5fc766b3d3e3664b9e116... Broker ID	288871247573080556 Instance ID	tanyagcp0212-mnx-data-... Host Name	cloudsync-dev-214020 Project Id
Network	us-east1-b Region	default Network	255.255.240.0 Subnet	10.142.0.37 Private IP
Software	linux Linux Distribution & Version	1.5.4 Vault Version	14.15.1 Node Version	1.3.0.18650-73f960d-integ Data Broker Version
Hardware	4 Available CPUs	62.22 MB Available RAM		
Configuration	50 Scanner Concurrency	4 Scanner CPUs	50 Transferrer Concurrency	4 Transferrer CPUs

# Remove a data broker from a group

You might remove a data broker from a group if it's no longer needed or if the initial deployment failed. This




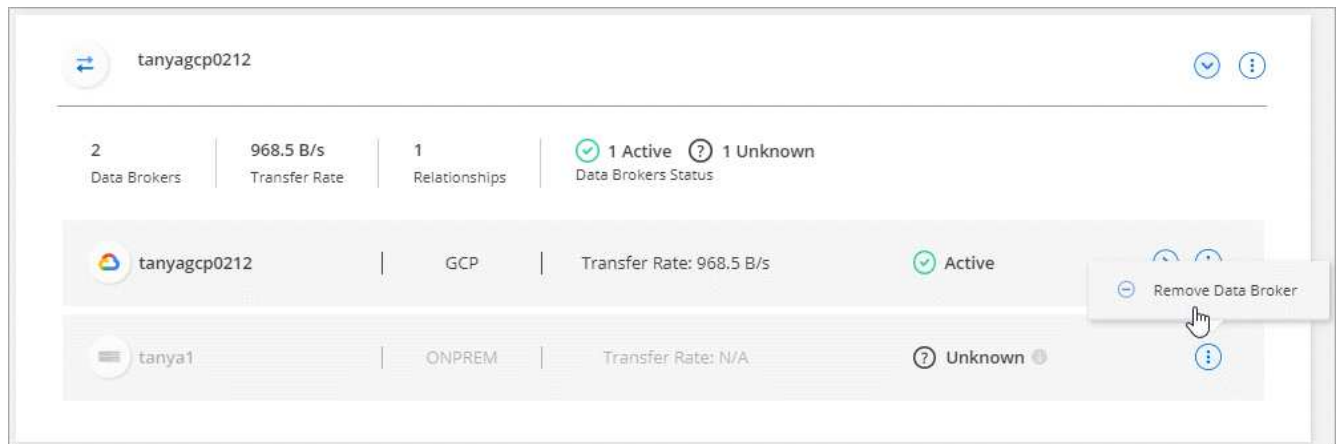
action only deletes the data broker from Cloud Sync's records. You'll need to manually delete the data broker and any additional cloud resources yourself.

### Things you should know

- Cloud Sync deletes a group when you remove the last data broker from the group.
- You can't remove the last data broker from a group if there is a relationship using that group.

### Steps

1. Click **Manage Data Brokers**.
2. Click  to expand the list of data brokers in a group.
3. Click the action menu for a data broker and select **Remove Data Broker**.



4. Click **Remove Data Broker**.

### Result


Cloud Sync removes the data broker from the group.

## Edit a group's name

Change the name of a data broker group at any time.

### Steps

1. Click **Manage Data Brokers**.
2. Click the action menu and select **Edit Group Name**.



3. Enter a new name and click **Save**.

### Result

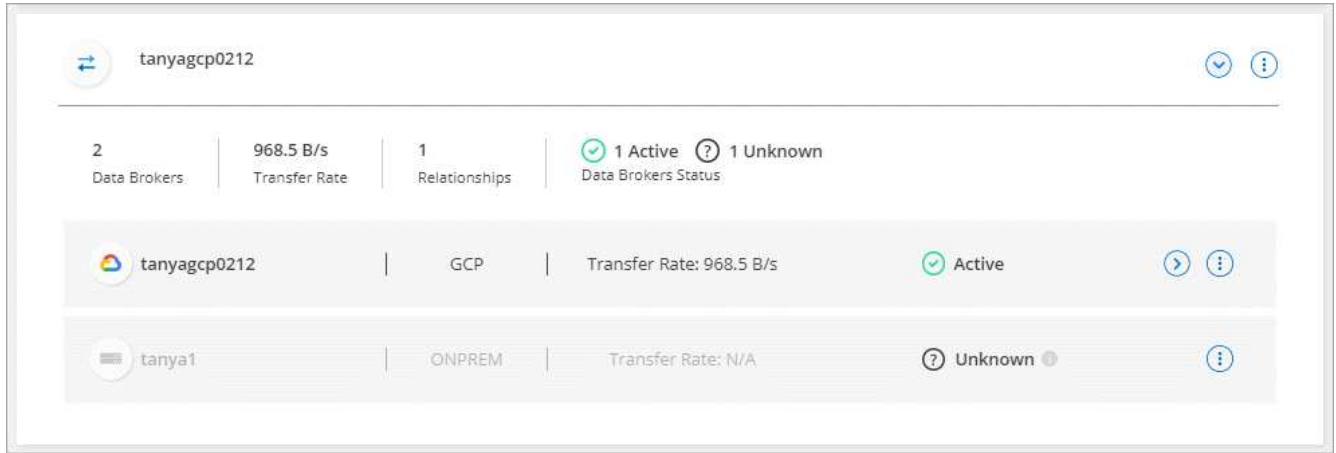
Cloud Sync updates the name of the data broker group.

# Address issues with a data broker

Cloud Sync displays a status for each data broker that can help you troubleshoot issues.

## Steps

1. Identify any data brokers that have a status of "Unknown" or "Failed."



2. Hover over the ⓘ to see the failure reason.
3. Correct the issue.

For example, you might need to simply restart the data broker if it's offline, or you might need to remove data broker if the initial deployment failed.

# Defining a unified configuration for a data broker group

If a sync relationship encounters errors during the sync process, unifying the concurrency of the data broker group can help to decrease the number of sync errors. Be aware that changes to the group's configuration can affect performance by slowing down the transfer.

We don't recommend changing the configuration on your own. You should consult with NetApp to understand when to change the configuration and how to change it.

## Steps

1. Click **Manage Data Brokers**.
2. Click the Settings icon for a data broker group.
3. Change the settings as needed and then click **Unify Configuration**.

Note the following:

- You can pick and choose which settings to change—you don't need to change all four at once.
- After a new configuration is sent to a data broker, the data broker automatically restarts and uses the new configuration.
- It can take up to a minute until this change takes place and is visible in the Cloud Sync interface.
- If a data broker isn't running, its configuration won't change because Cloud Sync can't communicate with it. The configuration will change after the data broker restarts.

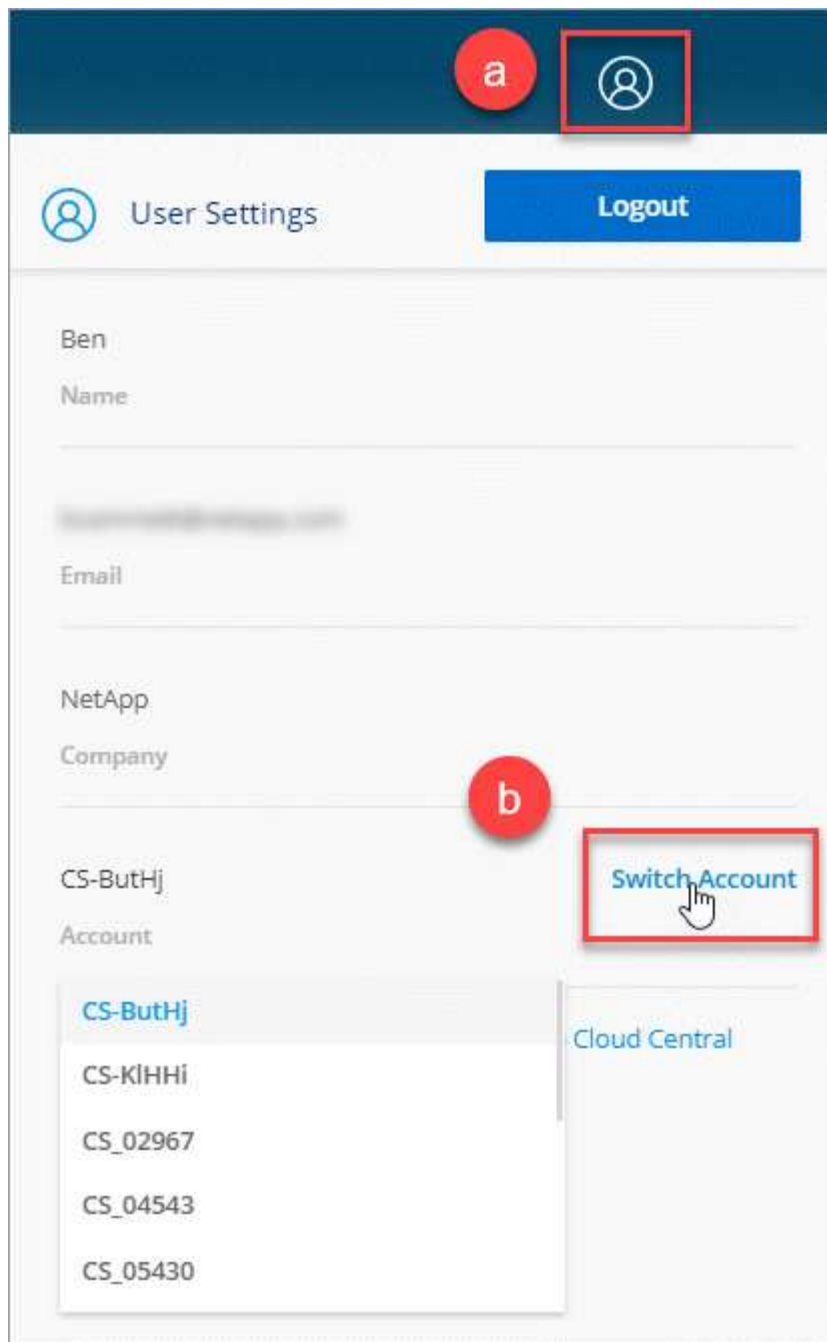
- After you set a unified configuration, any new data brokers will automatically use the new configuration.

# Associating users to an account

Associate additional users to a Cloud Central account so those users can see the same sync relationships and data brokers as other users in the account. These steps must be completed from Cloud Manager by an Account Admin.

## Steps

1. Ask the new user to go to [NetApp Cloud Central](#) and create a user account.
2. [Log in to NetApp Cloud Manager with a user who is an Account Admin](#).
3. If you have multiple Cloud Central accounts, switch to the account for Cloud Sync.
4. [Go into the account settings and associate the user](#).
5. The newly associated user should go to [Cloud Sync](#) and click the **User Settings** menu.
6. The user should then scroll down to Account, click **Switch Account**, and select the account that was just associated.



The user should now see the relationships for that account.

# Uninstalling the data broker

If needed, run an uninstall script to remove the data broker and the packages and directories that were created when the data broker was installed.

## Steps

1. Log in to the data broker host.
2. Change to the data broker directory: `/opt/netapp/databroker`
3. Run the following commands:

```
chmod +x uninstaller-DataBroker.sh  
./uninstaller-DataBroker.sh
```

4. Press 'y' to confirm the uninstallation.

# Cloud Sync APIs

The Cloud Sync capabilities that are available through the web UI are also available through the RESTful API.

## Getting started

To get started with the Cloud Sync API, you need to obtain a user token and your Cloud Central account ID. You'll need to add the token and account ID to the Authorization header when making API calls.

### Steps

1. Obtain a user token from NetApp Cloud Central.

```
POST https://netapp-cloud-account.auth0.com/oauth/token
Header: Content-Type: application/json
Body:
{
  "username": "<user_email>",
  "scope": "profile",
  "audience": "https://api.cloud.netapp.com",
  "client_id": "UaVhOIXMWQs5i1WdDxauXe5Mqkb34NJQ",
  "grant_type": "password",
  "password": "<user_password>"
}
```

2. Obtain your Cloud Central account ID.

```
GET https://api.cloudsync.netapp.com/api/accounts
Headers: Authorization: Bearer <user_token>
Content-Type: application/json
```

This API will return a response like the following:

```
[
  {
    "accountId": "account-JeL97Ry3",
    "name": "Test"
  }
]
```

3. Add the user token and account ID in the Authorization header of each API call.

### Example

The following example shows an API call to create a data broker in Microsoft Azure. You would simply replace `<user_token>` and `<accountId>` with the token and ID that you obtained in the previous steps.

```
POST https://api.cloudsync.netapp.com/api/data-brokers
Headers: Authorization: Bearer <user_token>
Content-Type: application/json
x-account-id: <accountId>
Body: { "name": "databroker1", "type": "AZURE" }
```

## What should I do when the token expires?

The user token from NetApp Cloud Central has an expiration date. To refresh the token, you need to call the API from step 1 again.

The API response includes an "expires\_in" field that states when the token expires.

## API reference

Documentation for each Cloud Sync API is available from <https://api.cloudsync.netapp.com/docs>.

## Using list APIs

List APIs are asynchronous APIs, so the result does not return immediately (for example: `GET /data-brokers/{id}/list-nfs-export-folders` and `GET /data-brokers/{id}/list-s3-buckets`). The only response from the server is HTTP status 202. To get the actual result, you must use the `GET /messages/client` API.

### Steps

1. Call the list API that you want to use.
2. Use the `GET /messages/client` API to view the result of the operation.
3. Use the same API by appending it with the ID that you just received: `GET http://api.cloudsync.netapp.com/api/messages/client?last=<id_from_step_2>`

Note that the ID changes each time that you call the `GET /messages/client` API.

### Example

When you call the `list-s3-buckets` API, a result is not immediately returned:

```
GET http://api.cloudsync.netapp.com/api/data-brokers/<data-broker-id>/list-s3-buckets
Headers: Authorization: Bearer <user_token>
Content-Type: application/json
x-account-id: <accountId>
```



The result is HTTP status code 202, which means the message was accepted, but was not processed yet.

To get the result of the operation, you need to use the following API:

```
GET http://api.cloudsync.netapp.com/api/messages/client
Headers: Authorization: Bearer <user_token>
Content-Type: application/json
x-account-id: <accountId>
```

The result is an array with one object that includes an ID field. The ID field represents the last message that the server sent. For example:

```
[
  {
    "header": {
      "requestId": "init",
      "clientId": "init",
      "agentId": "init"
    },
    "payload": {
      "init": {}
    },
    "id": "5801"
  }
]
```

You would now make the following API call using the ID that you just received:

```
GET
http://api.cloudsync.netapp.com/api/messages/client?last=<id_from_step_2>
Headers: Authorization: Bearer <user_token>
Content-Type: application/json
x-account-id: <accountId>
```

The result is an array of messages. Inside each message is a payload object, which consists of the name of the operation (as key) and its result (as value). For example:

```
[
  {
    "payload": {
      "list-s3-buckets": [
        {
          "tags": [
            {
              "Value": "100$",
              "Key": "price"
            }
          ],
          "region": {
            "displayName": "US West (Oregon)",
            "name": "us-west-2"
          },
          "name": "small"
        }
      ]
    },
    "header": {
      "requestId": "f687ac55-2f0c-40e3-9fa6-57fb8c4094a3",
      "clientId": "5beb032f548e6e35f4ed1ba9",
      "agentId": "5bed61f4489fb04e34a9aac6"
    },
    "id": "5802"
  }
]
```

# Cloud Sync technical FAQ

This FAQ can help if you're just looking for a quick answer to a question.

## Getting started

The following questions relate to getting started with Cloud sync.

### How does Cloud Sync work?

Cloud Sync uses the NetApp data broker software to sync data from a source to a target (this is called a *sync relationship*).

The data broker controls the sync relationships between your sources and targets. After you set up a sync relationship, Cloud Sync analyzes your source system and breaks it up into multiple replication streams to push to your selected target data.

After the initial copy, the service syncs any changed data based on the schedule that you set.

### How does the 14-day free trial work?

The 14-day free trial starts when you sign up for the Cloud Sync service. You're not subject to NetApp charges for Cloud Sync relationships you create for 14 days. However, all resource charges for any data broker that you deploy still applies.

### How much does Cloud Sync cost?

There are two types of costs associated with using Cloud Sync: service charges and resource charges.

#### Service charges

For pay-as-you-go pricing, Cloud Sync service charges are hourly, based on the number of sync relationships that you create.

- [View pay-as-you-go pricing in AWS](#)
- [View annual pricing in AWS](#)
- [View pricing in Azure](#)

Cloud Sync licenses are also available through your NetApp representative. Each license enables 20 sync relationships for 12 months.

[Learn more about licenses.](#)

#### Resource charges

The resource charges are related to the compute and storage costs for running the data broker in the cloud.

### How is Cloud Sync billed?

There are two ways to pay for sync relationships after your 14-day free trial ends. The first option is to subscribe from AWS or Azure, which enables you to pay-as-you-go or to pay annually. The second option is to purchase licenses directly from NetApp.

## Can I use Cloud Sync outside the cloud?

Yes, you can use Cloud Sync in a non-cloud architecture. The source and target can reside on-premises and so can the data broker.

Note the following key points about using Cloud Sync outside of the cloud:

- For on-premises synchronization, a private Amazon S3 bucket is available through NetApp StorageGRID.
- The data broker does need an internet connection to communicate with the Cloud Sync service.
- If you don't purchase a license directly from NetApp, you will need an AWS or Azure account for the PAYGO Cloud Sync service billing.

## How do I access Cloud Sync?

Go to the [Cloud Sync page on NetApp Cloud Central](#) and click **Start Free Trial**. Log in or sign up to Cloud Central. After you've authenticated, you're ready to get started using the Cloud Sync service.

## Supported sources and targets

The following questions related to the source and targets that are supported in a sync relationship.

### Which sources and targets does Cloud Sync support?

Cloud Sync supports many different types of sync relationships. [View the entire list](#).

### What versions of NFS and SMB does Cloud Sync support?

Cloud Sync supports NFS version 3 and later, and SMB version 1 and later.

[Learn more about sync requirements](#).

### When Amazon S3 is the target, can the data be tiered to a specific S3 storage class?

Yes, you can choose a specific S3 storage class when AWS S3 is the target:

- Standard (this is the default class)
- Intelligent-Tiering
- Standard-Infrequent Access
- One Zone-Infrequent Access
- Glacier
- Glacier Deep Archive

### What about storage tiers for Azure Blob storage?

You can choose a specific Azure Blob storage tier when a Blob container is the target:

- Hot storage
- Cool storage

# Networking

The following questions relate to networking requirements for Cloud Sync.

## What are the networking requirements for Cloud Sync?

The Cloud Sync environment requires that the data broker is connected with the source and the target through the selected protocol (NFS, SMB, EFS) or object storage API (Amazon S3, Azure Blob, IBM Cloud Object Storage).

In addition, the data broker needs an outbound internet connection over port 443 so it can communicate with the Cloud Sync service and contact a few other services and repositories.

For more details:

- [Review the networking overview](#)
- [Review the list of endpoints that the data broker contacts](#)

## Can I use a proxy server with the data broker?

Yes.

Cloud Sync supports proxy servers with or without basic authentication. If you specify a proxy server when you deploy a data broker, all HTTP and HTTPS traffic from the data broker is routed through the proxy. Note that non-HTTP traffic such as NFS or SMB can't be routed through a proxy server.

The only proxy server limitation is when using data-in-flight encryption with an NFS or Azure NetApp Files sync relationship. The encrypted data is sent over HTTPS and isn't routable through a proxy server.

# Data synchronization

The following questions relate to how data synchronization works.

## How often does synchronization occur?

The default schedule is set for daily synchronization. After the initial synchronization, you can:

- Modify the sync schedule to your desired number of days, hours, or minutes
- Disable the sync schedule
- Delete the sync schedule (no data will be lost; only the sync relationship will be removed)

## What is the minimum sync schedule?

You can schedule a relationship to sync data as often as every 1 minute.

## Does the data broker retry when a file fails to sync? Or does it timeout?

The data broker doesn't timeout when a single file fails to transfer. Instead, the data broker retries 3 times before skipping the file. The retry value is configurable in the settings for a sync relationship.

[Learn how to change the settings for a sync relationship.](#)

## What if I have a very large dataset?

If a single directory contains 600,000 files or more, [contact us](#) so we can help you configure the data broker to handle the payload. We might need to add additional memory to the data broker machine.

## Security

The following questions related to security.

### Is Cloud Sync secure?

Yes. All Cloud Sync service networking connectivity is done using [Amazon Simple Queue Service \(SQS\)](#).

All communication between the data broker and Amazon S3, Azure Blob, Google Cloud Storage, and IBM Cloud Object Storage is done through the HTTPS protocol.

If you're using Cloud Sync with on-premises (source or destination) systems, here's a few recommended connectivity options:

- An AWS Direct Connect, Azure ExpressRoute, or Google Cloud Interconnect connection, which is non-internet routed (and can only communicate with the cloud networks that you specify)
- A VPN connection between your on-premises gateway device and your cloud networks
- For extra secure data transfer with S3 buckets, Azure Blob storage, or Google Cloud Storage, an Amazon Private S3 Endpoint, Azure Virtual Network service endpoints, or Private Google Access may be established.

Any of these methods establishes a secure connection between your on-premises NAS servers and a Cloud Sync data broker.

### Is data encrypted by Cloud Sync?

- Cloud Sync supports data-in-flight encryption between source and target NFS servers. [Learn more.](#)
- Encryption is not supported with SMB.
- When an Amazon S3 bucket is the target in a sync relationship, you can choose whether to enable data encryption using AWS KMS encryption or AES-256 encryption.

## Permissions

The following questions relate to data permissions.

### Are SMB data permissions synced to the target location?

You can set up Cloud Sync to preserve access control lists (ACLs) between a source SMB share and a target SMB share. Or you can manually copy the ACLs yourself. [Learn how to copy ACLs between SMB shares.](#)

### Are NFS data permissions synced to the target location?

- NFS version 3: Cloud Sync copies the permissions and the user group owner.
- NFS version 4: You can set up Cloud Sync to preserve access control lists (ACLs) between a source NFS server and a target NFS server. [Learn how to copy ACLs between NFS servers.](#)

# Performance

The following questions relate to Cloud Sync performance.

## What does the progress indicator for a sync relationship represent?

The sync relationship shows the throughput of the data broker's network adapter. If you accelerated sync performance by using multiple data brokers, then the throughput is the sum of all traffic. This throughput refreshes every 20 seconds.

## I'm experiencing performance issues. Can we limit the number of concurrent transfers?

The data broker can sync 4 files at a time. If you have very large files (multiple TBs each), it can take a long time to complete the transfer process and performance might be impacted.

Limiting the number of concurrent transfers can help. [Contact us for help](#).

## Why am I experiencing low performance with Azure NetApp Files?

When you sync data to or from Azure NetApp Files, you might experience failures and performance issues if the disk service level is Standard.

Change the service level to Premium or Ultra to enhance the sync performance.

[Learn more about Azure NetApp Files service levels and throughput](#).

## Why am I experiencing low performance with Cloud Volumes Service for AWS?

When you sync data to or from a cloud volume, you might experience failures and performance issues if the level of performance for the cloud volume is Standard.

Change the Service level to Premium or Extreme to enhance the sync performance.

## How many data brokers are required?

When you create a new relationship, you start with a single data broker (unless you selected an existing data broker that belongs to an accelerated sync relationship). In many cases, a single data broker can meet the performance requirements for a sync relationship. If it doesn't, you can accelerate sync performance by adding additional data brokers. But you should first check other factors that can impact sync performance.

Multiple factors can impact data transfer performance. The overall sync performance might be impacted due to network bandwidth, latency, and network topology, as well as the data broker VM specs and storage system performance. For example, a single data broker in a sync relationship can reach 100 MB/s, while disk throughput on the target might only allow 64 MB/s. As a result, the data broker keeps trying to copy the data, but the target can't meet the performance of the data broker.

So be sure to check the performance of your networking and the disk throughput on the target.

Then you can consider accelerating sync performance by adding an additional data broker to share the load of that relationship. [Learn how to accelerate sync performance](#).

# Deleting things

The following questions relate to deleting sync relationships and data from sources and targets.

## What happens if I delete my Cloud Sync relationship?

Deleting a relationship stops all future data syncs and terminates payment. Any data that was synced to the target remains as-is.

## What happens if I delete something from my source server? Is it removed from the target too?

By default, if you have an active sync relationship, the item deleted on the source server is not deleted from the target during the next synchronization. But there is an option in the sync settings for each relationship, where you can define that Cloud Sync will delete files in the target location if they were deleted from the source.

[Learn how to change the settings for a sync relationship.](#)

## What happens if I delete something from my target? Is it removed from my source too?

If an item is deleted from the target, it will not be removed from the source. The relationship is one-way—from source to target. On the next sync cycle, Cloud Sync compares the source to the target, identifies that the item is missing, and Cloud Sync copies it again from the source to the target.

# Troubleshooting

[NetApp Knowledgebase: Cloud Sync FAQ: Support and Troubleshooting](#)

## Data broker deep dive

The following question relates to the data broker.

## Can you explain the architecture of the data broker?

Sure. Here are the most important points:

- The data broker is a node.js application running on a Linux host.
- Cloud Sync deploys the data broker as follows:
  - AWS: From an AWS CloudFormation template
  - Azure: From Azure Resource Manager
  - Google: From Google Cloud Deployment Manager
  - If you use your own Linux host, you need to manually install the software
- The data broker software automatically upgrades itself to the latest version.
- The data broker uses AWS SQS as a reliable and secure communication channel and for control and monitoring. SQS also provides a persistency layer.
- You can add additional data brokers to a relationship to increase transfer speed and add high availability. There is service resiliency if one data broker fails.



# How to get help and find more information

NetApp has several resources available for Cloud Sync assistance. You will need to register your support serial number to activate support before you can contact NetApp technical support.

## Self-support resources

Several resources are available to help you find the answers to your questions.

- [Frequently asked questions](#)

This FAQ can help if you're just looking for a quick answer to a question.

- [Cloud Sync knowledge base](#)

Search through the Cloud Sync knowledge base to find a number of useful articles.

- [NetApp Community forum: Cloud Data Services](#)

In this forum, use labels and filters to look at Cloud Sync topics. If you'd like to ask a question, click **Register** in the upper-right corner to sign up.

- [NetApp Cloud Central](#)

Find more information about Cloud Sync, as well as additional NetApp products and solutions for the cloud.


- [NetApp Product Documentation](#)

Search NetApp product documentation for instructions, resources, and answers.

## Chatting with NetApp cloud experts

Our inline chat is available for product or pricing related questions, or just general feedback with our NetApp cloud experts. While these experts can answer your questions in a timely manner, they aren't NetApp support personnel and there isn't a strict SLA. It's best to register for support and open a web ticket for technical support related issues.

Just click the chat icon available in the lower right of the interface and ask your question.

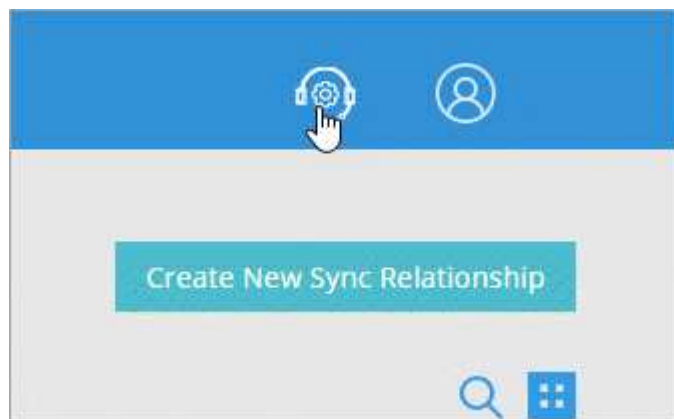


# Activating NetApp support

Registering your support serial number will activate support entitlement and enable you to contact NetApp technical support for help.



## Steps

1. In Cloud Sync, click the support icon in the upper-right of the interface. This is where your support serial number is located.



2. Follow these guidelines if you don't have a NetApp Support Site (NSS) account:
  - If you are a current NetApp customer with existing NetApp products or services, [create your NSS account here](#).
  - If you are new to NetApp with no previous products or services, then go to [Cloud Data Services Support Registration](#) to register your new Cloud Sync serial number first, then create your NSS account.
3. When you have a NetApp Support Site account, enter your user name and password and click **Activate Support**.

Note that your account can't be a temp or guest level account.


 Support

---

### Account Details

---

Account ID



---

NetApp Serial Number

91548350426159871855

---

### NetApp Support Site Credentials


---

Enter your NetApp Support Site (NSS) credentials to activate support for this subscription serial number.

NSS User Name

NSS Password

Activate Support





If you don't have NSS credentials, or if this is your first NetApp product/service, [click here to register](#).

The registration process takes a couple minutes. If you experience any issues registering from within Cloud Sync, [try registering from this NetApp website](#) instead.

## Result

Once the activation is complete, you will see the following:

 Support

User Details

Account ID

NetApp Serial Number


91548350426159871855

NetApp Support Site

☒ Your support is active

ASUP (AutoSupport)

Send AutoSupport to NetApp

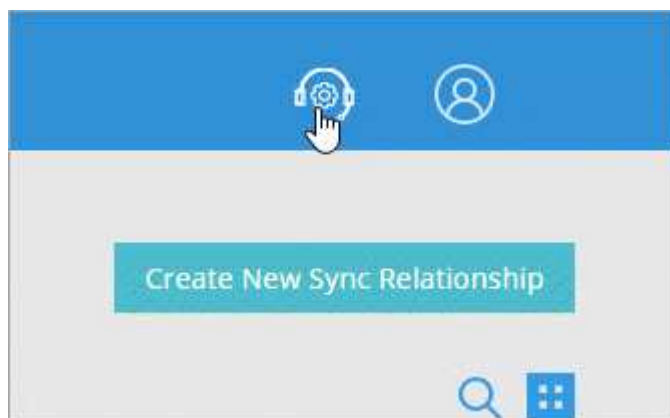
 Send ASUP

## Contacting NetApp support

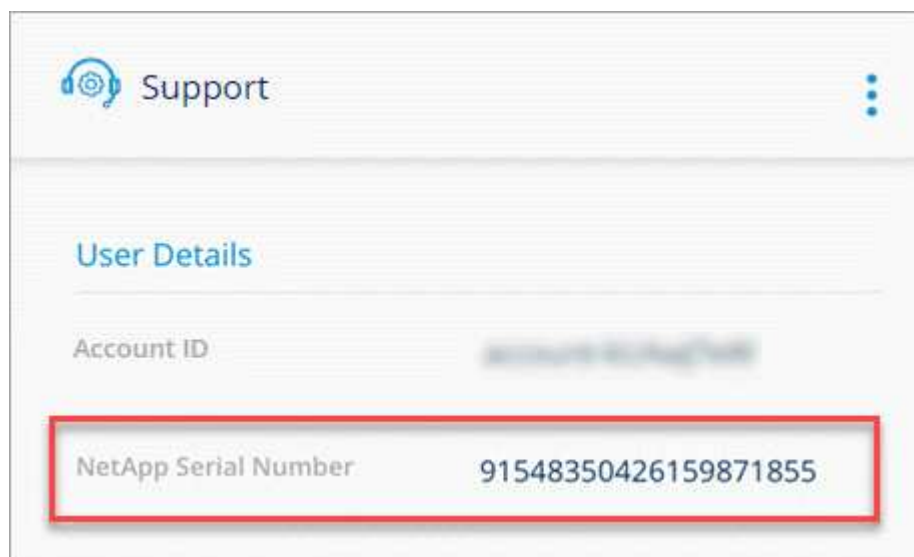
Contact NetApp technical support through our web ticketing system or by phone. Web tickets will get a call back. For more efficient support, open a web ticket first. If urgent, call NetApp using the case number from the web ticket.

### Steps

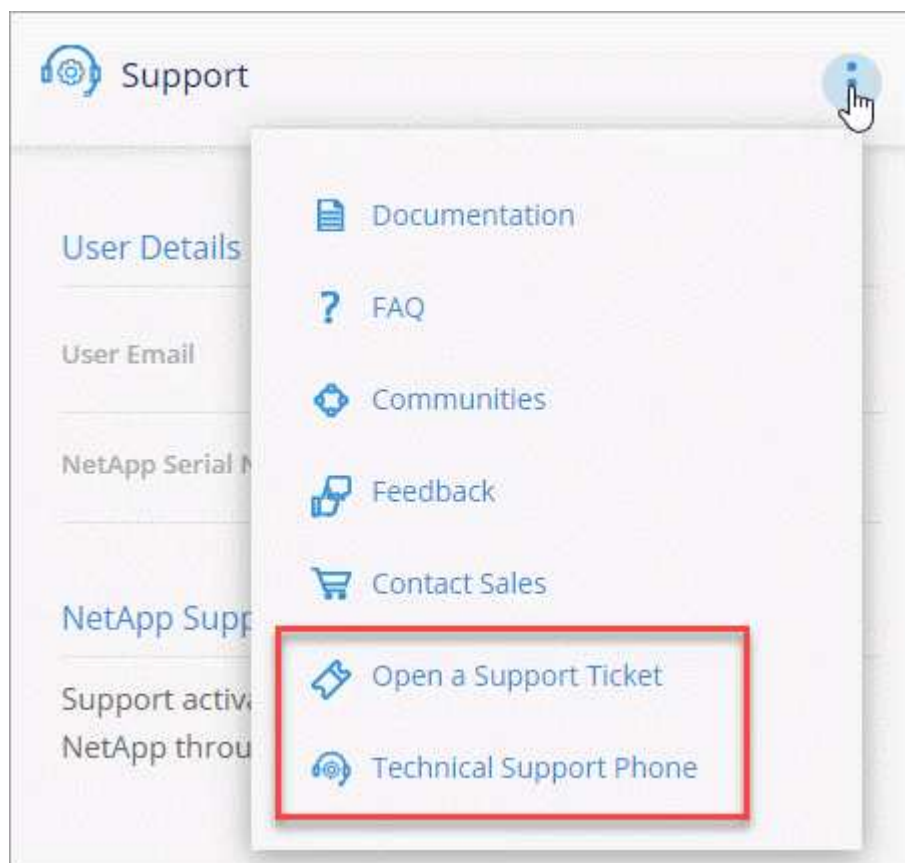
1. In Cloud Sync, click the support icon in the upper-right of the interface.



2. Make note of your NetApp serial number, which you'll need to provide to technical support.



3. Click the menu and select one of the available options to get support.



Here are the links for contacting support by web ticketing and phone:

- [Create Support Ticket](#)
- [Phone Support](#)

# Sending AutoSupport messages to NetApp

To assist in troubleshooting, Cloud Sync can remotely trigger AutoSupport messages from each data broker back to NetApp. This action is triggered only when you click the **Send ASUP** button. The data broker instance must have outbound HTTPS internet access to support.netapp.com to transmit troubleshooting information to NetApp.

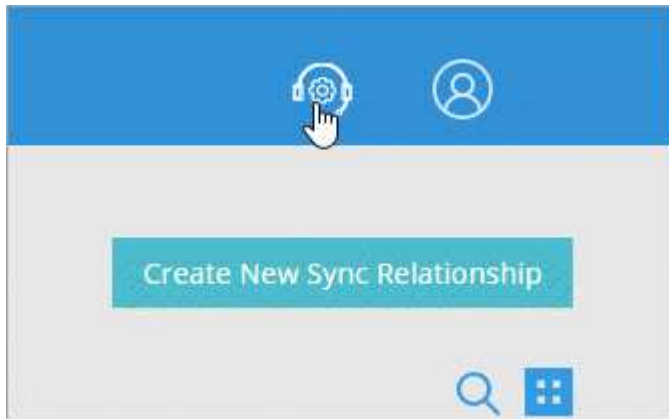
It's a good practice to send an AutoSupport message when creating a support ticket or when calling NetApp support. NetApp should be able to find your AutoSupport with the "915" serial number provided.



Access to *fedoraproject.org:443* is needed from the data broker virtual machine during data broker installation and software updates. This endpoint is contacted to install 7z, which is needed to send AutoSupport messages. [Learn more about endpoint requirements.](#)

## Steps

1. In Cloud Sync, click the support icon in the upper-right of the interface.



2. Click **Send ASUP**.



The **Send ASUP** button appears only after you [activate NetApp support](#).

# Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

## Copyright

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

## Trademarks

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

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

## Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/us/media/patents-page.pdf>

## Privacy policy

<https://www.netapp.com/us/legal/privacypolicy/index.aspx>

## Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

[Notice for Cloud Sync](#)

## 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.