NetApp Copy and Sync documentation

NetApp Copy and Sync

NetApp
December 16, 2025

Table of Contents

NetApp Copy and Sync documentation

Release notes

What’s new with NetApp Copy and Sync

Learn what’s new in NetApp Copy and Sync.

06 October 2025

BlueXP copy and sync is now NetApp Copy and Sync

BlueXP copy and sync has been renamed to NetApp Copy and Sync.

BlueXP is now NetApp Console

The NetApp Console, built on the enhanced and restructured BlueXP foundation, provides centralized management of NetApp storage and NetApp Data Services across on-premises and cloud environments at enterprise grade—delivering real-time insights, faster workflows, and simplified administration, that is highly secure and compliant.

For details on what’s changed, see the NetApp Console release notes.

02 February 2025

New OS support for data broker

The data broker is now supported on hosts running Red Hat Enterprise 9.4, Ubuntu 23.04, and Ubuntu 24.04.

View Linux host requirements.

27 October 2024

Bug fixes

We updated NetApp Copy and Sync and the data broker to fix a few bugs. The new data broker version is 1.0.56.

16 September 2024

Bug fixes

We updated NetApp Copy and Sync and the data broker to fix a few bugs. The new data broker version is 1.0.55.

11 August 2024

Bug fixes

We updated NetApp Copy and Sync and the data broker to fix a few bugs. The new data broker version is 1.0.54.

14 July 2024

Bug fixes

We updated Copy and Sync and the data broker to fix a few bugs. The new data broker version is 1.0.53.

02 June 2024

Bug fixes

NetApp Copy and Sync was updated to fix a few bugs. The data broker was also updated to apply security updates. The new data broker version is 1.0.52.

08 April 2024

Support for RHEL 8.9

The data broker is now supported on hosts running Red Hat Enterprise Linux 8.9.

View Linux host requirements.

11 February 2024

Filter directories by regex

Users now have the option to filter directories using regex.

Learn more about the **Exclude Directories** feature.

26 November 2023

Cold storage class support for Azure Blob

The cold storage Azure Blob tier is now available when creating a sync relationship.

Learn more about creating a sync relationship.

Support for Tel Aviv region in AWS data brokers

Tel Aviv is now a supported region when creating a data broker in AWS.

Learn more about creating a data broker in AWS.

Update to node version for data brokers

All new data brokers will now use node version 21.2.0. Data brokers that are not compatible with this update, such as CentOS 7.0 and Ubuntu Server 18.0, will no longer work with NetApp Copy and Sync.

03 September 2023

Exclude files by regex

Users now have the option to exclude files using regex.

Learn more about the **Exclude File Extensions** feature.

Add S3 keys when creating Azure data broker

Users may now add AWS S3 access keys and secret keys when creating an Azure data broker.

Learn more about creating a data broker in Azure.

06 August 2023

Use existing Azure security groups when creating a data broker

Users now have the option to use existing Azure security groups when creating a data broker.

The service account used when creating the data broker must have these permissions:

- "Microsoft.Network/networkSecurityGroups/securityRules/read"
- "Microsoft.Network/networkSecurityGroups/read"

Learn more about creating a data broker in Azure.

Encrypt data when syncing to Google Storage

Users now have the option to specify a customer-managed encryption key when creating a sync relationship with a Google Storage bucket as the target. You can manually input your key or choose from a list of your keys in a single region.

The service account used when creating the data broker must have these permissions:

- cloudkms.cryptoKeys.list
- cloudkms.keyRings.list

Learn more about Google Cloud Storage bucket requirements.

09 July 2023

Remove multiple sync relationships at once

Users are now able to delete more than one sync relationship at a time in the UI.

Learn more about deleting sync relelationships.

Copy only ACL

Users now have additional options for copying ACL information in CIF and NFS relationships. When creating or managing a sync relationship, you can copy files only, copy ACL information only, or copy files and ACL information.

Learn more about copying ACLs.

Updated to Node.js 20

Copy and Sync has updated to Node.js 20. All available data brokers will be updated. Operation systems incompatible with this update cannot be installed, and incompatible existing systems may experience performance issues.

11 June 2023

Support automatic abort by minutes

Active syncs that have not completed can now be aborted after fifteen minutes using the **Sync Timeout** feature.

Learn more about the Sync Timeout setting.

Copy access time metadata

In relationships including a file system, the **Copy for Objects** feature now copies access time metadata.

Learn more about the Copy For Objects setting.

08 May 2023

Hard link capabilities

Users can now include hard links for syncs involving unsecured NFS to NFS relationships.

Learn more about the File Types setting.

Ability to add user certificate for data brokers in secure NFS relationships

Users are now able to set their own certificate for the target data broker when creating a secure NFS relationship. They will need to set a server name and provide a private key and certificate ID when doing so. This feature is available for all data brokers.

Extended exclusion period for recently modified files

Users can now exclude files that were modified up to 365 days before the scheduled sync.

Learn more about the Recently Modified Files setting.

Filter relationships in UI by relationship ID

Those using the RESTful API can now filter relationships using relationship IDs.

Learn more about using the RESTful API with NetApp Copy and Sync.

Learn more about the Exclude Directories setting.

02 April 2023

Additional support for Azure Data Lake Storage Gen2 relationships

You can now create sync relationships with Azure Data Lake Storage Gen2 as a source and target with the following:

- Azure NetApp Files

- Amazon FSx for ONTAP
- Cloud Volumes ONTAP
- On-Prem ONTAP

Learn more about supported sync relationships.

Filter directories by full path

In addition to filtering directories out by name, you can now filter directories by their full path.

Learn more about the Exclude Directories setting.

07 March 2023

EBS Encryption for AWS data brokers

You can now encrypt AWS data broker volumes using a KMS key from your account.

Learn more about creating a data broker in AWS.

05 Feb 2023

Additional support for Azure Data Lake Storage Gen2, ONTAP S3 Storage, and NFS

Cloud Sync now supports additional sync relationships for ONTAP S3 Storage and NFS:

- ONTAP S3 Storage to NFS
- NFS to ONTAP S3 Storage

Cloud Sync also has additional support for Azure Data Lake Storage Gen2 as both a source and target to:

- NFS server
- SMB server
- ONTAP S3 Storage
- StorageGRID
- IBM Cloud Object Storage

Learn more about supported sync relationships.

Upgrade to Amazon Web Services data broker operating system

The operating system for AWS data brokers has been upgraded to the Amazon Linux 2022.

Learn more about the data broker instance in AWS.

03 Jan 2023

Show data broker local configuration on UI

There is now a **Show Configuration** option that allows users to view the local configuration of each data broker on the UI.

Learn more about managing data broker groups.

Upgrade to Azure and Google Cloud data broker operating system

The operating system for data brokers in Azure and Google Cloud has been upgraded to the Rocky Linux 9.0.

Learn more about the data broker instance in Azure.

Learn more about the data broker instance in Google Cloud.

11 Dec 2022

Filter directories by name

A new **Exclude Directory Names** setting is now available for sync relationships. Users can filter out a maximum of 15 directory names from their sync. The .copy-offload, .snapshot, ~snapshot directories are excluded by default.

Learn more about the Exclude Directory Names setting.

Additional Amazon S3 and ONTAP S3 Storage support

Cloud Sync now supports additional sync relationships for AWS S3 and ONTAP S3 Storage:

- AWS S3 to ONTAP S3 Storage
- ONTAP S3 Storage to AWS S3

Learn more about supported sync relationships.

30 Oct 2022

Continuous sync from Microsoft Azure

The Continuous Sync setting is now supported from a source Azure storage bucket to a cloud storage using an Azure data broker.

After the initial data sync, Cloud Sync listens for changes on the source Azure storage bucket and continuously syncs any changes to the target as they occur. This setting is available when syncing from an Azure storage bucket to Azure Blob storage, CIFS, Google Cloud Storage, IBM Cloud Object Storage, NFS, and StorageGRID.

The Azure data broker needs a custom role and the following permissions to use this setting:

```
'Microsoft.Storage/storageAccounts/read',
'Microsoft.EventGrid/systemTopics/eventSubscriptions/write',
'Microsoft.EventGrid/systemTopics/eventSubscriptions/read',
'Microsoft.EventGrid/systemTopics/eventSubscriptions/delete',
'Microsoft.EventGrid/systemTopics/eventSubscriptions/getFullUrl/action',
'Microsoft.EventGrid/systemTopics/eventSubscriptions/getDeliveryAttributes
/action',
'Microsoft.EventGrid/systemTopics/read',
'Microsoft.EventGrid/systemTopics/write',
'Microsoft.EventGrid/systemTopics/delete',
'Microsoft.EventGrid/eventSubscriptions/write',
'Microsoft.Storage/storageAccounts/write'
```

Learn more about the Continuous Sync setting.

04 Sept 2022

Additional Google Drive support

- Cloud Sync now supports additional sync relationships for Google Drive:
    - Google Drive to NFS servers
    - Google Drive to SMB servers
- You can also generate reports for sync relationships that include Google Drive.

  Learn more about reports.

Continuous sync enhancement

You can now enable the Continuous Sync setting on the following types of sync relationships:

- S3 bucket to an NFS server
- Google Cloud Storage to an NFS server

Learn more about the Continuous Sync setting.

Email notifications

You can now receive Cloud Sync notifications by email.

In order to receive the notifications by email, you'll need to enable the **Notifications** setting on the sync relationship and then configure the Alerts and Notification settings in the NetApp Console.

Learn how to set up notifications.

31 July 2022

Google Drive

You can now sync data from an NFS server or SMB server to Google Drive. Both "My Drive" and "Shared Drives" are supported as targets.

Before you can create a sync relationship that includes Google Drive, you need to set up a service account that has the required permissions and a private key. Learn more about Google Drive requirements.

View the list of supported sync relationships.

Additional Azure Data Lake support

Cloud Sync now supports additional sync relationships for Azure Data Lake Storage Gen2:

- Amazon S3 to Azure Data Lake Storage Gen2
- IBM Cloud Object Storage to Azure Data Lake Storage Gen2
- StorageGRID to Azure Data Lake Storage Gen2

View the list of supported sync relationships.

New ways to set up sync relationships

We've added additional ways to set up sync relationships directly from the NetApp Console's Systems page.

Drag and drop

You can now set up a sync relationship from the Systems page by dragging and dropping one system on top of another.

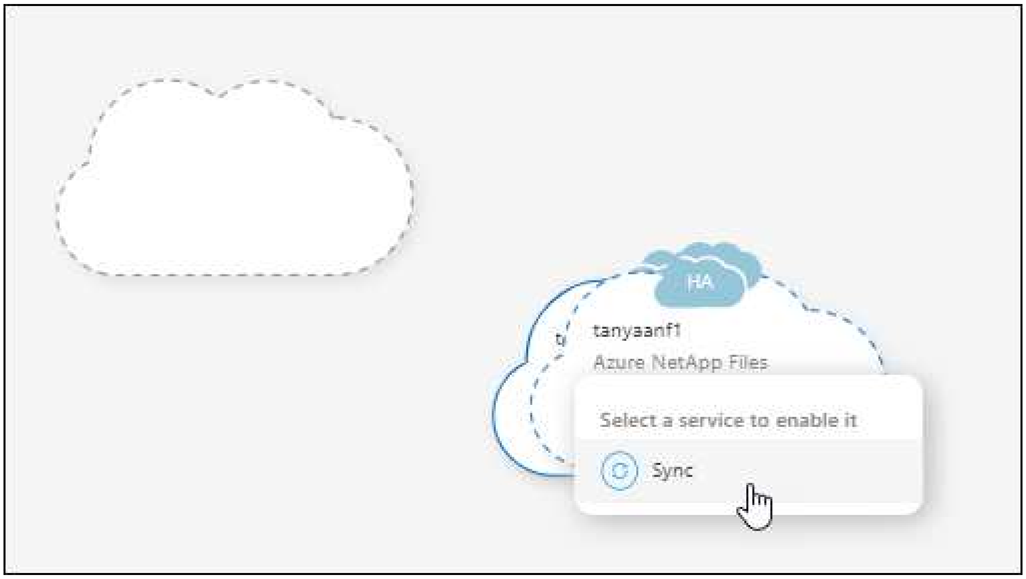

Right panel setup

You can now set up a sync relationship for Azure Blob storage or for Google Cloud Storage by selecting the system from the Systems page and then selecting the sync option from the right panel.

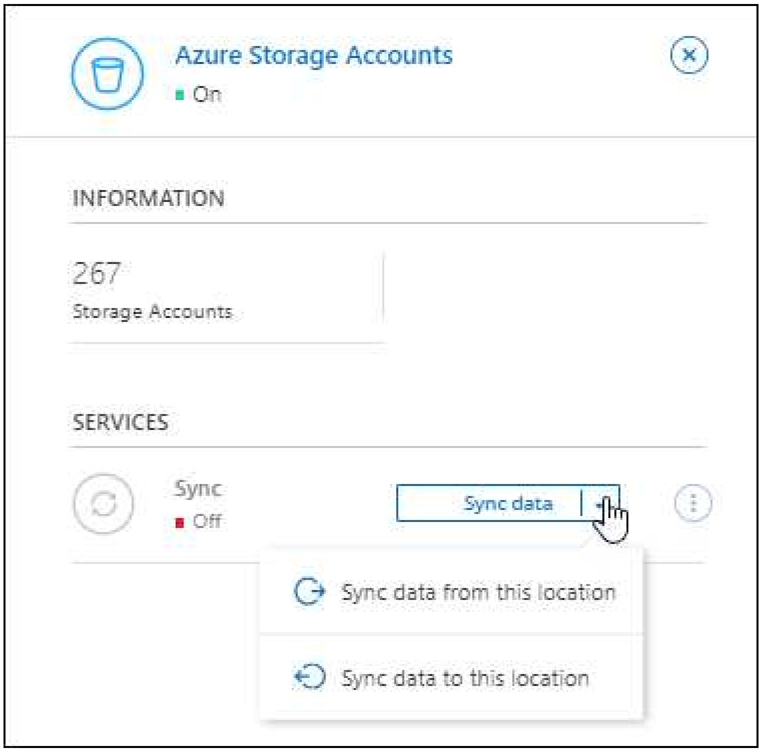

03 July 2022

Support for Azure Data Lake Storage Gen2

You can now sync data from an NFS server or SMB server to Azure Data Lake Storage Gen2.

When creating a sync relationship that includes Azure Data Lake, you need to provide Cloud Sync with the storage account connection string. It must be a regular connection string, not a shared access signature (SAS).

View the list of supported sync relationships.

Continuous sync from Google Cloud Storage

The Continuous Sync setting is now supported from a source Google Cloud Storage bucket to a cloud storage target.

After the initial data sync, Cloud Sync listens for changes on the source Google Cloud Storage bucket and continuously syncs any changes to the target as they occur. This setting is available when syncing from a Google Cloud Storage bucket to S3, Google Cloud Storage, Azure Blob storage, StorageGRID, or IBM Storage.

The service account associated with your data broker needs the following permissions to use this setting:

```
- pubsub.subscriptions.consume
- pubsub.subscriptions.create
- pubsub.subscriptions.delete
- pubsub.subscriptions.list
- pubsub.topics.attachSubscription
- pubsub.topics.create
- pubsub.topics.delete
- pubsub.topics.list
- pubsub.topics.setIamPolicy
- storage.buckets.update
```

Learn more about the Continuous Sync setting.

New Google Cloud region support

The Cloud Sync data broker is now supported in the following Google Cloud regions:

- Columbus (us-east5)
- Dallas (us-south1)
- Madrid (europe-southwest1)
- Milan (europe-west8)
- Paris (europe-west9)

New Google Cloud machine type

The default machine type for the data broker in Google Cloud is now n2-standard-4.

06 June 2022

Continuous sync

A new setting enables you to continuously sync changes from a source S3 bucket to a target.

After the initial data sync, Cloud Sync listens for changes on the source S3 bucket and continuously syncs any changes to the target as they occur. There’s no need to rescan the source at scheduled intervals. This setting is available only when syncing from an S3 bucket to S3, Google Cloud Storage, Azure Blob storage, StorageGRID, or IBM Storage.

Note that the IAM role associated with your data broker will need the following permissions to use this setting:

```
"s3:GetBucketNotification",
"s3:PutBucketNotification"
```

These permissions are automatically added to any new data brokers that you create.

Learn more about the Continuous Sync setting.

Show all ONTAP volumes

When you create a sync relationship, Cloud Sync now displays all volumes on a source Cloud Volumes ONTAP system, on-premises ONTAP cluster, or FSx for ONTAP file system.

Previously, Cloud Sync would only display the volumes that matched the selected protocol. Now all of the volumes display, but any volumes that don’t match the selected protocol or that don’t have a share or export are greyed out and not selectable.

Copying tags to Azure Blob

When you create a sync relationship where Azure Blob is the target, Cloud Sync now enables you to copy tags to the Azure Blob container:

- On the **Settings** page, you can use the **Copy for Objects** setting to copy tags from the source to the Azure Blob container. This is in addition to copying metadata.
- On the **Tags/Metadata** page, you can specify Blob index tags to set on the objects that are copied to the Azure Blob container. Previously, you could only specify relationship metadata.

These options are supported when Azure Blob is the target and the source is either Azure Blob or an S3-compatible endpoint (S3, StorageGRID, or IBM Cloud Object Storage).

01 May 2022

Sync timeout

A new **Sync Timeout** setting is now available for sync relationships. This setting enables you to define whether Cloud Sync should cancel a data sync if the sync hasn’t completed in the specified number of hours or days.

Learn more about changing the settings for a sync relationship.

Notifications

A new **Notifications** setting is now available for sync relationships. This setting enables you to choose whether to receive Cloud Sync notifications in the NetApp Console’s Notification Center. You can enable notifications for successful data syncs, failed data syncs, and canceled data syncs.

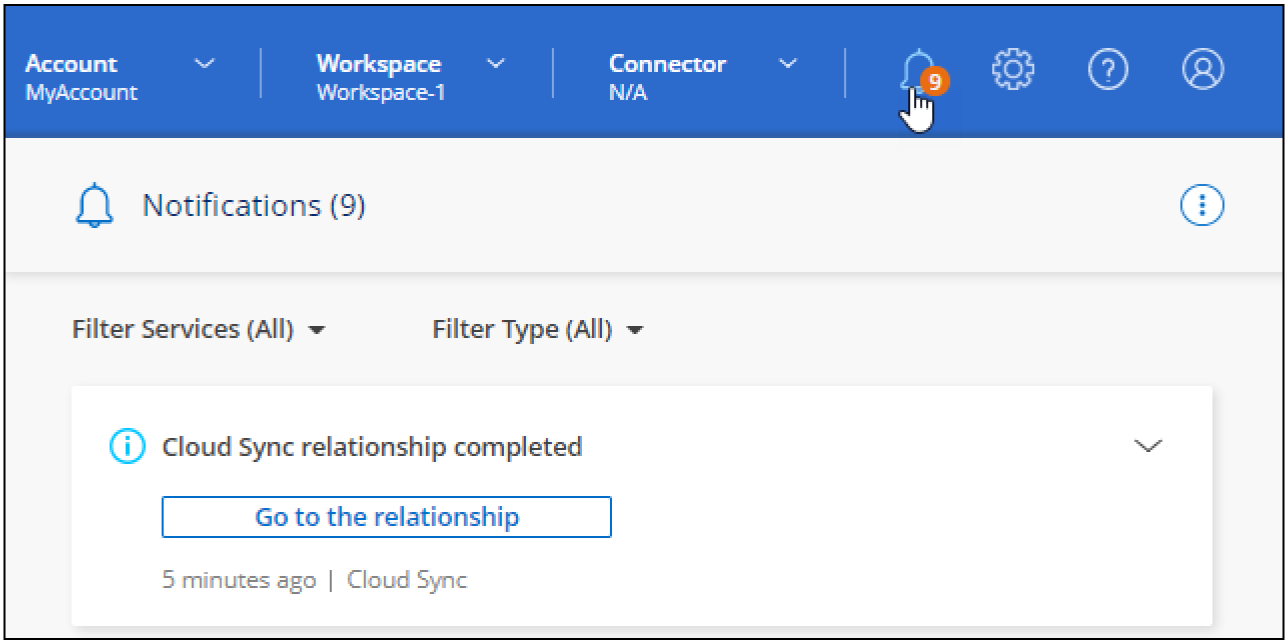

Notifications (9)

Cloud Sync relationship completed

5 minutes ago | Cloud Sync

Learn more about changing the settings for a sync relationship.

03 April 2022

Data broker group enhancements

We made several enhancements to data broker groups:

- You can now move a data broker to a new or existing group.
- You can now update the proxy configuration for a data broker.
- Finally, you can also delete data broker groups.

Learn how to manage data broker groups.

Dashboard filter

You can now filter the contents of the Sync Dashboard to more easily find sync relationships that match a certain status. For example, you can filter on sync relationships that have a failed status

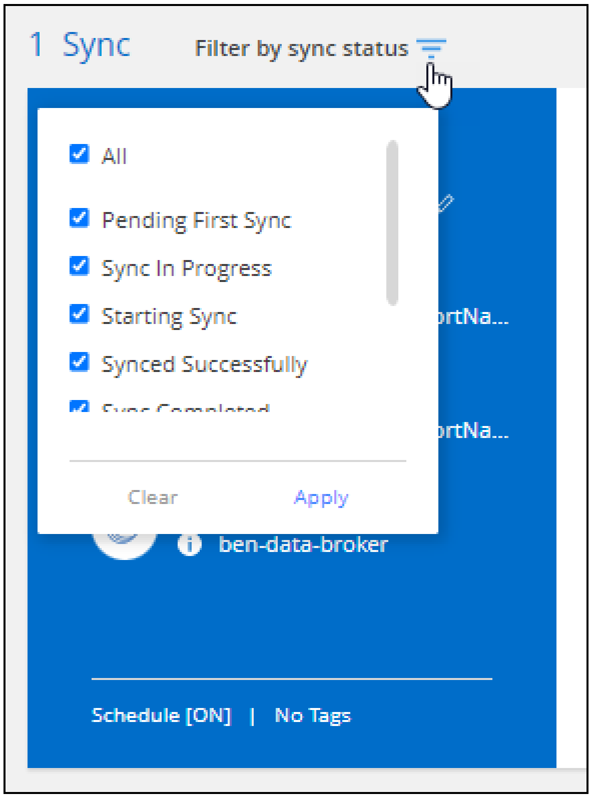

03 March 2022

Sorting in the dashboard

You now sort the dashboard by sync relationship name.

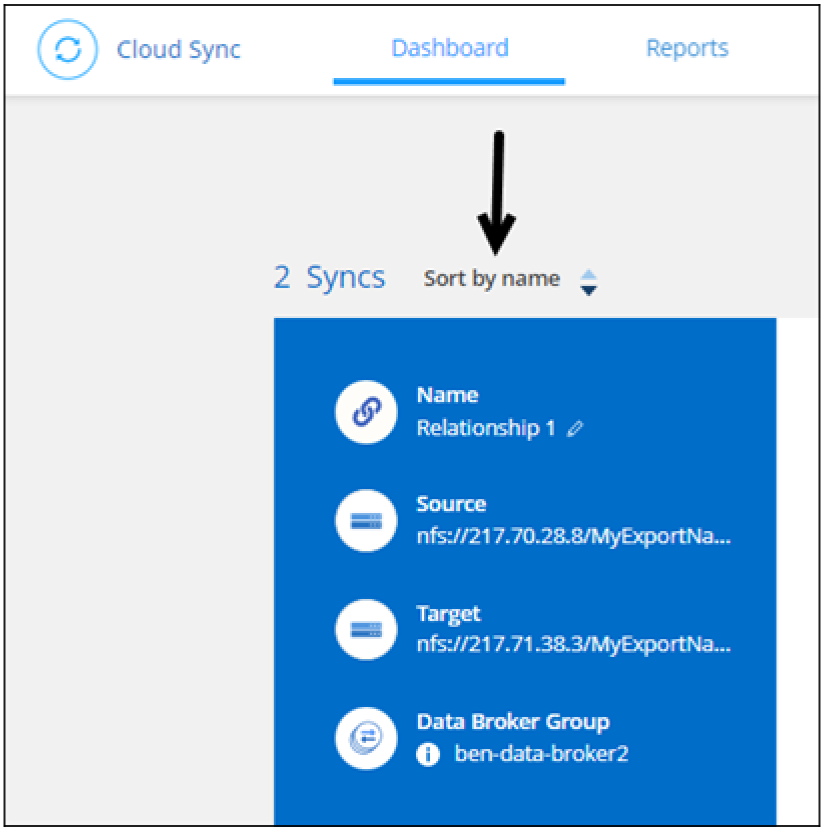

Enhancement to Data Sense integration

In the previous release, we introduced Cloud Sync integration with Cloud Data Sense. In this update, we enhanced the integration by making it easier to create the sync relationship. After you initiate a data sync from Cloud Data Sense, all of the source information is contained in a single step and only requires you to enter a few key details.

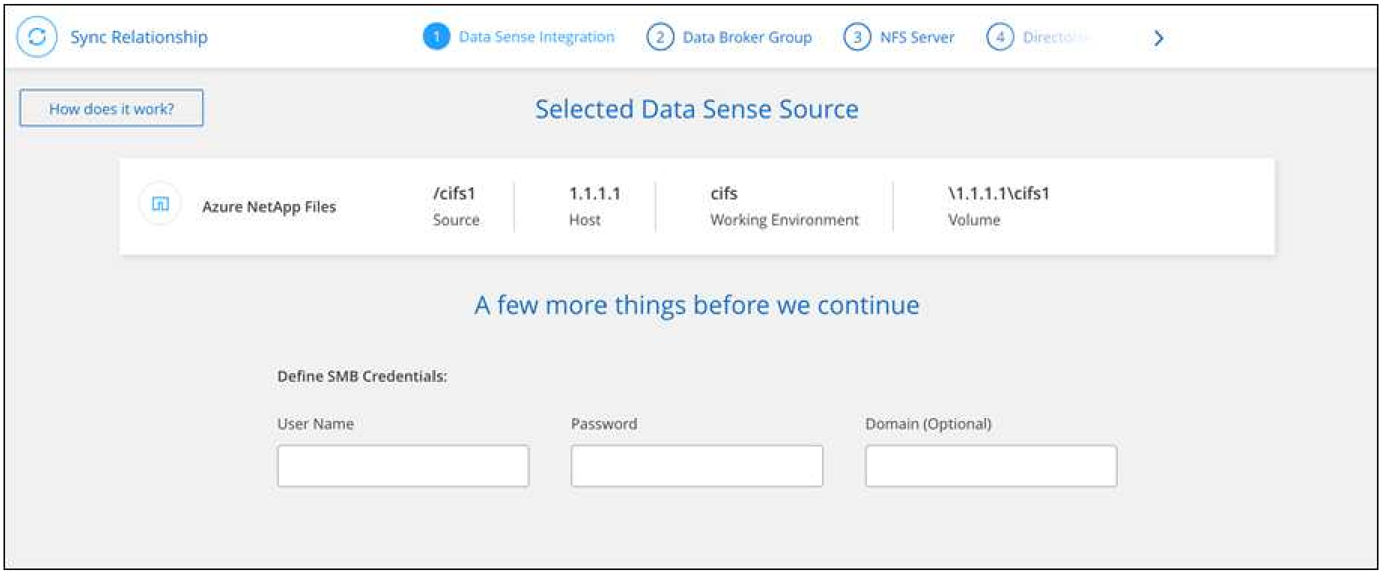

Selected Data Sense Source

Azure NetApp Files	/cifs1	1.1.1.1	cifs	\1.1.1.1\cifs1
	Source	Host	Working Environment	Volume

A few more things before we continue

Define SMB Credentials:

User Name

Password

Domain (Optional)

06 February 2022

Enhancement to data broker groups

We changed how you interact with data brokers by emphasizing data broker *groups*.

For example, when you create a new sync relationship, you select the data broker *group* to use with the relationship, rather than a specific data broker.

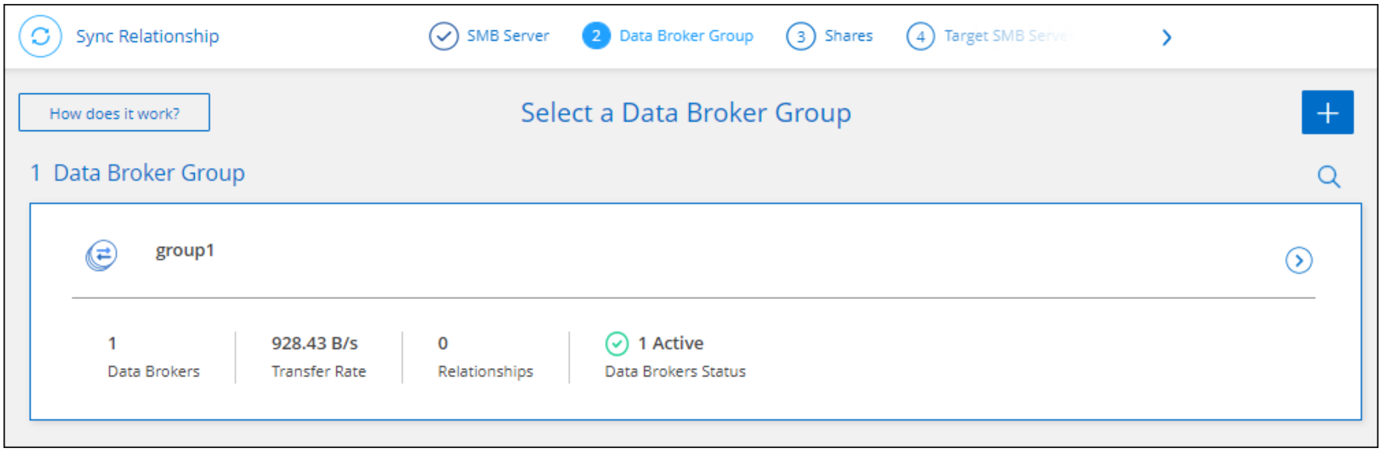

Select a Data Broker Group

1 Data Broker Group

group1

1	928.43 B/s	0	1 Active
Data Brokers	Transfer Rate	Relationships	Data Brokers Status

In the **Manage Data Brokers** tab, we also show the number of sync relationships that a data broker group is managing.

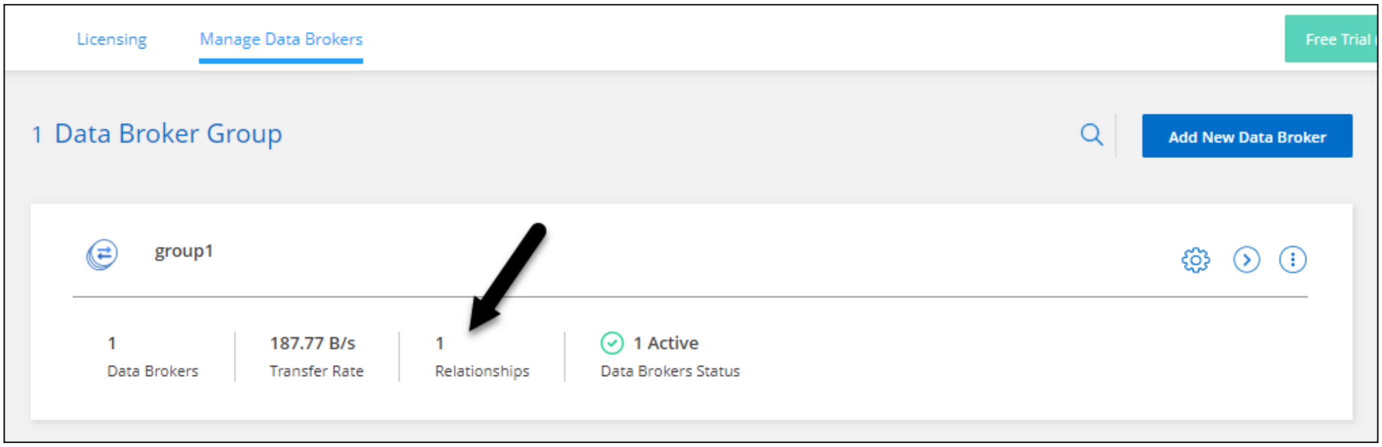

1 Data Brokers

187.77 B/s Transfer Rate

1 Relationships

1 Active Data Brokers Status

Download PDF reports

You can now download a PDF of a report.

Learn more about reports.

02 January 2022

New Box sync relationships

Two new sync relationships are supported:

- Box to Azure NetApp Files
- Box to Amazon FSx for ONTAP

View the list of supported sync relationships.

Relationship names

You can now provide a meaningful name to each of your sync relationships to more easily identify the purpose of each relationship. You can add the name when you create the relationship and any time after.

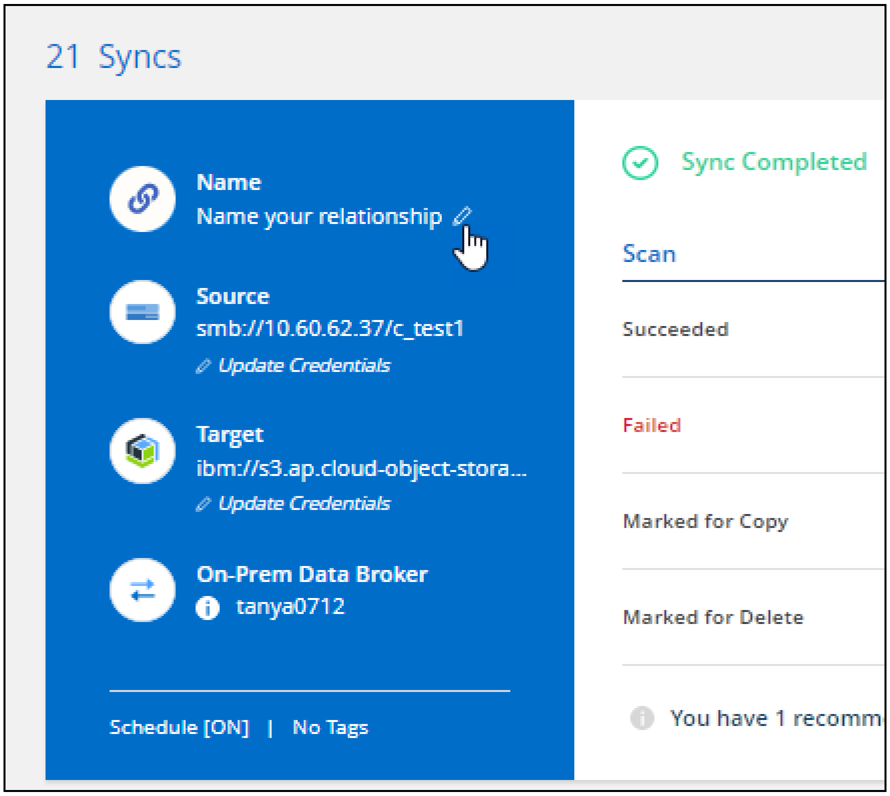
21 Syncs

Name
Name your relationship

Source
smb://10.60.62.37/c_test1
Update Credentials

Target
ibm://s3.ap.cloud-object-stora...
Update Credentials

On-Prem Data Broker
tanya0712

Schedule [ON] | No Tags

Sync Completed

Scan

Succeeded

Failed

Marked for Copy

Marked for Delete

You have 1 recomm

S3 private links

When you sync data to or from Amazon S3, you can choose whether to use an S3 private link. When the data broker copies data from the source to the target, it goes through the private link.

Note that the IAM role associated with your data broker will need the following permission to use this feature:

```
"ec2:DescribeVpcEndpoints"
```

This permission is automatically added to any new data brokers that you create.

Glacier Instant Retrieval

You can now choose the *Glacier Instant Retrieval* storage class when Amazon S3 is the target in a sync relationship.

ACLs from object storage to SMB shares

Cloud Sync now supports copying ACLs from object storage to SMB shares. Previously, we only supported copying ACLs from an SMB share to object storage.

SFTP to S3

Creating a sync relationship from SFTP to Amazon S3 is now supported in the user interface. This sync relationship was previously supported with the API only.

Table view enhancement

We redesigned the table view on the Dashboard for ease of use. If you select **More info**, Cloud Sync filters the dashboard to show you more information about that specific relationship.

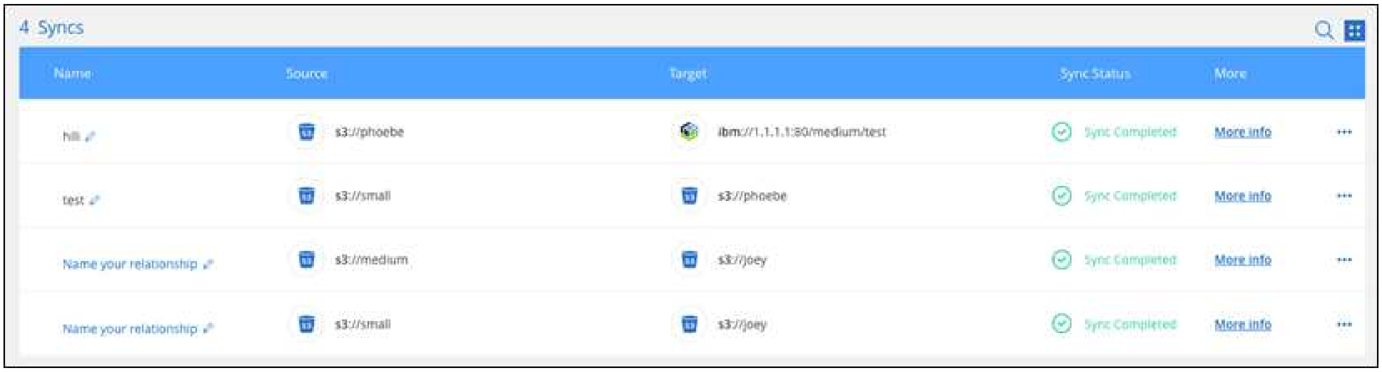

4 Syncs

Name	Source	Target	Sync Status	More	
hill	s3://phoebe	ibm://1.1.1.1:80/medium/test	Sync Completed	More info	•••
test	s3://small	s3://phoebe	Sync Completed	More info	•••
Name your relationship	s3://medium	s3://joey	Sync Completed	More info	•••
Name your relationship	s3://small	s3://joey	Sync Completed	More info	•••

Support for Jarkarta region

Cloud Sync now supports deploying the data broker in the AWS Asia Pacific (Jakarta) region.

28 November 2021

ACLs from SMB to object storage

Cloud Sync can now copy access control lists (ACLs) when setting up a sync relationship from a source SMB share to object storage (except for ONTAP S3).

Cloud Sync doesn’t support copying ACLs from object storage to SMB shares.

Learn how to copy ACLs from an SMB share.

Update licenses

You can now update Cloud Sync licenses that you extended.

If you extended a Cloud Sync license that you purchased from NetApp, you can add the license again to refresh the expiration date.

Learn how to update a license.

Update Box credentials

You can now update the Box credentials for an existing sync relationship.

Learn how to update credentials.

31 October 2021

Box support

Box support is now available in the Cloud Sync user interface as a preview.

Box can be the source or target in several types of sync relationships. View the list of supported sync relationships.

Date Created setting

When an SMB server is the source, a new sync relationship setting called *Date Created* enables you to sync files that were created after a specific date, before a specific date, or between a specific time range.

Learn more about Cloud Sync settings.

04 October 2021

Additional Box support

Cloud Sync now supports additional sync relationships for Box when using the Cloud Sync API:

- Amazon S3 to Box
- IBM Cloud Object Storage to Box
- StorageGRID to Box
- Box to an NFS server
- Box to an SMB server

Learn how to set up a sync relationship using the API.

Reports for SFTP paths

You can now create a report for SFTP paths.

02 September 2021

Support for FSx for ONTAP

You can now sync data to or from an Amazon FSx for ONTAP file system.

- Learn about Amazon FSx for ONTAP
- View supported sync relationships
- Learn how to create a sync relationship for Amazon FSx for ONTAP

01 August 2021

Update credentials

Cloud Sync now enables you to update the data broker with the latest credentials of the source or target in an existing sync relationship.

This enhancement can help if your security policies require you to update credentials on a periodic basis. Learn how to update credentials.

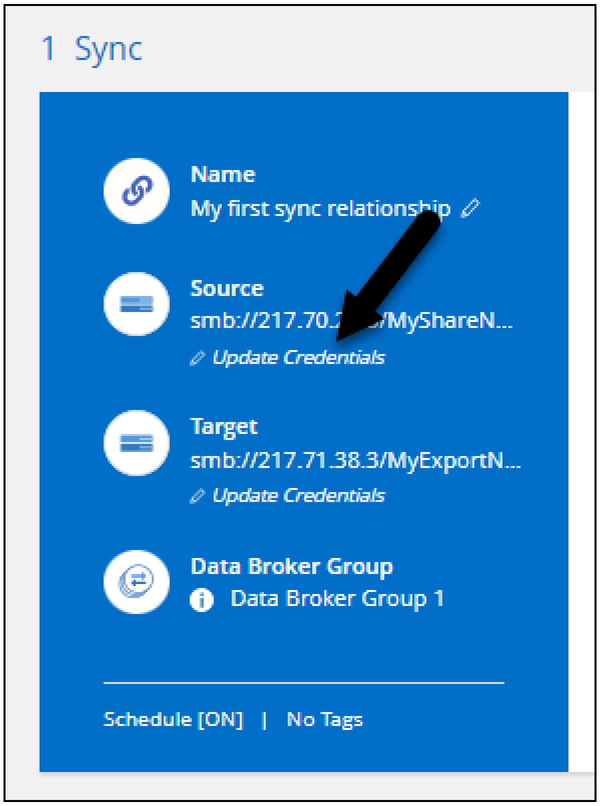

1 Sync

Name
My first sync relationship

Source
smb://217.70.[illegible]/MyShareN...
Update Credentials

Target
smb://217.71.38.3/MyExportN...
Update Credentials

Data Broker Group
Data Broker Group 1

Schedule [ON] | No Tags

Tags for object storage targets

When creating a sync relationship, you can now add tags to the object storage target in a sync relationship.

Adding tags is supported with Amazon S3, Azure Blob, Google Cloud Storage, IBM Cloud Object Storage, and StorageGRID.

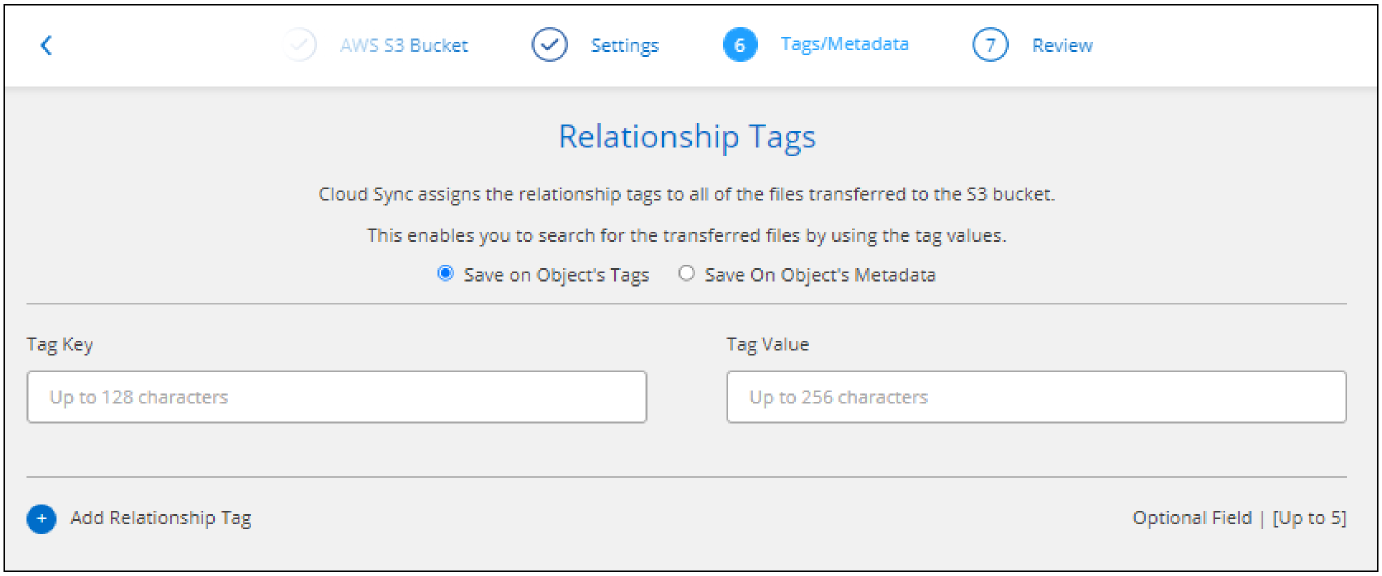

Relationship Tags

Cloud Sync assigns the relationship tags to all of the files transferred to the S3 bucket.

This enables you to search for the transferred files by using the tag values.

● Save on Object's Tags ○ Save On Object's Metadata

Tag Key: Up to 128 characters

Tag Value: Up to 256 characters

Optional Field | [Up to 5]

Support for Box

Cloud Sync now supports Box as the source in a sync relationship to Amazon S3, StorageGRID, and IBM Cloud Object Storage when using the Cloud Sync API.

Learn how to set up a sync relationship using the API.

Public IP for data broker in Google Cloud

When you deploy a data broker in Google Cloud, you can now choose whether to enable or disable a public IP address for the virtual machine instance.

Learn how to deploy a data broker in Google Cloud.

Dual-protocol volume for Azure NetApp Files

When you choose the source or target volume for Azure NetApp Files, Cloud Sync now displays a dual-protocol volume no matter which protocol you chose for the sync relationship.

07 July 2021

ONTAP S3 Storage and Google Cloud Storage

Cloud Sync now supports sync relationships between ONTAP S3 Storage and a Google Cloud Storage bucket from the user interface.

View the list of supported sync relationships.

Object metadata tags

Cloud Sync can now copy object metadata and tags between object-based storage when you create a sync relationship and enable a setting.

Learn more about the Copy for Objects setting.

Support for HashiCorp Vaults

You can now set up the data broker to access credentials from an external HashiCorp Vault by authenticating with a Google Cloud service account.

Learn more about using a HashiCorp Vault with a data broker.

Define tags or metadata for S3 bucket

When setting up a sync relationship to an Amazon S3 bucket, the Sync Relationship wizard now enables you to define the tags or metadata that you want to save on the objects in the target S3 bucket.

The tagging option was previously part of the sync relationship’s settings.

07 June 2021

Storage classes in Google Cloud

When a Google Cloud Storage bucket is the target in a sync relationship, you can now choose the storage class that you want to use. Cloud Sync supports the following storage classes:

- Standard
- Nearline
- Coldline
- Archive

02 May 2021

Errors in reports

You can now view the errors found in reports and you can delete the last report or all reports.

Learn more about creating and viewing reports to tune your configuration.

Compare attributes

A new **Compare by** setting is now available for each sync relationship.

This advanced setting enables you to choose whether Cloud Sync should compare certain attributes when determining whether a file or directory has changed and should be synced again.

Learn more about changing the settings for a sync relationship.

11 Apr 2021

Standalone Cloud Sync service is retired

The standalone Cloud Sync service has been retired. You should now access Cloud Sync directly from the NetApp Console where all of the same features and functionality are available.

After logging in to the NetApp Console, you can switch to the Sync tab at the top and view your relationships, just like before.

Google Cloud buckets in different projects

When setting up a sync relationship, you can choose from Google Cloud buckets in different projects, if you provide the required permissions to the data broker’s service account.

Learn how to set up the service account.

Metadata between Google Cloud Storage and S3

Cloud Sync now copies metadata between Google Cloud Storage and S3 providers (AWS S3, StorageGRID, and IBM Cloud Object Storage).

Restart data brokers

You can now restart a data broker from Cloud Sync.

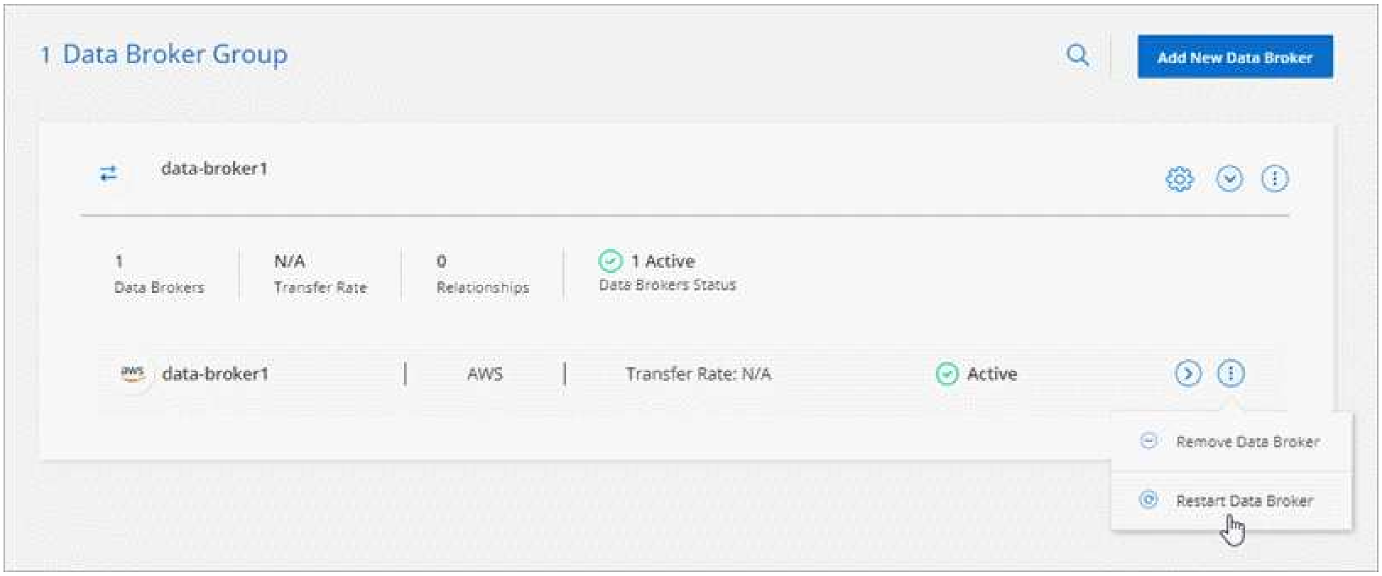

Message when not running the latest release

Cloud Sync now identifies when a data broker isn’t running the latest software release. This message can help to ensure that you’re getting the latest features and functionalities.

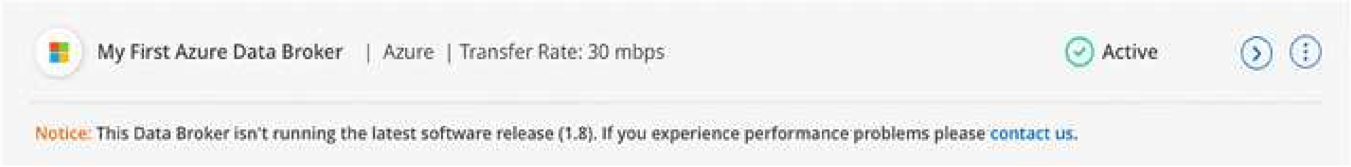

Notice: This Data Broker isn't running the latest software release (1.8). If you experience performance problems please contact us.

Limitations in NetApp Copy and Sync

Known limitations identify platforms, devices, or functions that are not supported by this release of the product, or that do not interoperate correctly with it. Review these limitations carefully.

NetApp Copy and Sync is not supported in the following regions:

- AWS Government regions
- Azure Government regions
- China

Get started

Learn about NetApp Copy and Sync

NetApp Copy and Sync 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, Copy and Sync can convert and move it for you.

NetApp Console

NetApp Copy and Sync is accessible through the NetApp Console.

The NetApp Console provides centralized management of NetApp storage and data services across on-premises and cloud environments at enterprise grade. The Console is required to access and use NetApp data services. As a management interface, it enables you to manage many storage resources from one interface. Console administrators can control access to storage and services for all systems within the enterprise.

You don’t need a license or subscription to start using NetApp Console and you only incur charges when you need to deploy Console agents in your cloud to ensure connectivity to your storage systems or NetApp data services. However, some NetApp data services accessible from the Console are licensed or subscription-based.

Learn more about the NetApp Console.

How NetApp Copy and Sync works

NetApp Copy and Sync is a software-as-a-service (SaaS) platform that consists of a data broker group, a cloud-based interface available through the NetApp Console, and a source and target.

The following image shows the relationship between Copy and Sync components:

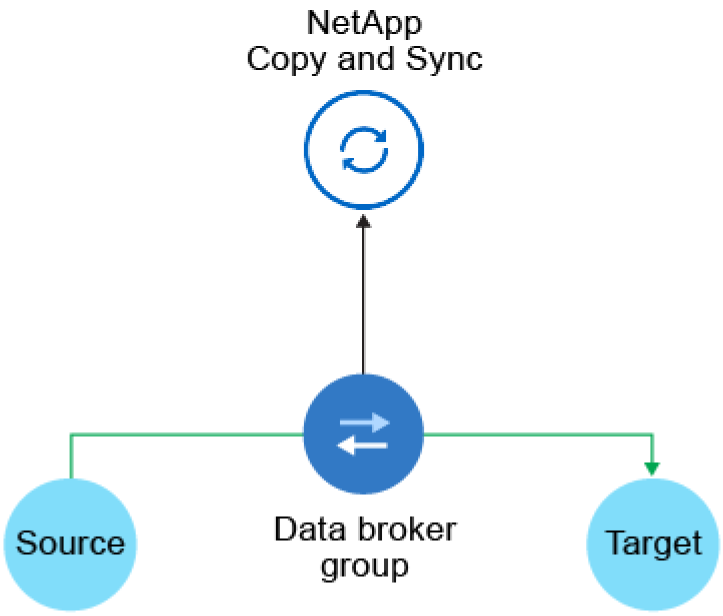

The NetApp data broker software syncs data from a source to a target (this is called a *sync relationship*). You can run the data broker in AWS, Azure, Google Cloud Platform, or on your premises. A data broker group, which consists of one or more data brokers, needs an outbound internet connection over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

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

Supported storage types

Copy and Sync supports the following storage types:

- Any NFS server
- Any SMB server
- Amazon EFS
- Amazon FSx for ONTAP
- Amazon S3
- Azure Blob
- Azure Data Lake Storage Gen2
- Azure NetApp Files
- Box (available as a preview)
- Cloud Volumes ONTAP
- Google Cloud Storage
- Google Drive
- IBM Cloud Object Storage
- On-premises ONTAP cluster
- ONTAP S3 Storage
- SFTP (using API only)
- StorageGRID

View the supported sync relationships.

Costs

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

Resource charges

Resource charges are related to the compute and storage costs for running one or more data brokers in the cloud.

Service charges

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 hourly or annually. The second option is to purchase licenses directly from NetApp.

Learn how licensing works.

Quick start for NetApp Copy and Sync

Getting started with NetApp Copy and Sync includes a few steps.

1 **Log in and set up the NetApp Console**

You should have gotten started with the NetApp Console, which includes logging in, setting up an account, and possibly deploying a Console agent and creating systems.

If you want to create sync relationships for any of the following, then you first need to create or discover a system:

- Amazon FSx for ONTAP
- Azure NetApp Files
- Cloud Volumes ONTAP
- On-prem ONTAP clusters

A Console agent is required for Cloud Volumes ONTAP, on-prem ONTAP clusters, and Amazon FSx for ONTAP.

- Learn how to get started with the NetApp Console
- Learn more about Console agents

2 **Prepare your source and target**

Verify that your source and target are supported and set up. The most important requirement is to verify connectivity between the data broker group and the source and target locations.

- View supported relationships
- Prepare the source and target

3 **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. A data broker group, which consists of one or more data brokers, needs an outbound internet connection over port 443 so it can communicate with NetApp Copy and Sync and contact a few other services and repositories. View the list of endpoints.

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

- Review AWS installation
- Review Azure installation
- Review Google Cloud installation
- Review Linux host installation

4

Create your first sync relationship

Log in to the NetApp Console, select **Sync**, and then drag and drop your selections for the source and target. Follow the prompts to complete the setup. Learn more.

5

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 NetApp Copy and Sync to set it up. Learn more.

Supported sync relationships in NetApp Copy and Sync

NetApp Copy and 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
Amazon EFS	• Amazon EFS • Amazon FSx for ONTAP • Amazon S3 • Azure Blob • Azure NetApp Files • Cloud Volumes ONTAP • Google Cloud Storage • IBM Cloud Object Storage • NFS server • On-premises ONTAP cluster (NFS or SMB) • SMB server • StorageGRID

Source location	Supported target locations
Amazon FSx for ONTAP	• Amazon EFS • Amazon FSx for ONTAP • Amazon S3 • Azure Blob • Azure Data Lake Storage Gen2 • Azure NetApp Files • Cloud Volumes ONTAP • Google Cloud Storage • IBM Cloud Object Storage • NFS server • On-premises ONTAP cluster (NFS or SMB) • SMB server • StorageGRID
Amazon S3	• Amazon EFS • Amazon FSx for ONTAP • Amazon S3 • Azure Blob • Azure Data Lake Storage Gen2 • Azure NetApp Files • Box [1] • Cloud Volumes ONTAP • Google Cloud Storage • IBM Cloud Object Storage • NFS server • On-premises ONTAP cluster (NFS or SMB) • ONTAP S3 Storage • SMB server • StorageGRID

Source location	Supported target locations
Azure Blob	• Amazon EFS • Amazon FSx for ONTAP • Amazon S3 • Azure Blob • Azure NetApp Files • Cloud Volumes ONTAP • Google Cloud Storage • IBM Cloud Object Storage • NFS server • On-premises ONTAP cluster (NFS or SMB) • SMB server • StorageGRID
Azure Data Lake Storage Gen2	• Azure NetApp Files • Cloud Volumes ONTAP • FSx for ONTAP • IBM Cloud Object Storage • NFS server • On-Prem ONTAP • ONTAP S3 Storage • SMB server • StorageGRID
Azure NetApp Files	• Amazon EFS • Amazon FSx for ONTAP • Amazon S3 • Azure Blob • Azure Data Lake Storage Gen2 • Azure NetApp Files • Cloud Volumes ONTAP • Google Cloud Storage • IBM Cloud Object Storage • NFS server • On-premises ONTAP cluster (NFS or SMB) • SMB server • StorageGRID

Source location	Supported target locations
Box [1]	• Amazon FSx for ONTAP • Amazon S3 • Azure NetApp Files • Cloud Volumes ONTAP • IBM Cloud Object Storage • NFS server • SMB server • StorageGRID
Cloud Volumes ONTAP	• Amazon EFS • Amazon FSx for ONTAP • Amazon S3 • Azure Blob • Azure Data Lake Storage Gen2 • Azure NetApp Files • Cloud Volumes ONTAP • Google Cloud Storage • IBM Cloud Object Storage • NFS server • On-premises ONTAP cluster (NFS or SMB) • SMB server • StorageGRID
Google Cloud Storage	• Amazon EFS • Amazon FSx for ONTAP • Amazon S3 • Azure Blob • Azure NetApp Files • Cloud Volumes ONTAP • Google Cloud Storage • IBM Cloud Object Storage • NFS server • On-premises ONTAP cluster (NFS or SMB) • ONTAP S3 Storage • SMB server • StorageGRID

Source location	Supported target locations
Google Drive	• NFS server • SMB server
IBM Cloud Object Storage	• Amazon EFS • Amazon FSx for ONTAP • Amazon S3 • Azure Blob • Azure Data Lake Storage Gen2 • Azure NetApp Files • Box [1] • Cloud Volumes ONTAP • Google Cloud Storage • IBM Cloud Object Storage • NFS server • On-premises ONTAP cluster (NFS or SMB) • SMB server • StorageGRID
NFS server	• Amazon EFS • Amazon FSx for ONTAP • Amazon S3 • Azure Blob • Azure Data Lake Storage Gen2 • Azure NetApp Files • Cloud Volumes ONTAP • Google Cloud Storage • Google Drive • IBM Cloud Object Storage • NFS server • On-premises ONTAP cluster (NFS or SMB) • ONTAP S3 Storage • SMB server • StorageGRID

Source location	Supported target locations
On-premises ONTAP cluster (NFS or SMB)	• Amazon EFS • Amazon FSx for ONTAP • Amazon S3 • Azure Blob • Azure Data Lake Storage Gen2 • Azure NetApp Files • Cloud Volumes ONTAP • Google Cloud Storage • IBM Cloud Object Storage • NFS server • On-premises ONTAP cluster (NFS or SMB) • SMB server • StorageGRID
ONTAP S3 Storage	• Amazon S3 • Azure Data Lake Storage Gen2 • Google Cloud Storage • NFS server • SMB server • StorageGRID • ONTAP S3 Storage
SFTP [2]	S3

Source location	Supported target locations
SMB server	• Amazon EFS • Amazon FSx for ONTAP • Amazon S3 • Azure Blob • Azure Data Lake Storage Gen2 • Azure NetApp Files • Cloud Volumes ONTAP • Google Cloud Storage • Google Drive • IBM Cloud Object Storage • NFS server • On-premises ONTAP cluster (NFS or SMB) • ONTAP S3 Storage • SMB server • StorageGRID
StorageGRID	• Amazon EFS • Amazon FSx for ONTAP • Amazon S3 • Azure Blob • Azure Data Lake Storage Gen2 • Azure NetApp Files • Box [1] • Cloud Volumes ONTAP • Google Cloud Storage • IBM Cloud Object Storage • NFS server • On-premises ONTAP cluster (NFS or SMB) • ONTAP S3 Storage • SMB server • StorageGRID

Notes:

1. Box support is available as a preview.
2. Sync relationships with this source/target are supported by using the Copy and Sync API only.
3. You can choose a specific Azure Blob storage tier when a Blob container is the target:

    - Hot storage
    - Cool storage

4. You can choose a specific S3 storage class when Amazon S3 is the target:
    - Standard (this is the default class)
    - Intelligent-Tiering
    - Standard-Infrequent Access
    - One Zone-Infrequent Access
    - Glacier Deep Archive
    - Glacier Flexible Retrieval
    - Glacier Instant Retrieval
5. You can choose a specific storage class when a Google Cloud Storage bucket is the target:
    - Standard
    - Nearline
    - Coldline
    - Archive

Prepare the source and target in NetApp Copy and Sync

Verify that your source and targets meet the following requirements in NetApp Copy and Sync.

Networking

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

  For example, if an NFS server is in your data center and a 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, the target, and data brokers to use a Network Time Protocol (NTP) service. The time difference between the three components should not exceed 5 minutes.

Target directory

When you create a sync relationship, Copy and Sync enables you to select an existing target directory and then optionally create a new folder inside that directory. So be sure that your preferred target directory already exists.

Permissions to read directories

In order to show every directory or folder in a source or target, Copy and Sync needs read permissions on the directory or folder.

NFS

Permissions must be defined on the source/target with uid/gid on files and directories.

Object storage

- For AWS and Google Cloud, a data broker must have list object permissions (these permissions are provided by default if you follow the data broker installation steps).
- For Azure, StorageGRID, and IBM, the credentials that you enter when setting up a sync relationship must have list object permissions.

SMB

The SMB credentials that you enter when setting up a sync relationship must have list folder permissions.

> The data broker ignores the following directories by default: .snapshot, ~snapshot, .copy-offload

> When copying SMB data into Cloud Volumes ONTAP using Copy and Sync, file and folder ownership from the source system is not preserved. This behavior occurs because Copy and Sync uses a Linux SMB client, which assigns ownership to the user or service account used to authenticate the transfer. While access control lists may be retained, ownership and audit information can differ from the source system. This is expected behavior.

Amazon S3 bucket requirements

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

Supported data broker locations for Amazon S3

Sync relationships that include S3 storage require a data broker deployed in AWS or on your premises. In either case, Copy and 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 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 a data broker.

The permissions included in this JSON file must be applied to that S3 bucket so a 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 a data broker.
2. *<RoleARN>* should be replaced with one of the following:
    - If a 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 a data broker.
    - If a 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 then selecting the IAM role from the Description tab. You should then see the Summary page in the IAM console that contains the Role ARN.

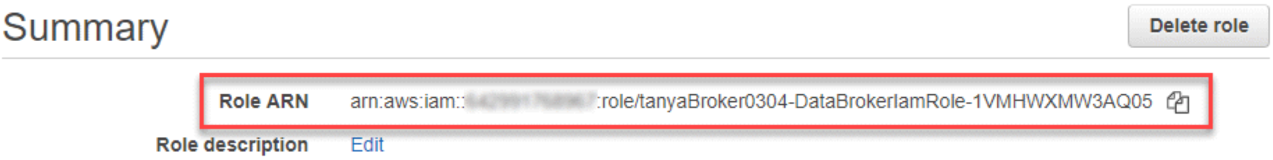

Summary

Role ARN: arn:aws:iam::[redacted]:role/tanyaBroker0304-DataBrokerIamRole-1VMHWXMW3AQ05

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

A 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 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 Copy and Sync with the storage account connection string:

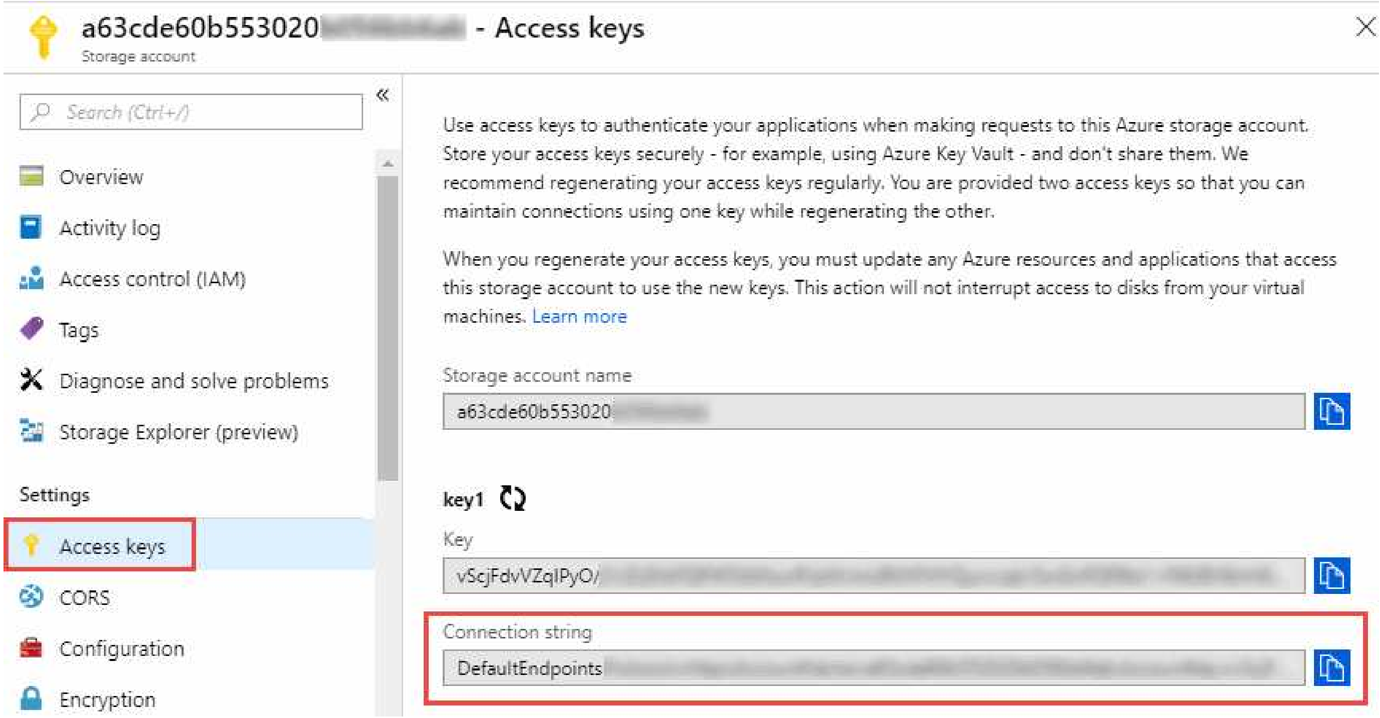

a63cde60b553020[redacted] - Access keys
Storage account

Use access keys to authenticate your applications when making requests to this Azure storage account. Store your access keys securely - for example, using Azure Key Vault - and don't share them. We recommend regenerating your access keys regularly. You are provided two access keys so that you can maintain connections using one key while regenerating the other.

When you regenerate your access keys, you must update any Azure resources and applications that access this storage account to use the new keys. This action will not interrupt access to disks from your virtual machines. Learn more

Storage account name: a63cde60b553020[redacted]

key1

Key: vScjFdvVZqlPyO/[redacted]

Connection string: DefaultEndpoints[redacted]

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

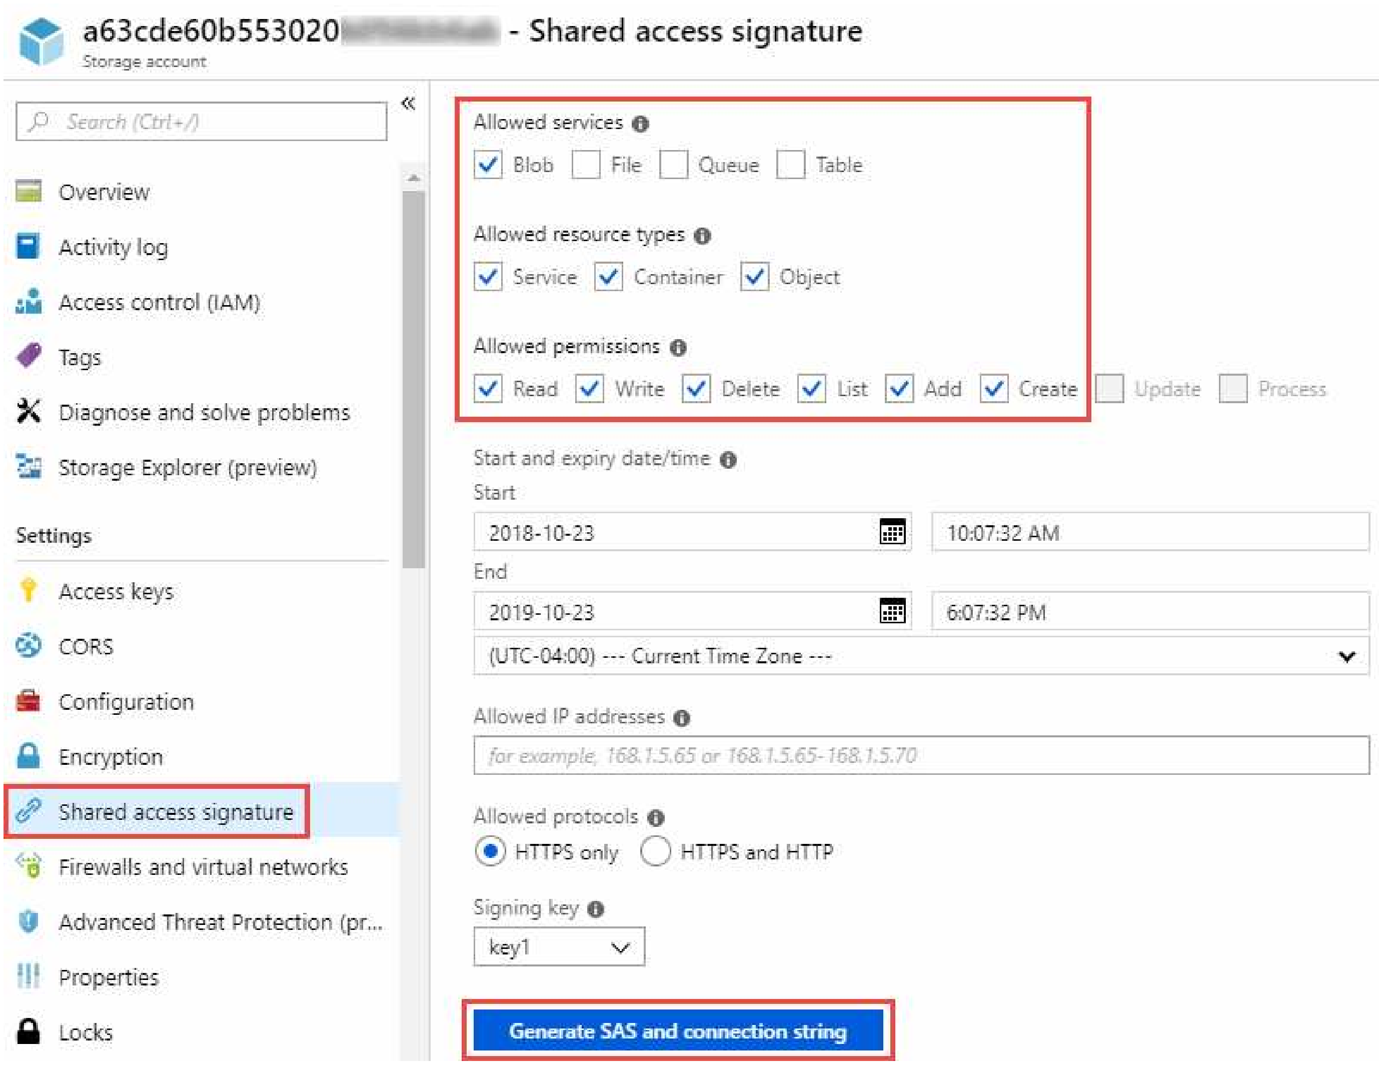

a63cde60b553020[redacted] - Shared access signature

Storage account

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

If you choose to implement a Continuous Sync relationship that includes an Azure Blob container, you can use a regular connection string or SAS connection string. If using a SAS connection string, it must not be set to expire in the near future.

Azure Data Lake Storage Gen2

When creating a sync relationship that includes Azure Data Lake, you need to provide Copy and Sync with the storage account connection string. It must be a regular connection string, not a shared access signature (SAS).

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.

Box requirements

- To create a sync relationship that includes Box, you'll need to provide the following credentials:
  - Client ID
  - Client secret
  - Private key
  - Public key ID
  - Passphrase
  - Enterprise ID
- If you create a sync relationship from Amazon S3 to Box, you must use a data broker group that has a unified configuration where the following settings are set to 1:
  - Scanner Concurrency
  - Scanner Processes Limit
  - Transferrer Concurrency
  - Transferrer Processes Limit

  Learn how to define a unified configuration for a data broker group.

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 Google Cloud or on your premises. Copy and Sync guides you through the data broker installation process when you create a sync relationship.

- Learn how to deploy the Google Cloud data broker
- Learn how to install the data broker on a Linux host

Supported Google Cloud regions

All regions are supported.

Permissions for buckets in other Google Cloud projects

When setting up a sync relationship, you can choose from Google Cloud buckets in different projects, if you provide the required permissions to the data broker's service account. Learn how to set up the service account.

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.

Encrypting a Google Cloud bucket

You can encrypt a target Google Cloud bucket with a customer-managed KMS key or the default, Google-managed key. If the bucket already has a KMS encryption added to it, it will override the default Google-managed encryption.

To add a customer-managed KMS key, you will need to use a data broker with the correct permissions, and the key must be in the same region as the bucket.

Google Drive

When you set up a sync relationship that includes Google Drive, you'll need to provide the following:

- The email address for a user who has access to the Google Drive location where you want to sync data
- The email address for a Google Cloud service account that has permissions to access Google Drive
- A private key for the service account

To set up the service account, follow the instructions in Google documentation:

- Create the service account and credentials
- Delegate domain-wide authority to your service account

When you edit the OAuth Scopes field, enter the following scopes:

- https://www.googleapis.com/auth/drive
- https://www.googleapis.com/auth/drive.file

NFS server requirements

- The NFS server can be a NetApp system or a non-NetApp system.
- The file server must allow a data broker host to access the exports over the required ports.
    - 111 TCP/UDP
    - 2049 TCP/UDP
    - 5555 TCP/UDP
- 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.

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.

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

SMB server requirements

- The SMB server can be a NetApp system or a non-NetApp system.
- You need to provide Copy and Sync with credentials that have permissions on the SMB server.
  - For a source SMB server, the following permissions are required: list and read.

    Members of the Backup Operators group are supported with a source SMB server.
  - For a target SMB server, the following permissions are required: list, read, and write.
- The file server must allow a data broker host to access the exports over the required ports.
  - 139 TCP
  - 445 TCP
  - 137-138 UDP
- 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"

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.

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 Copy and 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 Copy and Sync tries to copy the directory named "b" to the target, Copy and Sync receives an error that says the directory already exists. As a result, Copy and 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.

Networking overview for NetApp Copy and Sync

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

Data broker location

A data broker group consists of one or more data brokers installed in the cloud or on your premises.

Data broker in the cloud

The following image shows a data broker running in the cloud, in either AWS, Google Cloud, 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 Copy and Sync deploys the data broker in AWS, Azure, or Google Cloud, it creates a security group that enables the required outbound communication.

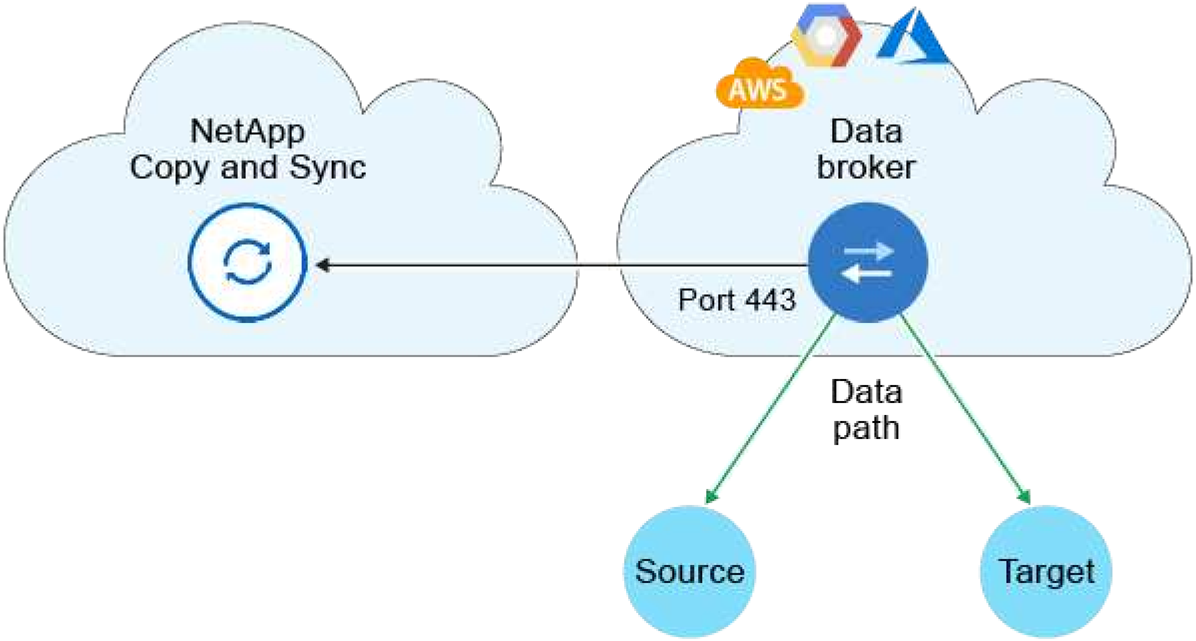

Data broker on your premises

The following image shows the data broker running on-premises 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.

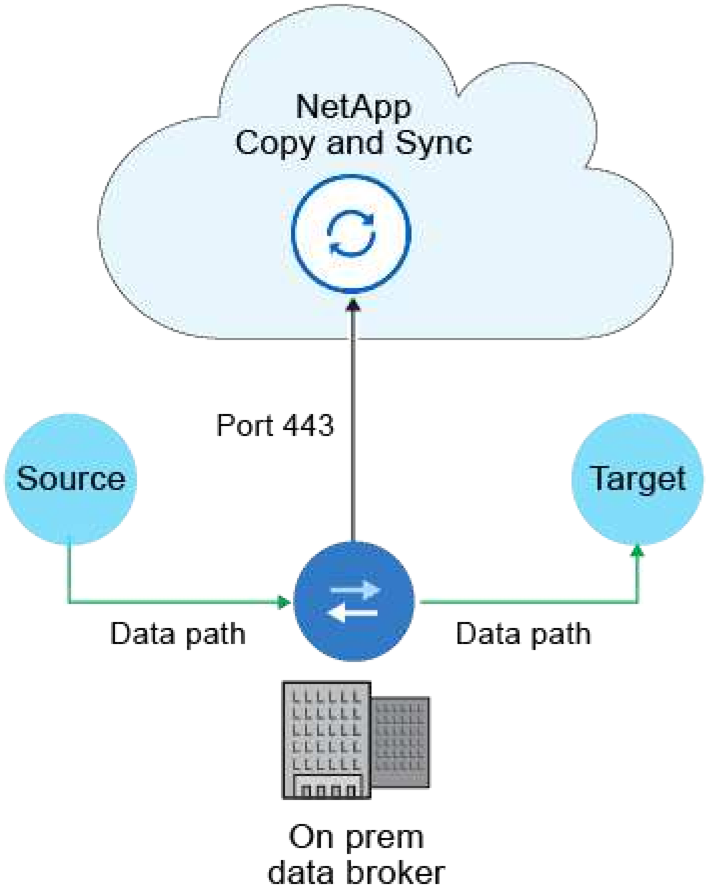

Networking requirements

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

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

- A data broker needs an outbound internet connection so it can poll Copy and Sync for tasks over port 443.
- NetApp recommends configuring the source, target, and data brokers to use a Network Time Protocol (NTP) service. The time difference between the three components should not exceed 5 minutes.

Networking endpoints

The NetApp data broker requires outbound internet access over port 443 to communicate with Copy and Sync 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

A data broker contacts the following endpoints:

Endpoints	Purpose
https://olcentgbl.trafficmanager.net	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.
https://rpm.nodesource.com https://registry.npmjs.org https://nodejs.org:	To contact repositories for updating Node.js, npm, and other 3rd party packages used in development.
https://tgz.pm2.io	To access a repository for updating PM2, which is a 3rd party package used to monitor Copy and Sync.
https://sqs.us-east-1.amazonaws.com https://kinesis.us-east-1.amazonaws.com	To contact the AWS services that Copy and Sync uses for operations (queuing files, registering actions, and delivering updates to the data broker).
https://s3.*region*.amazonaws.com For example: s3.us-east-2.amazonaws.com:443 See AWS documentation for a list of S3 endpoints	To contact Amazon S3 when a sync relationship includes an S3 bucket.
https://s3.amazonaws.com/	When you download data broker logs from Copy and Sync, the data broker zips its logs directory and uploads the logs to a predefined S3 bucket in the us-east-1 region.
https://storage.googleapis.com/	To contact Google Cloud when a sync relationship uses a GCP bucket.
https://*storage-account*.blob.core.windows.net If using Azure Data Lake Gen2: https://*storage-account*.dfs.core.windows.net Where *storage-account* is the user's source storage account.	To open the proxy to a user's Azure storage account address.
https://cf.cloudsync.netapp.com https://repo.cloudsync.netapp.com	To contact Copy and Sync.
https://support.netapp.com	To contact NetApp support when using a BYOL license for sync relationships.
https://fedoraproject.org	To install 7z on the data broker virtual machine during installation and updates. 7z is needed to send AutoSupport messages to NetApp technical support.
https://sts.amazonaws.com https://sts.us-east-1.amazonaws.com	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.
https://api.bluexp.netapp.com https://netapp-cloud-account.auth0.com	To contact NetApp Data Classification when you use classification to select the source files for a new sync relationship.

Endpoints	Purpose
https://pubsub.googleapis.com	If creating a continuous sync relationship from a Google storage account.
https://*storage-account*.queue.core.windows.net https://management.azure.com/subscriptions/${*subscriptionId*}/resourceGroups/${*resourceGroup*}/providers/Microsoft.EventGrid/* Where *storage-account* is the user's source storage account, *subscriptionid* is the is the source subscription ID, and *resourceGroup* is the source resource group.	If creating a continuous sync relationship from an Azure storage account.

Web browser endpoints

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

logs.cloudsync.netapp.com:443

Log in to NetApp Copy and Sync

Use the NetApp Console to log in to NetApp Copy and Sync.

To log in to the Console, you can use your NetApp Support Site credentials or you can sign up for a NetApp cloud login using your email and a password. Learn more about logging in.

NetApp Copy and Sync uses identity access management to govern the access that each user has to specific actions.

Required NetApp Console role
Organization admin role. Learn about NetApp Console access roles.

Steps

1. Open a web browser and go to the NetApp Console.

   The NetApp Console login page appears.

2. Log in to the Console.
3. From the Console left navigation, select **Mobility** > **Copy and Sync**.

Install a data broker

Create a new data broker in AWS for NetApp Copy and Sync

When you create a new data broker group for NetApp Copy and Sync, choose Amazon Web Services to deploy the data broker software on a new EC2 instance in a VPC. NetApp Copy and 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 regions.

Root privileges

The data broker software automatically runs as root on the Linux host. Running as root is a requirement for data broker operations. For example, to mount shares.

Networking requirements

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

  When Copy and 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 Copy and 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.

Create the data broker

There are a few ways to create a new data broker. These steps describe how to install a data broker in AWS when creating a sync relationship.

Steps

1. Log in to Copy and Sync.

2. Select **Create New Sync**.
3. On the **Define Sync Relationship** page, choose a source and target and select **Continue**.

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

4. On the **Data Broker Group** page, select **Create Data Broker** and then select **Amazon Web Services**.

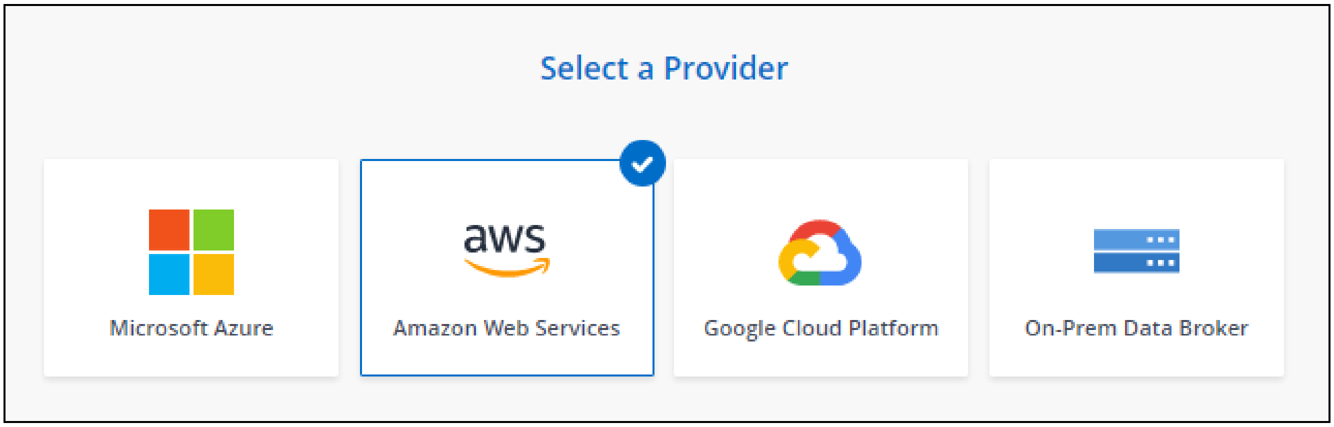

Select a Provider

5. Enter a name for the data broker and select **Continue**.
6. Enter an AWS access key so Copy and 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, select 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:

   Launch a data broker from an AWS CloudFormation template

7. 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 select an existing IAM role, or leave the field blank so Copy and Sync creates the role for you. You also have the option of encrypting your data broker using a KMS key.

   If you choose your own IAM role, you’ll need to provide the required permissions.

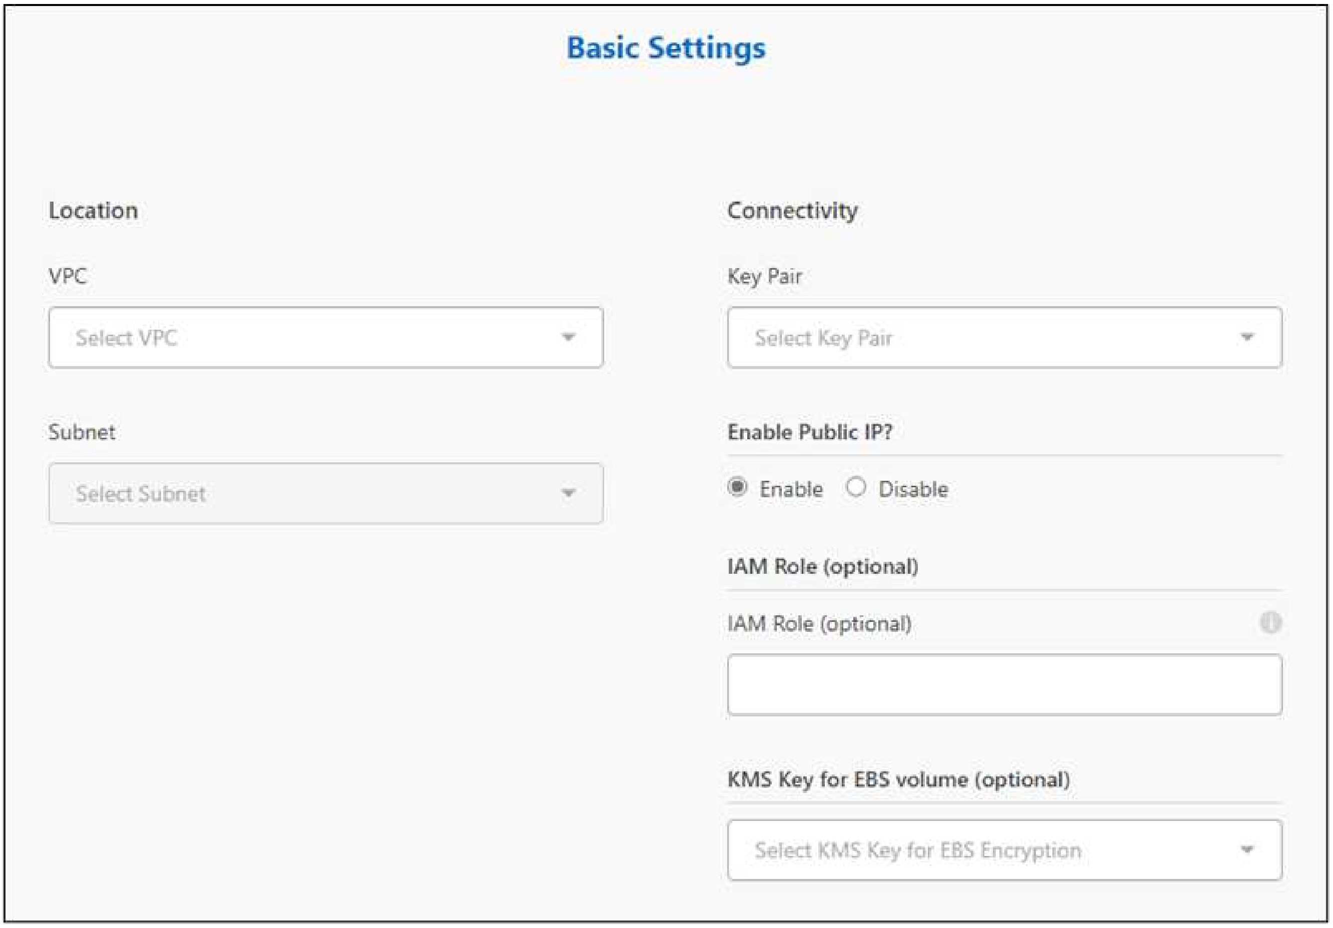

Basic Settings

Location

VPC

Select VPC

Subnet

Select Subnet

Connectivity

Key Pair

Select Key Pair

Enable Public IP?

● Enable ○ Disable

IAM Role (optional)

IAM Role (optional)

KMS Key for EBS volume (optional)

Select KMS Key for EBS Encryption

8. Specify a proxy configuration, if a proxy is required for internet access in the VPC.
9. After the data broker is available, select **Continue** in Copy and Sync.

   The following image shows a successfully deployed instance in AWS:

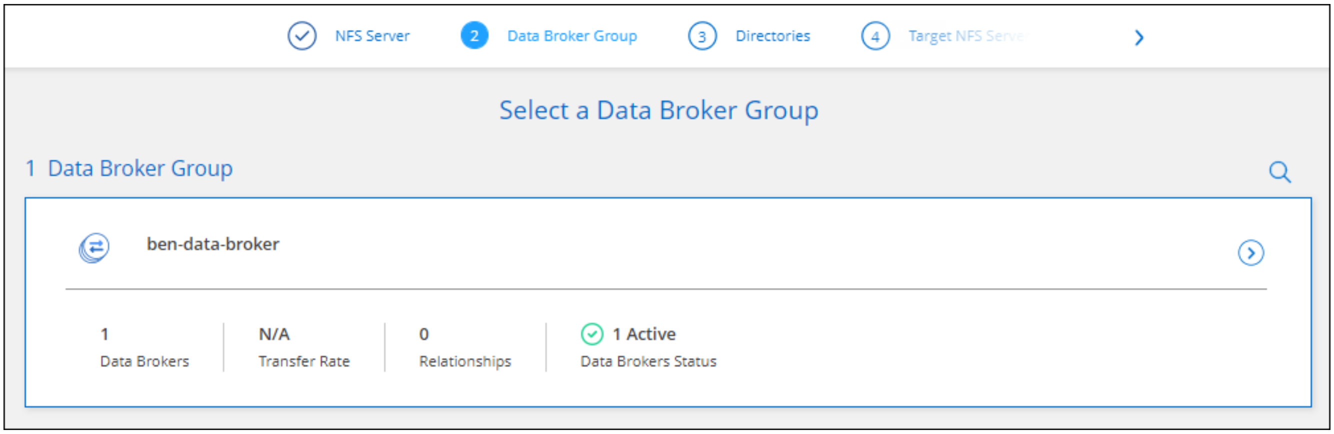

Select a Data Broker Group

1 Data Broker Group

ben-data-broker

1	N/A	0	✓ 1 Active
Data Brokers	Transfer Rate	Relationships	Data Brokers Status

10. 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 group with additional sync relationships.

Details about the data broker instance

Copy and Sync creates a data broker in AWS using the following configuration.

Node.js compatibility
v21.2.0

Instance type
m5n.xlarge when available in the region, otherwise m5.xlarge

vCPUs
4

RAM
16 GB

Operating system
Amazon Linux 2023

Disk size and type
10 GB GP2 SSD

Create a new data broker in Azure for NetApp Copy and Sync

When you create a new data broker group for NetApp Copy and Sync, choose the Microsoft Azure to deploy the data broker software on a new virtual machine in a VNet. NetApp Copy and 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.

Root privileges

The data broker software automatically runs as root on the Linux host. Running as root is a requirement for data broker operations. For example, to mount shares.

Networking requirements

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

  When Copy and 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.

Permissions required to deploy the data broker in Azure

Ensure that the Azure user account that you use to deploy the data broker has the following permissions:

```
{
    "Name": "Azure Data Broker",
    "Actions": [
                    "Microsoft.Resources/subscriptions/read",

"Microsoft.Resources/deployments/operationstatuses/read",
                    "Microsoft.Resources/subscriptions/locations/read",
                    "Microsoft.Network/networkInterfaces/read",
                    "Microsoft.Network/virtualNetworks/subnets/read",

"Microsoft.Resources/subscriptions/resourceGroups/write",

"Microsoft.Resources/subscriptions/resourceGroups/delete",
                    "Microsoft.Resources/deployments/write",
                    "Microsoft.Resources/deployments/validate/action",

"Microsoft.Resources/deployments/operationStatuses/read",
                    "Microsoft.Resources/deployments/cancel/action",
                    "Microsoft.Compute/virtualMachines/read",
                    "Microsoft.Compute/virtualMachines/delete",
                    "Microsoft.Compute/disks/delete",
                    "Microsoft.Network/networkInterfaces/delete",
                    "Microsoft.Network/publicIPAddresses/delete",

"Microsoft.Network/networkSecurityGroups/securityRules/delete",

"Microsoft.Resources/subscriptions/resourceGroups/write",
                    "Microsoft.Compute/virtualMachines/delete",
                    "Microsoft.Network/networkSecurityGroups/write",
                    "Microsoft.Network/networkSecurityGroups/join/action",
                    "Microsoft.Compute/disks/write",
                    "Microsoft.Network/networkInterfaces/write",
                    "Microsoft.Network/virtualNetworks/read",
                    "Microsoft.Network/publicIPAddresses/write",
                    "Microsoft.Compute/virtualMachines/write",
                    "Microsoft.Compute/virtualMachines/extensions/write",
                    "Microsoft.Resources/deployments/read",
                    "Microsoft.Network/networkSecurityGroups/read",
                    "Microsoft.Network/publicIPAddresses/read",

"Microsoft.Network/virtualNetworks/subnets/join/action",
                    "Microsoft.Network/publicIPAddresses/join/action",
                    "Microsoft.Network/networkInterfaces/join/action",
                    "Microsoft.Storage/storageAccounts/read",

"Microsoft.EventGrid/systemTopics/eventSubscriptions/write",

"Microsoft.EventGrid/systemTopics/eventSubscriptions/read",

"Microsoft.EventGrid/systemTopics/eventSubscriptions/delete",

"Microsoft.EventGrid/systemTopics/eventSubscriptions/getFullUrl/action",

"Microsoft.EventGrid/systemTopics/eventSubscriptions/getDeliveryAttributes
/action",
                    "Microsoft.EventGrid/systemTopics/read",
                    "Microsoft.EventGrid/systemTopics/write",
                    "Microsoft.EventGrid/systemTopics/delete",
                    "Microsoft.EventGrid/eventSubscriptions/write",
                    "Microsoft.Storage/storageAccounts/write"

"Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreemen
ts/read"

"Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreemen
ts/write"

"Microsoft.Network/networkSecurityGroups/securityRules/read",
                    "Microsoft.Network/networkSecurityGroups/read",
```

```
    ],
    "NotActions": [],
    "AssignableScopes": [],
    "Description": "Azure Data Broker",
    "IsCustom": "true"
}
```

Note:

1. The following permissions are only required if you plan to enable the Continuous Sync setting on a sync relationship from Azure to another cloud storage location:

   - 'Microsoft.Storage/storageAccounts/read',
   - 'Microsoft.EventGrid/systemTopics/eventSubscriptions/write',
   - 'Microsoft.EventGrid/systemTopics/eventSubscriptions/read',
   - 'Microsoft.EventGrid/systemTopics/eventSubscriptions/delete',

    - 'Microsoft.EventGrid/systemTopics/eventSubscriptions/getFullUrl/action',
    - 'Microsoft.EventGrid/systemTopics/eventSubscriptions/getDeliveryAttributes/action',
    - 'Microsoft.EventGrid/systemTopics/read',
    - 'Microsoft.EventGrid/systemTopics/write',
    - 'Microsoft.EventGrid/systemTopics/delete',
    - 'Microsoft.EventGrid/eventSubscriptions/write',
    - 'Microsoft.Storage/storageAccounts/write'

    Additionally, the assignable scope must be set to subscription scope and **not** resource group scope if you plan to implement Continuous Sync in Azure.

2. The following permissions are only required if you plan to choose your own security for data broker creation:
    - "Microsoft.Network/networkSecurityGroups/securityRules/read"
    - "Microsoft.Network/networkSecurityGroups/read"

Authentication method

When you deploy the data broker, you'll need to choose an authentication method for the virtual machine: 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.

Create the data broker

There are a few ways to create a new data broker. These steps describe how to install a data broker in Azure when you create a sync relationship.

Steps

1. Log in to Copy and Sync.
2. Select **Create New Sync**.
3. On the **Define Sync Relationship** page, choose a source and target and select **Continue**.

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

4. On the **Data Broker Group** page, select **Create Data Broker** and then select **Microsoft Azure**.

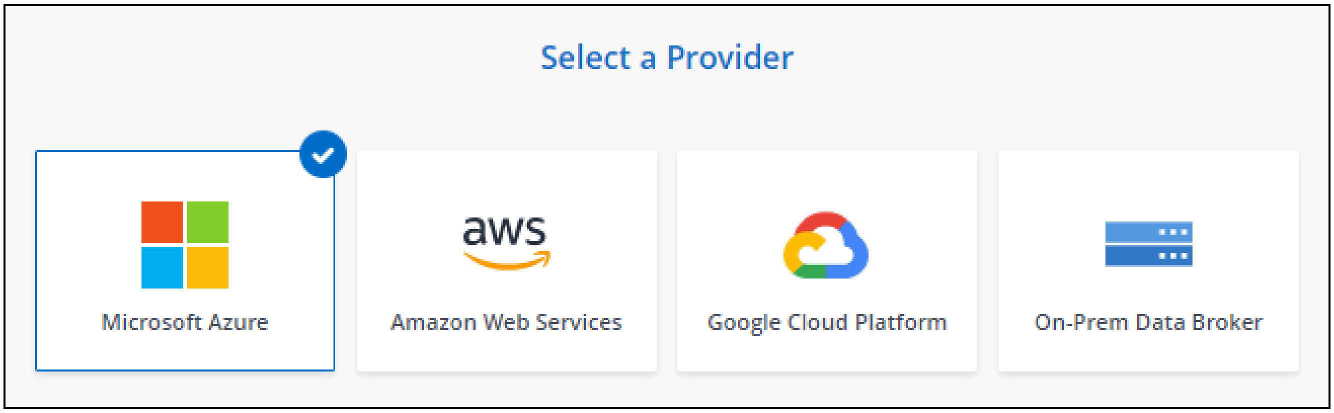

5. Enter a name for the data broker and select **Continue**.
6. If you're prompted, log in to your Microsoft account. If you're not prompted, select **Log in to Azure**.

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

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

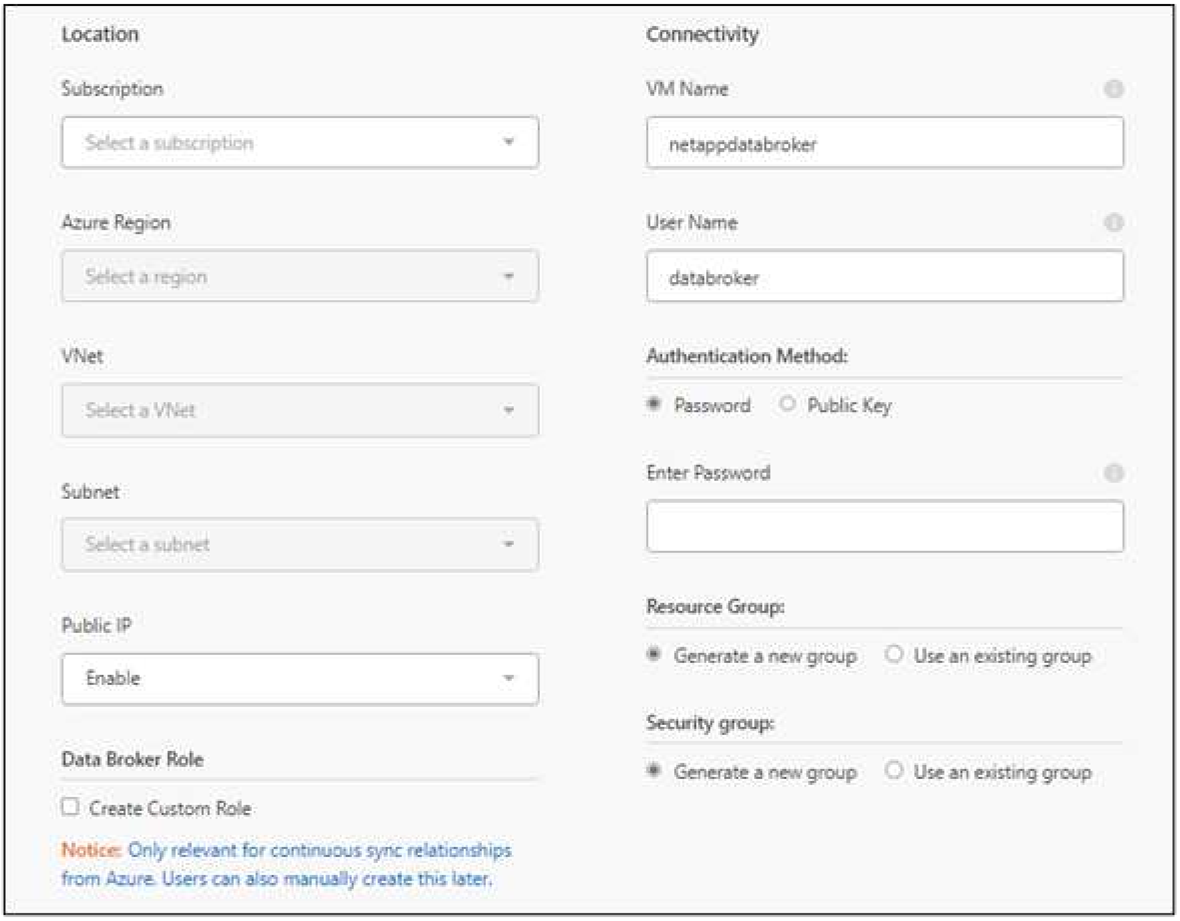

Location

Subscription

Select a subscription ▾

Azure Region

Select a region ▾

VNet

Select a VNet ▾

Subnet

Select a subnet ▾

Public IP

Enable ▾

Data Broker Role

□ Create Custom Role

Notice: Only relevant for continuous sync relationships from Azure. Users can also manually create this later.

Connectivity

VM Name

netappdatabroker

User Name

databroker

Authentication Method:

● Password ○ Public Key

Enter Password

Resource Group:

● Generate a new group ○ Use an existing group

Security group:

● Generate a new group ○ Use an existing group

> If you plan to implement a Continuous Sync relationship, you must assign a custom role to your data broker. This can also be done manually after the broker is created.

8. Specify a proxy configuration, if a proxy is required for internet access in the VNet.
9. Select **Continue**. If you would like to add S3 permissions to your data broker, enter your AWS access and secret keys.
10. Select **Continue** and keep the page open until the deployment is complete.

    The process can take up to 7 minutes.

11. In Copy and Sync, select **Continue** once the data broker is available.
12. 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 Copy and 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.

Details about the data broker VM

Copy and Sync creates a data broker in Azure using the following configuration.

Node.js compatibility
v21.2.0

VM type
Standard DS4 v2

vCPUs
8

RAM
28 GB

Operating system
Rocky Linux 9.0

Disk size and type
64 GB Premium SSD

Create a new data broker in Google Cloud for NetApp Copy and Sync

When you create a new data broker group for NetApp Copy and Sync, choose Google Cloud Platform to deploy the data broker software on a new virtual machine instance in a Google Cloud VPC. NetApp Copy and 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 Google Cloud regions

All regions are supported.

Root privileges

The data broker software automatically runs as root on the Linux host. Running as root is a requirement for data broker operations. For example, to mount shares.

Networking requirements

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

  When Copy and Sync deploys the data broker in Google Cloud, 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 Google Cloud

Ensure that the Google Cloud 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.create
- storage.objects.delete
- storage.objects.get
- storage.objects.getIamPolicy
- storage.objects.list
- storage.objects.setIamPolicy
- storage.objects.update
- iam.serviceAccounts.signJwt
- pubsub.subscriptions.consume
- pubsub.subscriptions.create
- pubsub.subscriptions.delete
- pubsub.subscriptions.list
- pubsub.topics.attachSubscription
- pubsub.topics.create
- pubsub.topics.delete
- pubsub.topics.list
- pubsub.topics.setIamPolicy
- storage.buckets.update
- cloudkms.cryptoKeys.list
- cloudkms.keyRings.list
```

Notes:

1. The "iam.serviceAccounts.signJwt" permission is required only if you’re planning to set up the data broker to use an external HashiCorp vault.
2. The "pubsub.*" and "storage.buckets.update" permissions are required only if you plan to enable the Continuous Sync setting on a sync relationship from Google Cloud Storage to another cloud storage location. Learn more about the Continuous Sync option.
3. The "cloudkms.cryptoKeys.list" and "cloudkms.keyRings.list" permissions are required only if you plan to use a customer-managed KMS key on a target Google Cloud Storage bucket.

Create the data broker

There are a few ways to create a new data broker. These steps describe how to install a data broker in Google Cloud when you create a sync relationship.

Steps

1. Log in to Copy and Sync.
2. Select **Create New Sync**.
3. On the **Define Sync Relationship** page, choose a source and target and select **Continue**.

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

4. On the **Data Broker Group** page, select **Create Data Broker** and then select **Google Cloud Platform**.

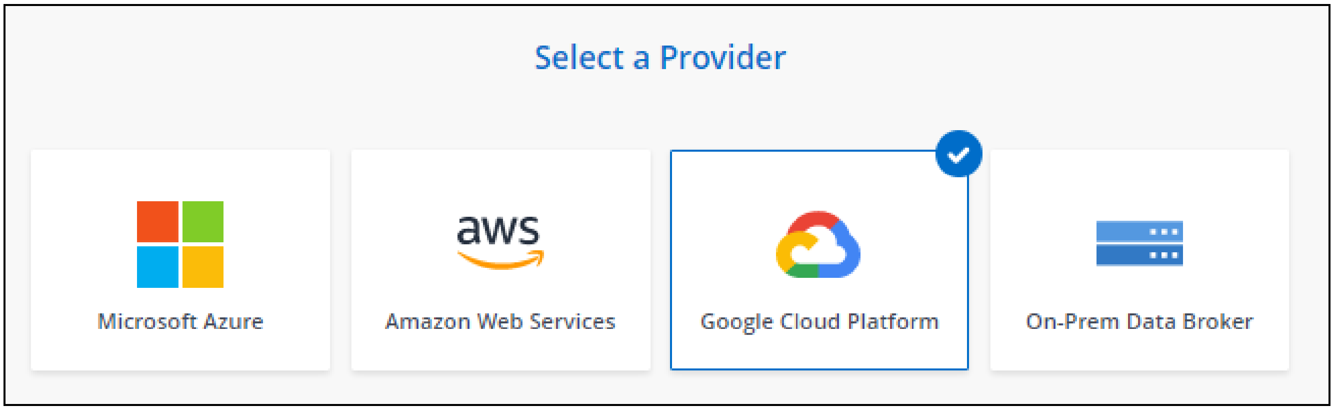

Select a Provider

5. Enter a name for the data broker and select **Continue**.
6. 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.

7. Select a project and service account and then choose a location for the data broker, including whether you want to enable or disable a public IP address.

   If you don't enable a public IP address, then you'll need to define a proxy server in the next step.

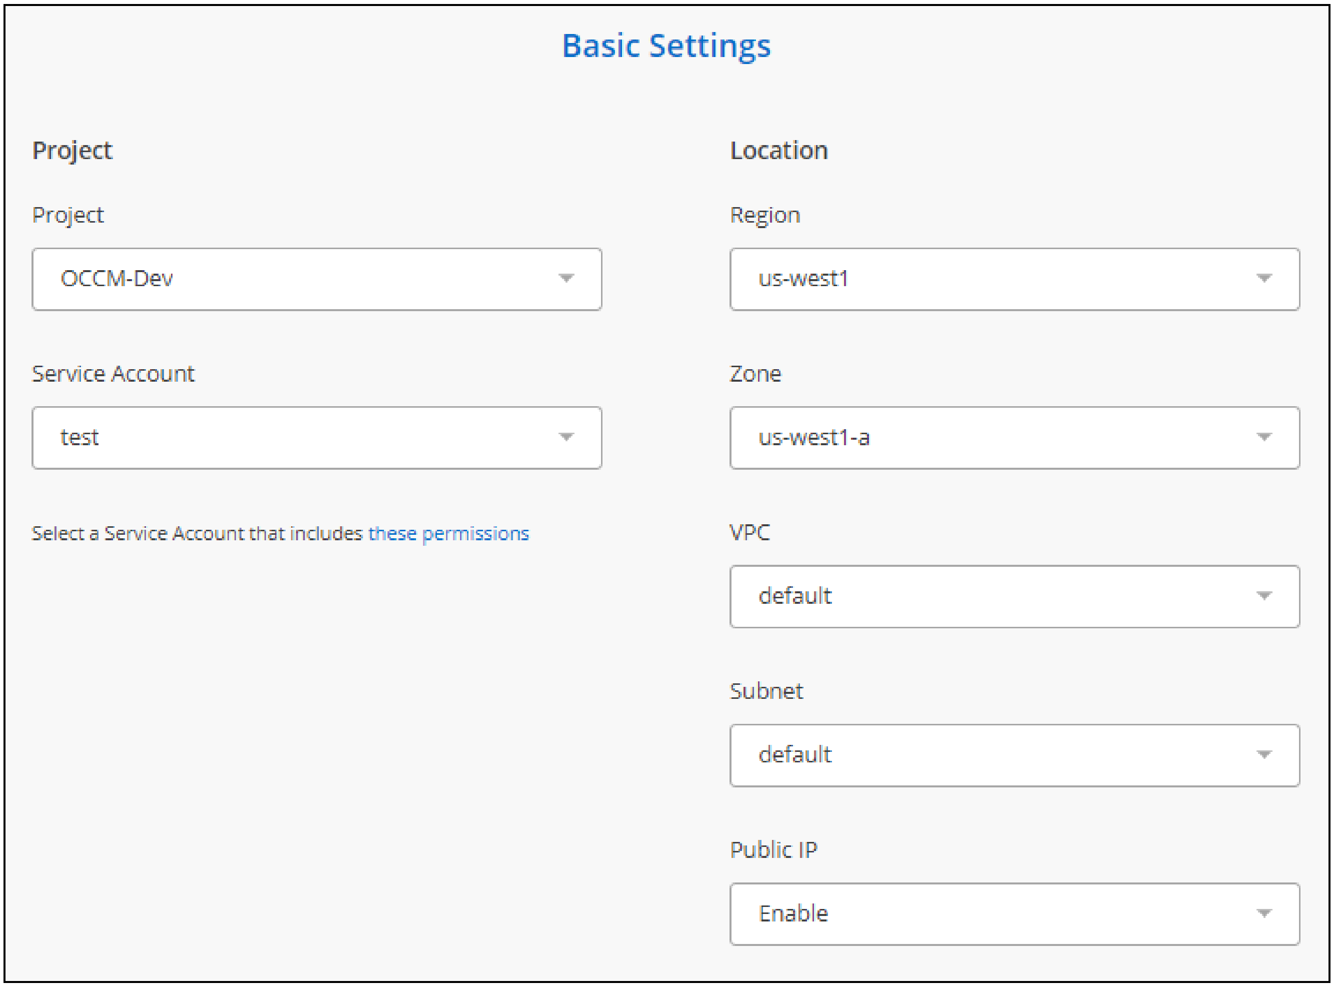

Basic Settings

Project

Project: OCCM-Dev

Service Account: test

Select a Service Account that includes these permissions

Location

Region: us-west1

Zone: us-west1-a

VPC: default

Subnet: default

Public IP: Enable

8. 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.

9. Once the data broker is available, select **Continue** in Copy and Sync.

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

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

Result

You’ve deployed a data broker in Google Cloud and created a new sync relationship. You can use this data broker with additional sync relationships.

Provide permissions to use buckets in other Google Cloud projects

When you create a sync relationship and choose Google Cloud Storage as the source or target, Copy and Sync enables you to choose from the buckets that the data broker’s service account has permissions to use. By default, this includes the buckets that are in the *same* project as the data broker service account. But you can choose buckets from *other* projects if you provide the required permissions.

Steps

1. Open the Google Cloud Platform console and load the Cloud Storage service.
2. Select the name of the bucket that you’d like to use as a source or target in a sync relationship.
3. Select **Permissions**.
4. Select **Add**.
5. Enter the name of the data broker’s service account.
6. Select a role that provides the same permissions as shown above.
7. Select **Save**.

Result

When you set up a sync relationship, you can now choose that bucket as the source or target in the sync relationship.

Details about the data broker VM instance

Copy and Sync creates a data broker in Google Cloud using the following configuration.

Node.js compatibility

v21.2.0

Machine type

n2-standard-4

vCPUs

4

RAM

15 GB

Operating system

Rocky Linux 9.0

Disk size and type

20 GB HDD pd-standard

Install the data broker on a Linux host for NetApp Copy and Sync

When you create a new data broker group for NetApp Copy and Sync, 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. NetApp Copy and 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

- **Node.js compatibility**: v21.2.0
- **Operating system**:
  - CentOS 8.0 and 8.5

    CentOS Stream is not supported.
  - Red Hat Enterprise Linux 8.5, 8.8, 8.9, and 9.4
  - Rocky Linux 9
  - Ubuntu Server 20.04 LTS, 23.04 LTS, and 24.04 LTS
  - SUSE Linux Enterprise Server 15 SP1

    The command `yum update` 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.

Root privileges

The data broker software automatically runs as root on the Linux host. Running as root is a requirement for data broker operations. For example, to mount shares.

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.

Enable 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.

Enable 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 Google Cloud 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 Google Cloud 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 Google Cloud 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.

Enable 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.

Install the data broker

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

Steps

1. Log in to Copy and Sync.
2. Select **Create New Sync**.
3. On the **Define Sync Relationship** page, choose a source and target and select **Continue**.

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

4. On the **Data Broker Group** page, select **Create Data Broker** and then select **On-Prem Data Broker**.

   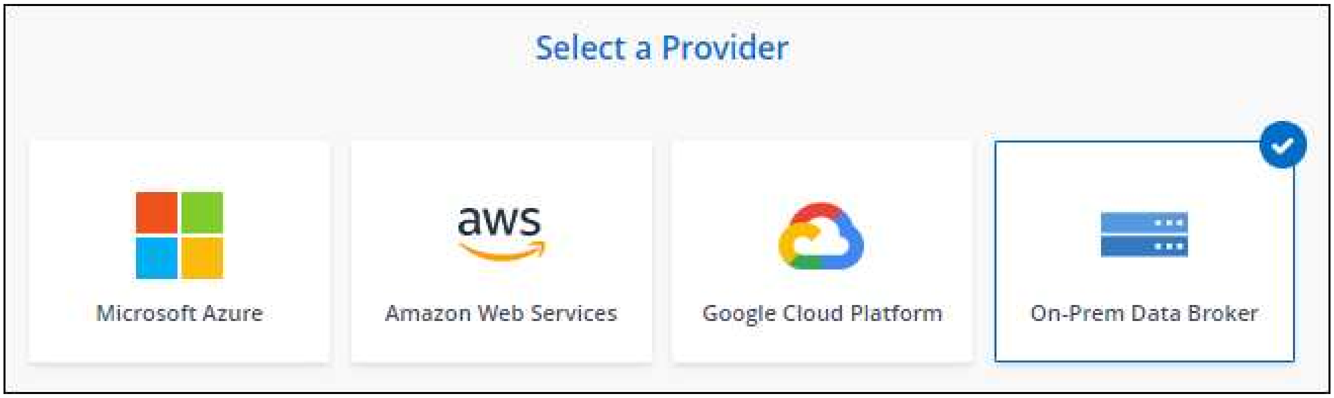

   ⓘ Even though the option is labeled ***On-Prem* Data Broker**, it applies to a Linux host on your premises or in the cloud.

5. Enter a name for the data broker and select **Continue**.

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

6. 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**.

      ⓘ The user must be a local user. Domain users are not supported.

   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

  Copy and 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 Copy and Sync.

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

> Each command listed below includes parameters for AWS access and Google Cloud 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.

7. Once the data broker is available, select **Continue** in Copy and Sync.
8. Complete the pages in the wizard to create the new sync relationship.

Use NetApp Copy and Sync

Sync data between a source and target

Prepare a data broker to sync data between object storage in NetApp Copy and Sync

If you're planning to sync data from object storage to object storage (for example, Amazon S3 to Azure Blob) in NetApp Copy and Sync, then you need to prepare the data broker group before you create the sync relationship.

About this task

To prepare the data broker group, you'll need to modify the configuration of the scanner. If you don't modify the configuration, you might notice performance issues for this sync relationship.

Before you begin

The data broker group that you use to sync data from object storage to object storage should only manage these types of sync relationships. If the data broker group manages a different type of sync relationship (for example, NFS to NFS or object storage to SMB), then the performance of those sync relationships might be negatively affected.

Steps

1. Log in to Copy and Sync.
2. From Copy and Sync, select **Manage Data Brokers**.
3. Select ⚙
4. Update the scanner configuration:
   a. Change **Scanner Concurrency** to **1**.
   b. Change **Scanner Processes Limit** to **1**.
5. Select **Unify Configuration**.

Result

Copy and Sync updates the configuration of the data broker group.

What's next?

You can now create the sync relationship between object storage using the data broker group that you just configured.

Create sync relationships in NetApp Copy and Sync

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

Before you can create some types of sync relationships, you'll first need to create a system in the NetApp Console.

Create sync relationships for specific types of systems

If you want to create sync relationships for any of the following, then you first need to create or discover the system:

- Amazon FSx for ONTAP
- Azure NetApp Files
- Cloud Volumes ONTAP
- On-prem ONTAP clusters

Steps

1. Log in to Copy and Sync.
2. Create or discover the system.
   - Create an Amazon FSx for ONTAP system
   - Setting up and discovering Azure NetApp Files
   - Launching Cloud Volumes ONTAP in AWS
   - Launching Cloud Volumes ONTAP in Azure
   - Launching Cloud Volumes ONTAP in Google Cloud
   - Adding existing Cloud Volumes ONTAP systems
   - Discovering ONTAP clusters
3. Select **Systems page**.
4. Select a system that matches any of the types listed above.
5. Select the action menu next to Sync.

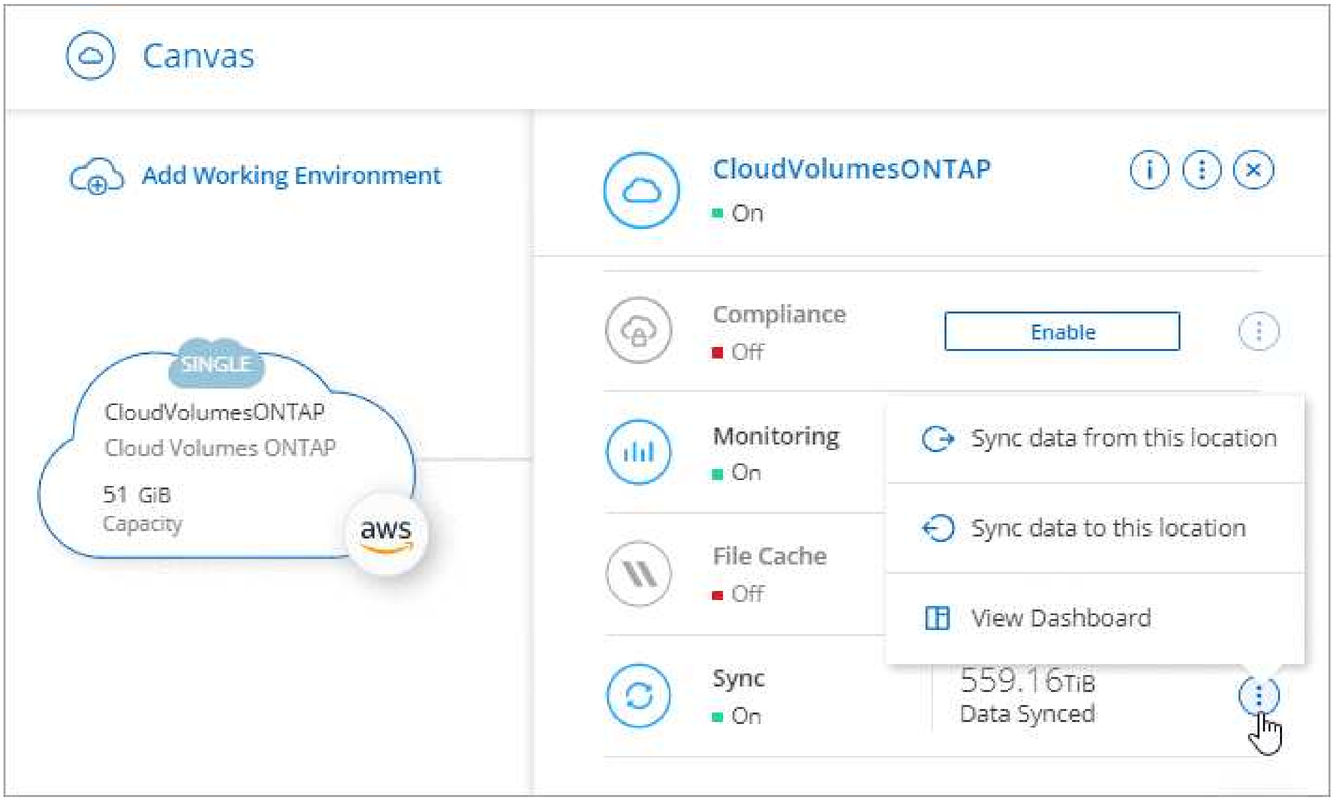

6. Select **Sync data from this location** or **Sync data to this location** and follow the prompts to set up the sync relationship.

Create other types of sync relationships

Use these steps to sync data to or from a supported storage type other than Amazon FSx for ONTAP, Azure NetApp Files, Cloud Volumes ONTAP, or on-prem ONTAP clusters. The steps below provide an example that shows how to set up a sync relationship from an NFS server to an S3 bucket.

1. In the NetApp Console, select **Sync**.
2. On the **Define Sync Relationship** 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.

   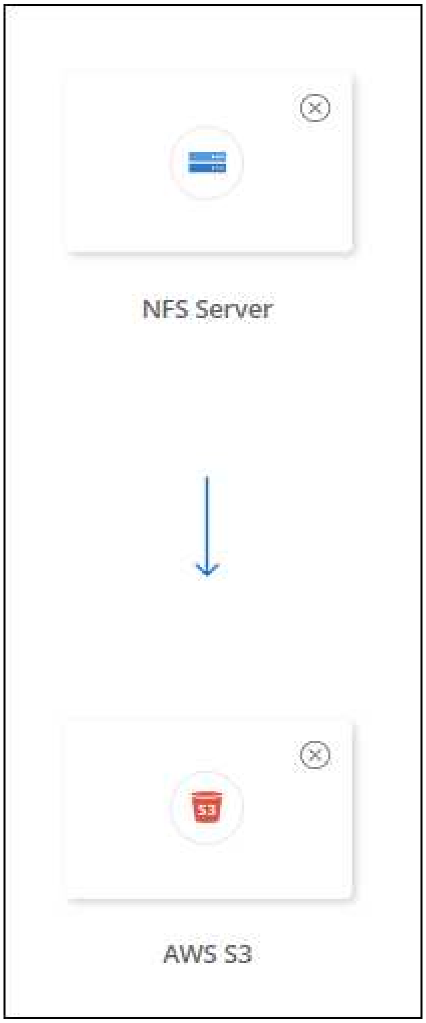

3. On the **NFS Server** page, enter the IP address or fully qualified domain name of the NFS server that you want to sync to AWS.
4. On the **Data Broker Group** page, follow the prompts to create a data broker virtual machine in AWS, Azure, or Google Cloud Platform, or to install the data broker software an existing Linux host.

   For more details, refer to the following pages:

   - Create a data broker in AWS

- Create a data broker in Azure
- Create a data broker in Google Cloud
- Installing the data broker on a Linux host

5. After you install the data broker, select **Continue**.

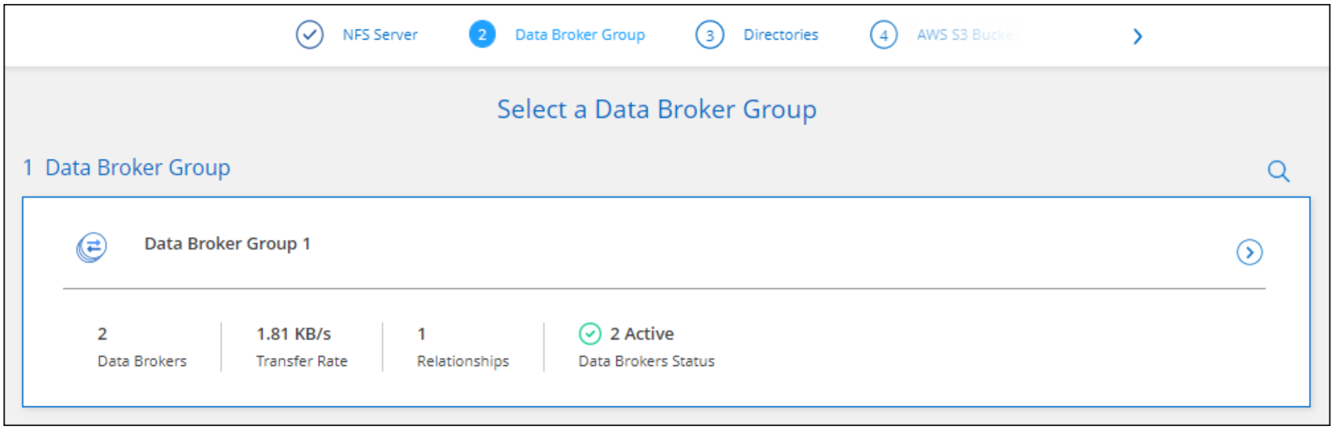

Select a Data Broker Group

1 Data Broker Group

Data Broker Group 1

2	1.81 KB/s	1	✓ 2 Active
Data Brokers	Transfer Rate	Relationships	Data Brokers Status

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

   If Copy and Sync is unable to retrieve the exports, select **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.

7. 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.
   - Select **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.

8. 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.

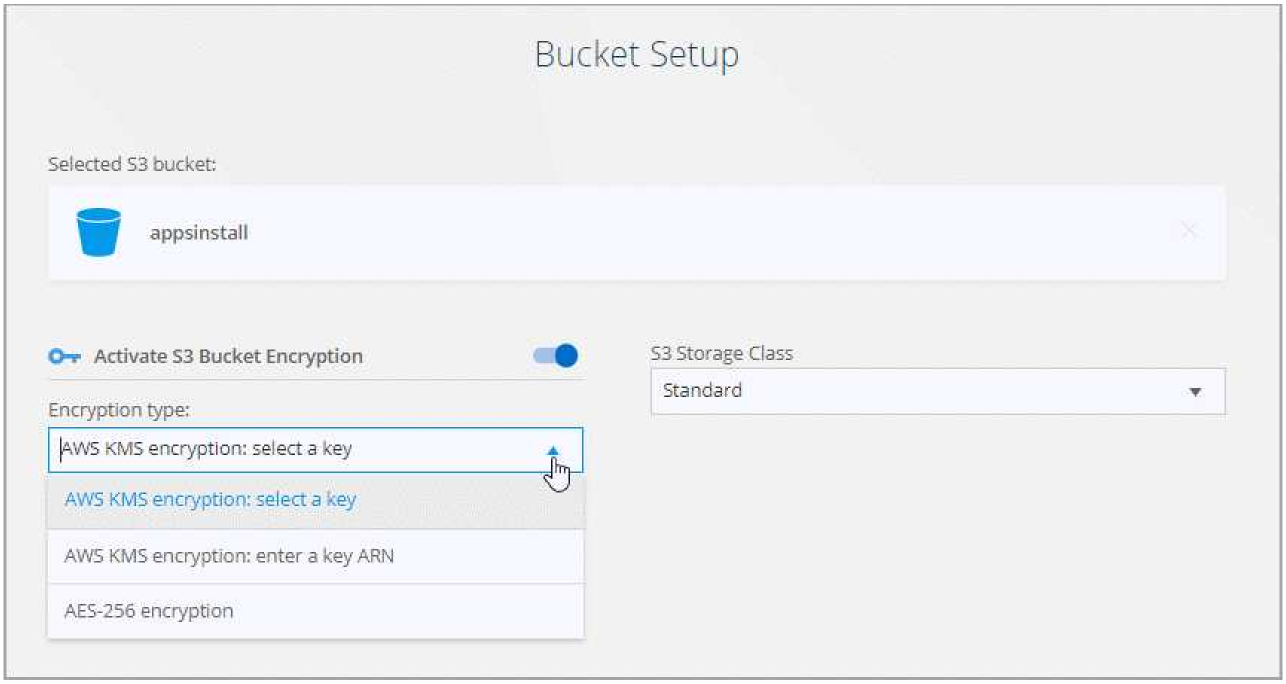

Bucket Setup

Selected S3 bucket: appsinstall

Activate S3 Bucket Encryption

S3 Storage Class: Standard

Encryption type: AWS KMS encryption: select a key

AWS KMS encryption: select a key
AWS KMS encryption: enter a key ARN
AES-256 encryption

9. 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.

Sync Timeout

Define whether Copy and Sync should cancel a data sync if the sync hasn't completed in the specified number of minutes,hours, or days.

Notifications

Enables you to choose whether to receive Copy and Sync notifications in the NetApp Console's Notification Center. You can enable notifications for successful data syncs, failed data syncs, and canceled data syncs.

Retries

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

Continuous Sync

After the initial data sync, Copy and Sync listens for changes on the source S3 bucket or Google Cloud Storage bucket and continuously syncs any changes to the target as they occur. There's no need to rescan the source at scheduled intervals.

This setting is available only when creating a sync relationship and when you sync data from an S3 bucket or Google Cloud Storage to Azure Blob storage, CIFS, Google Cloud Storage, IBM Cloud Object Storage, NFS, S3, and StorageGRID **or** from Azure Blob storage to Azure Blob storage, CIFS, Google Cloud Storage, IBM Cloud Object Storage, NFS, and StorageGRID.

If you enable this setting, it affects other features as follows:

- The sync schedule is disabled.

- The following settings are reverted to their default values: Sync Timeout, Recently Modified Files, and Date Modified.
- If S3 is the source, filter by size will be active only on copy events (not on delete events).
- After the relationship is created, you can only accelerate or delete the relationship. You can’t abort syncs, modify settings, or view reports.

  It is possible to create a Continuous Sync relationship with an external bucket. To do so, follow these steps:

  a. Go to the Google Cloud console for the external bucket’s project.

  b. Go to **Cloud Storage > Settings > Cloud Storage Service Account**.

  c. Update the local.json file:

  ```
  {
  "protocols": {
      "gcp": {
          "storage-account-email": <storage account email>
  }
  }
  }
  ```

  d. Restart the data broker:

     i. sudo pm2 stop all

     ii. sudo pm2 start all

  e. Create a Continuous Sync relationship with the relevant external bucket.

  A data broker used to create a continuous sync relationship with an external bucket will not be able to create another Continuous Sync relationship with a bucket in its project.

Compare By

Choose whether Copy and Sync should compare certain attributes when determining whether a file or directory has changed and should be synced again.

Even if you uncheck these attributes, Copy and Sync still compares the source to the target by checking the paths, file sizes, and file names. If there are any changes, then it syncs those files and directories.

You can choose to enable or disable Copy and Sync from comparing the following attributes:

- **mtime**: The last modified time for a file. This attribute isn’t valid for directories.
- **uid**, **gid**, and **mode**: Permission flags for Linux.

Copy for Objects

Enable this option to copy object storage metadata and tags. If a user changes the metadata on the source, Copy and Sync copies this object in the next sync, but if a user changes the tags on the source (and not the data itself), Copy and Sync doesn’t copy the object in the next sync.

You can’t edit this option after you create the relationship.

Copying tags is supported with sync relationships that include Azure Blob or an S3-compatible endpoint (S3, StorageGRID, or IBM Cloud Object Storage) as the target.

Copying metadata is supported with "cloud-to-cloud" relationships between any of the following endpoints:

- AWS S3
- Azure Blob
- Google Cloud Storage
- IBM Cloud Object Storage
- StorageGRID

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 Copy and 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 update it as follows:

```
{
"workers":{
"transferrer":{
"delete-on-source": true
}
}
}
```

After updating the local.json file, you should do a restart: `pm2 restart all`.

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.

File Types

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

> ⓘ Hard links are only available for unsecured NFS to NFS relationships. Users will be limited to one scanner process and one scanner concurrency, and scans must be run from a root directory.

Exclude File Extensions

Specify the regex or 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:

Exclude file extensions for a sync relationship

Regex, or regular expressions, differ from wildcards or glob expressions. This feature **only** works with regex.

Exclude Directories

Specify a maximum of 15 regex or directories to exclude from the sync by typing their name or directory full path and pressing **Enter**. The .copy-offload, .snapshot, ~snapshot directories are excluded by default.

Regex, or regular expressions, differ from wildcards or glob expressions. This feature **only** works with regex.

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.

Date Created

When an SMB server is the source, this setting enables you to sync files that were created after a specific date, before a specific date, or between a specific time range.

ACL - Access Control List

Copy ACLs only, files only, or ACLs and files from an SMB server by enabling a setting when you create a relationship or after you create a relationship.

10. On the **Tags/Metadata** page, choose whether to save a key-value pair as a tag on all files transferred to the S3 bucket or to assign a metadata key-value pair on all files.

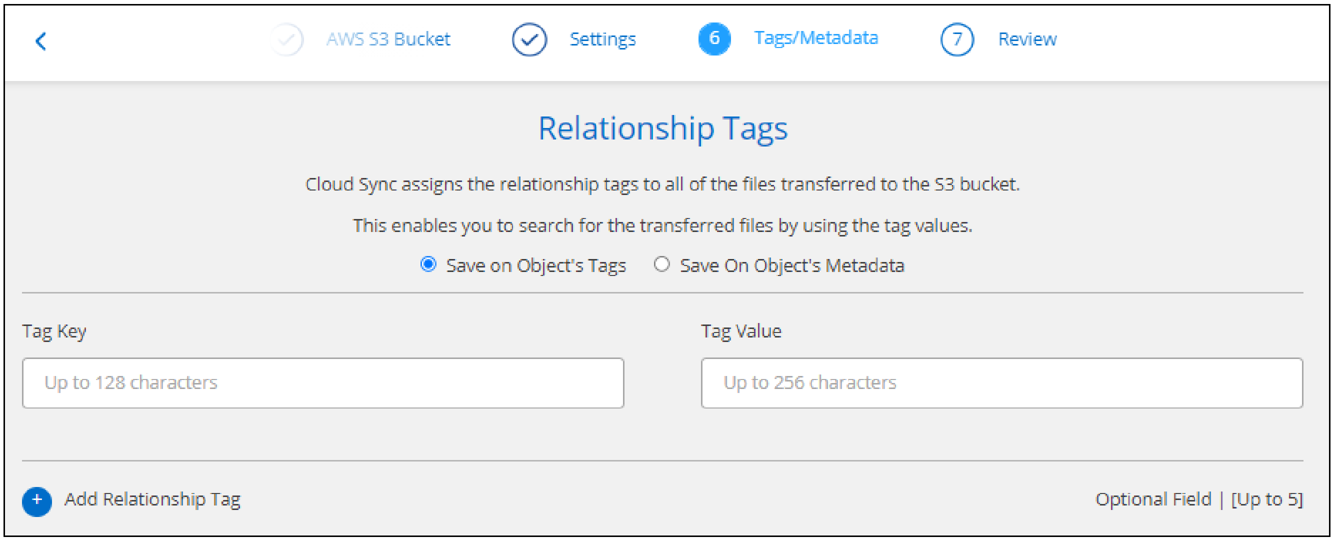

Relationship Tags

Cloud Sync assigns the relationship tags to all of the files transferred to the S3 bucket.

This enables you to search for the transferred files by using the tag values.

● Save on Object's Tags ○ Save On Object's Metadata

Tag Key: Up to 128 characters

Tag Value: Up to 256 characters

+ Add Relationship Tag

Optional Field | [Up to 5]

This same feature is available when syncing data to StorageGRID and IBM Cloud Object Storage. For Azure and Google Cloud Storage, only the metadata option is available.

11. Review the details of the sync relationship and then select **Create Relationship**.

Result

Copy and Sync starts syncing data between the source and target. Sync statistics on how long the sync took, whether it halted, and how many files were copied, scanned, or deleted are available. You can then manage your sync relationships, manage your data brokers, or create reports to optimize your performance and configuration.

Create sync relationships from NetApp Data Classification

Copy and Sync is integrated with NetApp Data Classification. From within NetApp Data Classification, you can select the source files that you'd like to sync to a target location using Copy and Sync.

After you initiate a data sync from NetApp Data Classification, all of the source information is contained in a single step and only requires you to enter a few key details. You then choose the target location for the new sync relationship.

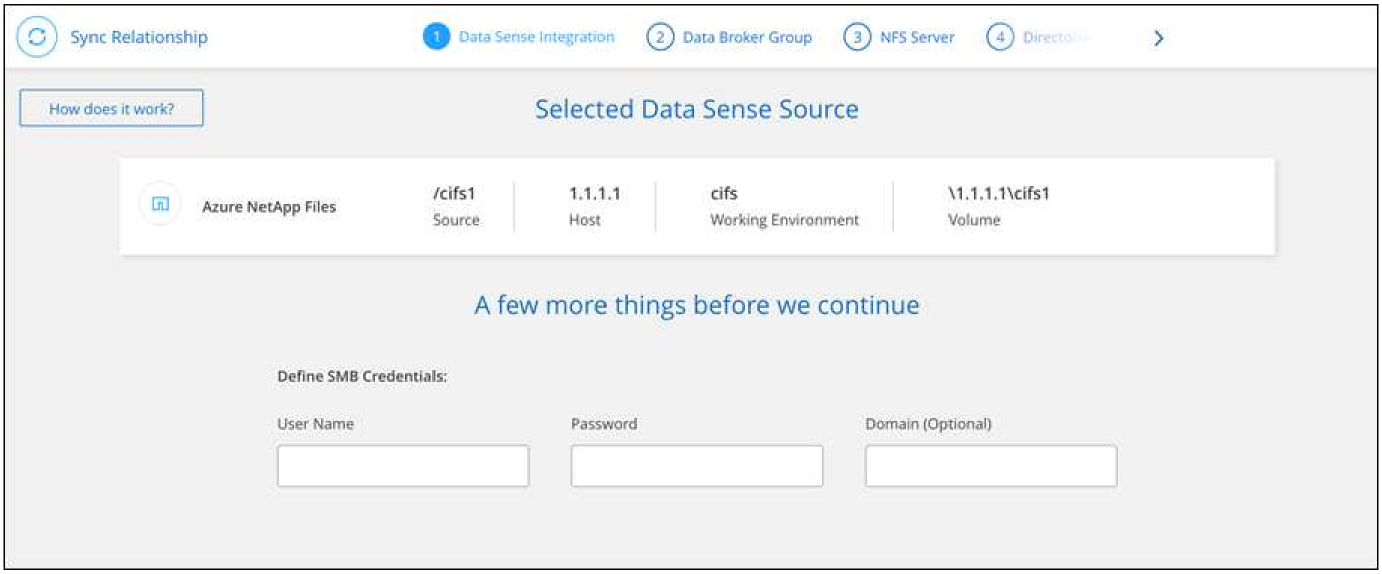

Selected Data Sense Source

Azure NetApp Files	/cifs1 Source	1.1.1.1 Host	cifs Working Environment	\1.1.1.1\cifs1 Volume

A few more things before we continue

Define SMB Credentials:

User Name

Password

Domain (Optional)

Learn how to start a sync relationship from NetApp Data Classification.

Copy ACLs from SMB shares in NetApp Copy and Sync

NetApp Copy and Sync can copy access control lists (ACLs) between SMB shares and between an SMB share and object storage (except for ONTAP S3). If needed, you also have the option to manually preserve ACLs between SMB shares by using robocopy.

Choices

- Set up Copy and Sync to automatically copy ACLs
- Manually copy the ACLs between SMB shares

Set up Copy and Sync to copy ACLs

Copy ACLs between SMB shares and between SMB shares and object storage by enabling a setting when you create a relationship or after you create a relationship.

Before you begin

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.

Steps for a new relationship

1. Log in to Copy and Sync.
2. From Copy and Sync, select **Create New Sync**.
3. Drag and drop an SMB server or object storage as the source and an SMB server or object storage as the target, and select **Continue**.
4. On the **SMB Server** page:
   a. Enter a new SMB server or select an existing server and select **Continue**.
   b. Enter credentials for the SMB server.
   c. Choose to either **Copy only files**, **Copy only ACL**, or **Copy files and ACL** and select **Continue**.

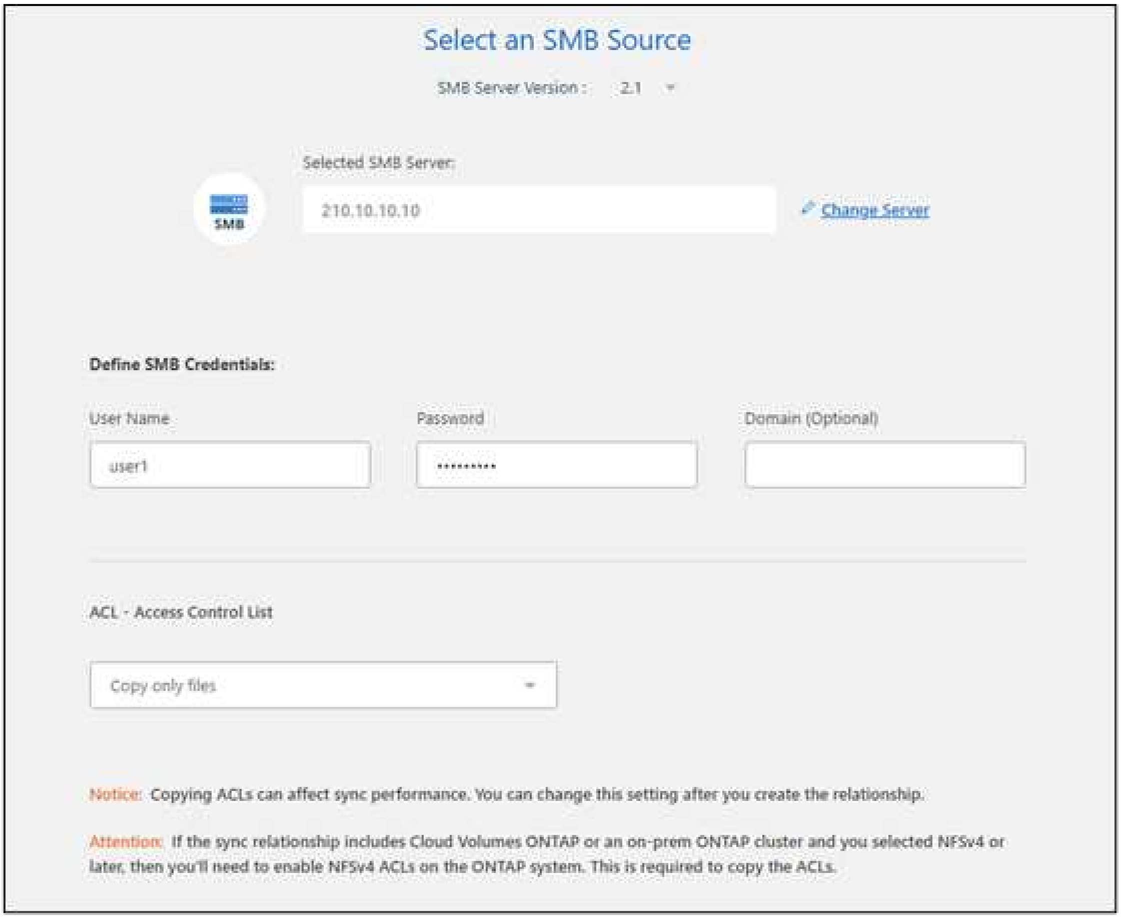

Select an SMB Source

SMB Server Version : 2.1

Selected SMB Server: 210.10.10.10

Define SMB Credentials:

User Name: user1

Password: ••••••••••

Domain (Optional)

ACL - Access Control List: Copy only files

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

Attention: 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.

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

   When you copy ACLs from SMB to object storage, you can choose to copy the ACLs to the object's tags or on the object's metadata, depending on the target. For Azure and Google Cloud Storage, only the metadata option is available.

   The following screenshot shows an example of the step where you can make this choice.

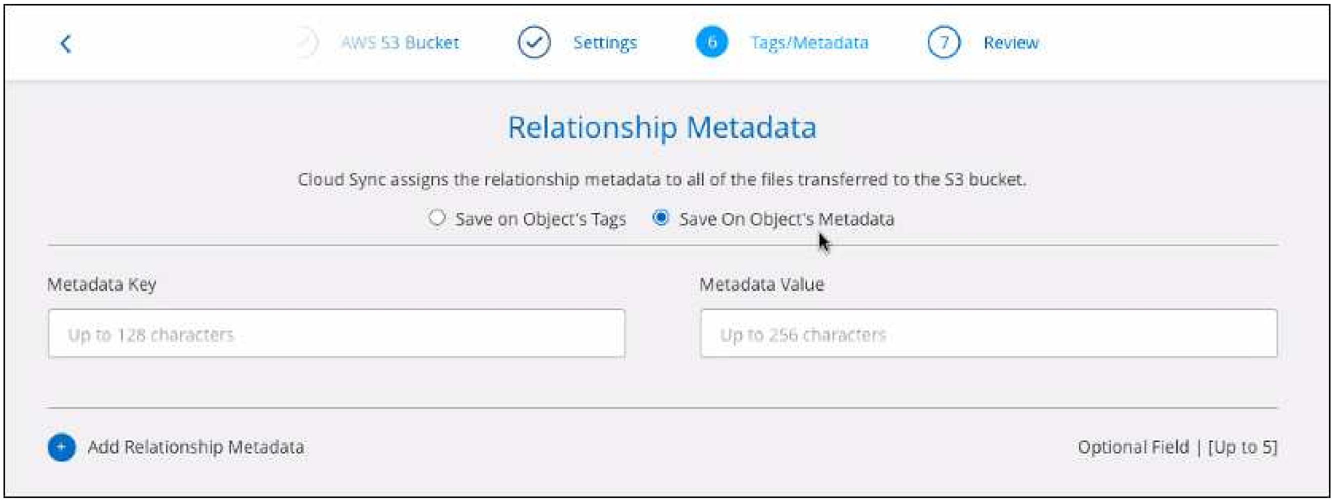

Relationship Metadata

Cloud Sync assigns the relationship metadata to all of the files transferred to the S3 bucket.

○ Save on Object's Tags ● Save On Object's Metadata

Metadata Key: Up to 128 characters

Metadata Value: Up to 256 characters

Steps for an existing relationship

1. Hover over the sync relationship and select the action menu.
2. Select **Settings**.
3. Choose to either **Copy only files**, **Copy only ACL**, or **Copy files and ACL** and select **Continue**.
4. Select **Save Settings**.

> Copy and Sync preserves SMB ACLs (permissions), but does not copy file or folder ownership. Ownership is not included in the SMB ACL transfer operation.

Result

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

Manually copy ACLs between SMB shares

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

> If you need to preserve ownership (Owner and Group) in addition to ACLs, you can use the `robocopy` command.
> Using the `/copyall` flag copies ACLs, ownership, and audit information.

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 Copy and 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>

Sync NFS data using data-in-flight encryption in NetApp Copy and Sync

If your business has strict security policies, you can sync NFS data using data-in-flight encryption in NetApp Copy and Sync. 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:

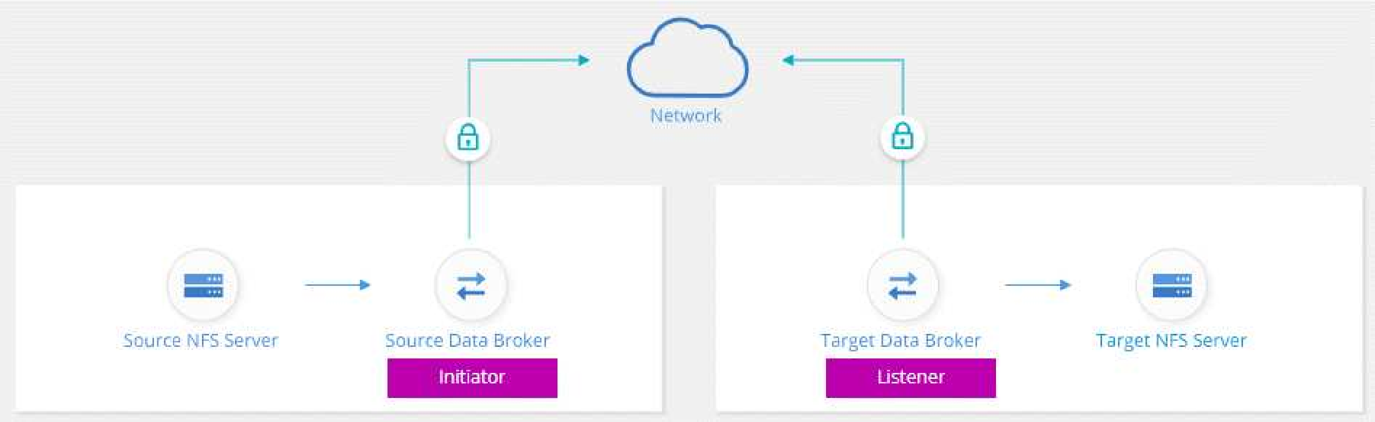

One data broker functions as the *initiator*. When it's time to sync data, it sends a connection request to the other data broker, which is the *listener*. That data broker listens for requests on port 443. You can use a different port, if needed, but be sure to check that the port is not in use by another 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, Copy and Sync 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 want to use an existing data broker group, the group must have only one data broker. Multiple data brokers in a group aren't supported with encrypted sync relationships.

  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 Google Cloud installation
  - Review Linux host installation

Sync 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 Copy and Sync.
2. Select **Create New Sync**.
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.
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.

   Note the following:

   - If you want to use an existing data broker group, the group must have only one data broker. Multiple data brokers in a group aren't supported with encrypted sync relationships.
   - If the source data broker acts as the listener, then it must be a new data broker.
   - If you need a new data broker, Copy and 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.

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

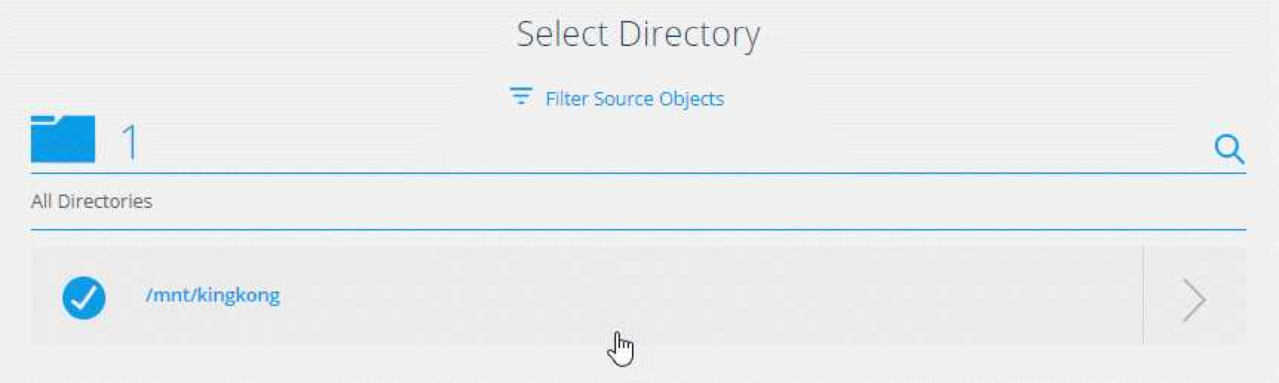

Select Directory

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.

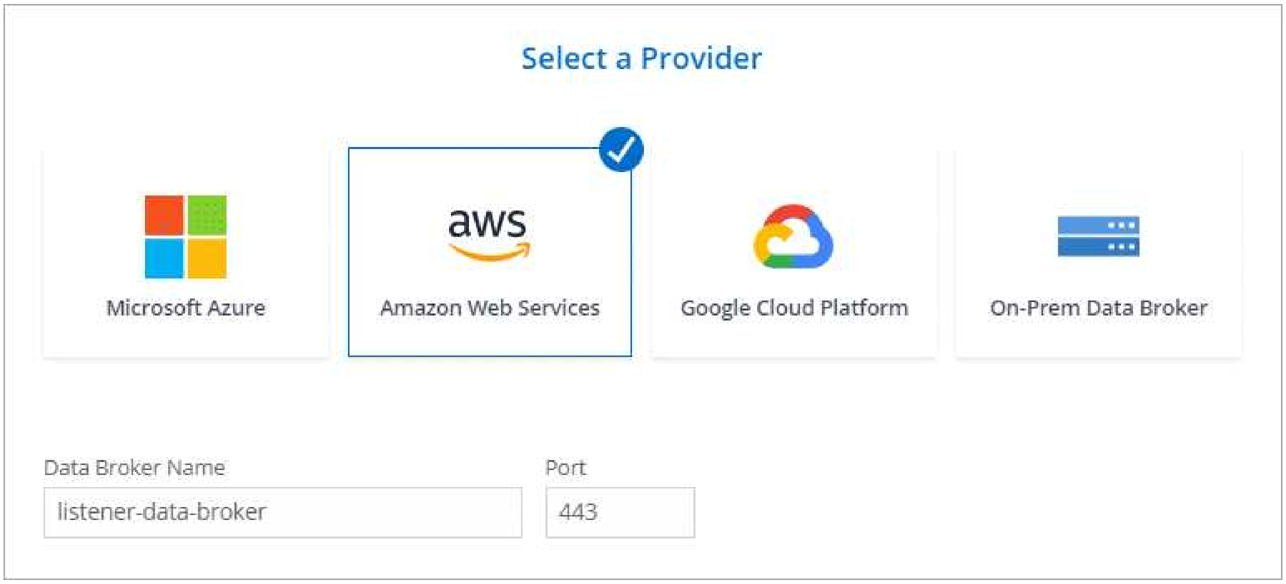

Data Broker Name: listener-data-broker

Port: 443

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 select **Create Relationship**.

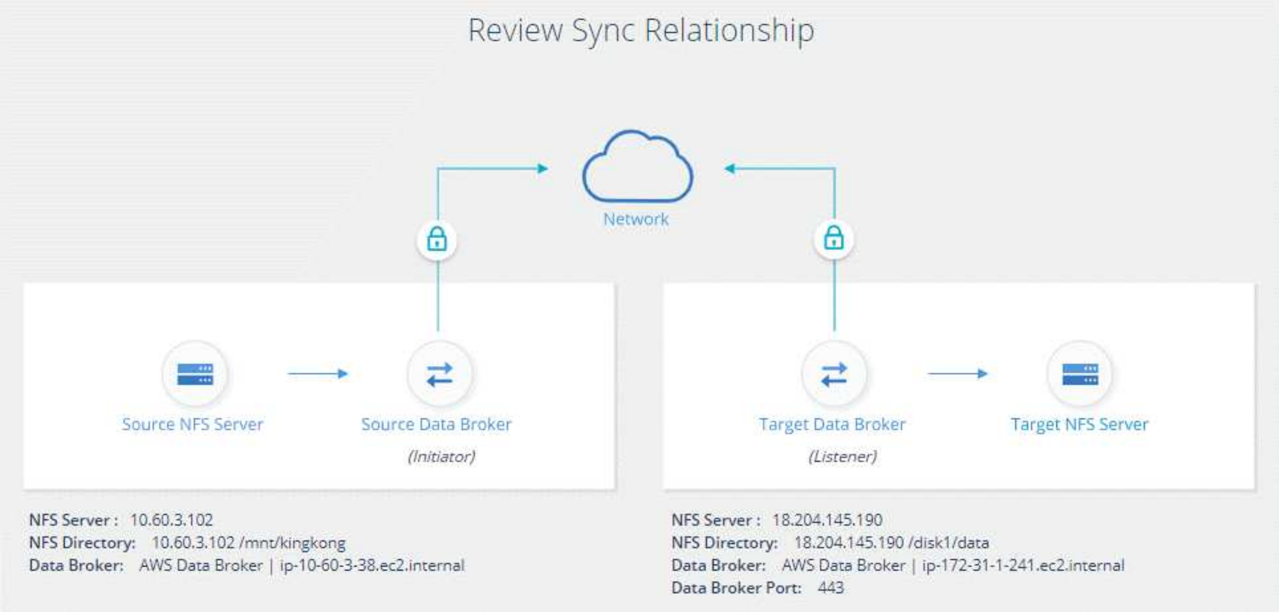

NFS Server : 10.60.3.102
NFS Directory: 10.60.3.102 /mnt/kingkong
Data Broker: AWS Data Broker | ip-10-60-3-38.ec2.internal

NFS Server : 18.204.145.190
NFS Directory: 18.204.145.190 /disk1/data
Data Broker: AWS Data Broker | ip-172-31-1-241.ec2.internal
Data Broker Port: 443

Result

Copy and Sync starts creating the new sync relationship. When it’s done, select **View in Dashboard** to view details about the new relationship.

Set up a data broker group to use an external HashiCorp Vault in NetApp Copy and Sync

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

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

1 **Prepare the vault**

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

2 **Prepare the data broker group**

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

3 **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.

Prepare the vault

You'll need to provide Copy and 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/hb312vdasr2/S3Creds

As you can see in the example, the prefix path is */my-path/all-secrets/*, the request ID is *hb312vdasr2* 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.

Prepare the data broker group

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

Steps

1. SSH to a data broker in the group.
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

   - The authentication method that the data broker should use to access credentials from the external vault
   - Type: string
   - Valid values: “aws-iam” / “role-app” / "gcp-iam"

   role-name

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

Secretid & rootid

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

Namespace

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

4. Repeat these steps for any other data brokers in the group.

Example for aws-role authentication

```
{
          “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”
                         }
                }
       }
}
```

Example for gcp-iam authentication

```
{
"external-integrations": {
    "hashicorp": {
      "enabled": true,
      "url": http://ip-10-20-30-55.ec2.internal:8200,
      "path": "v1/secret",
      "namespace": "",
      "reject-unauthorized": true,
      "auth-method": "gcp-iam",
      "aws-iam": {
        "role-name": ""
      },
      "app-role": {
        "root_id": "",
        "secret_id": ""
      },
"gcp-iam": {
          "role-name": "my-iam-role"
      }
    }
  }
}
```

Set up permissions when using gcp-iam authentication

If you’re using the *gcp-iam* authentication method, then the data broker must have the following GCP permission:

```
- iam.serviceAccounts.signJwt
```

Learn more about GCP permission requirements for the data broker.

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 Copy and 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 NetApp Console 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": "5e6e111d578dtyuu1555sa60",
"source": {
        "protocol": "s3",
        "s3": {
                "provider": "sgws",
                "host": "1.1.1.1",
                "port": "443",
                "bucket": "my-source"
     },
"target": {
        "protocol": "s3",
        "s3": {
                "bucket": "my-target-bucket"
        }
    }
}
```

Pay for sync relationships after your NetApp Copy and Sync free trial ends

There are two ways to pay for sync relationships after your 14-day free trial ends for NetApp Copy and Sync. 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 subscribe from either the AWS Marketplace or the Azure Marketplace. You can’t subscribe from both.

You have the option to use licenses from NetApp with a marketplace 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.

If you don’t immediately pay after your 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.

Licenses should be managed through NetApp Copy and Sync or the applicable website and **not** through the NetApp Console subscriptions.

Subscribe from AWS

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

Steps to pay-as-you-go

1. From the NetApp Console navigation menu, select **Mobility > Copy and Sync**.
2. Select **Licensing**.
3. Select **AWS**.
4. Select **Subscribe** and then select **Continue**.
5. Subscribe from the AWS Marketplace, and then log back in to Copy and Sync to complete the registration.

   The following video shows the process:

   Subscribe to Copy and Sync from the AWS Marketplace

Steps to pay annually

1. Go to the AWS Marketplace page.
2. Select **Continue to Subscribe**.
3. Select your contract options and then select **Create contract**.

Subscribe 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. From the NetApp Console navigation menu, select **Mobility > Copy and Sync**.
2. Select **Licensing**.
3. Select **Azure**.
4. Select **Subscribe** and then select **Continue**.
5. In the Azure portal, select **Create**, select your options, and then select **Subscribe**.

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

6. When deployment is complete, select the name of the SaaS resource in the notification pop-up.
7. Select **Configure Account** to return to Copy and Sync.

   The following video shows the process:

   Subscribe to Copy and Sync from the Azure Marketplace

Purchase licenses from NetApp and add them to Copy and Sync

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

What you'll need

You'll need the serial number for your license and the user name and password for the NetApp Support Site account that the license is associated with.

Steps

1. Purchase a license by contacting NetApp.
2. Log in to Copy and Sync.
3. Select **Licensing**.
4. Select **Add License** and add the required information:
   a. Enter the serial number.
   b. Select the NetApp Support Site account that is associated with the license that you're adding:
      - If your account was already added to the NetApp Console, select it from the drop-down list.
      - If your account wasn't added yet, select **Add NSS Credentials**, enter the user name and password, select **Register**, and then select it from the drop-down list.
   c. Select **Add**.

Update a license

If you extended a Copy and Sync license that you purchased from NetApp, the new expiration date won't update automatically in Copy and Sync. You need to add the license again to refresh the expiration date. Licenses should be managed through Copy and Sync or the applicable website and **not** through the NetApp Console subscriptions.

Steps

1. From the NetApp Console navigation menu, select **Mobility > Copy and Sync**.

2. Select **Licensing**.
3. Select **Add License** and add the required information:
    a. Enter the serial number.
    b. Select the NetApp Support Site account that is associated with the license that you're adding.
    c. Select **Add**.

Result

Copy and Sync updates the existing license with the new expiration date.

Managing sync relationships in NetApp Copy and Sync

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

Perform an immediate data sync

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

Steps

1. Log in to Copy and Sync.
2. From the **Dashboard**, navigate to the sync relationship and select ⋮
3. Select **Sync Now** and then select **Sync** to confirm.

Result

Copy and Sync starts the data sync process for the relationship.

Accelerate sync performance

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

How this works

If the data broker group manages other sync relationships, then the new data broker that you add to the group also accelerates the performance of those sync relationships.

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

- Relationship 1 is managed by data broker group A
- Relationship 2 is managed by data broker group B
- Relationship 3 is managed by data broker group A

You want to accelerate the performance of relationship 1 so you add a new data broker to data broker group A. Because group A also manages sync relationship 3, the sync performance for relationship is automatically accelerated as well.

Steps

1. Ensure that at least one of the existing data brokers in the relationship are online.

2. From the **Dashboard**, navigate to the sync relationship and select ⋮
3. Select **Accelerate**.
4. Follow the prompts to create a new data broker.

Result

Copy and Sync adds the new data broker to the group. The performance of the next data sync should be accelerated.

Update credentials

You can update the data broker with the latest credentials of the source or target in an existing sync relationship. Updating the credentials can help if your security policies require you to update credentials on a periodic basis.

Updating credentials is supported with any source or target that Copy and Sync requires credentials for: Azure Blob, Box, IBM Cloud Object Storage, StorageGRID, ONTAP S3 Storage, SFTP, and SMB servers.

Steps

1. From the **Sync Dashboard**, go to a sync relationship that requires credentials and then select **Update Credentials**.

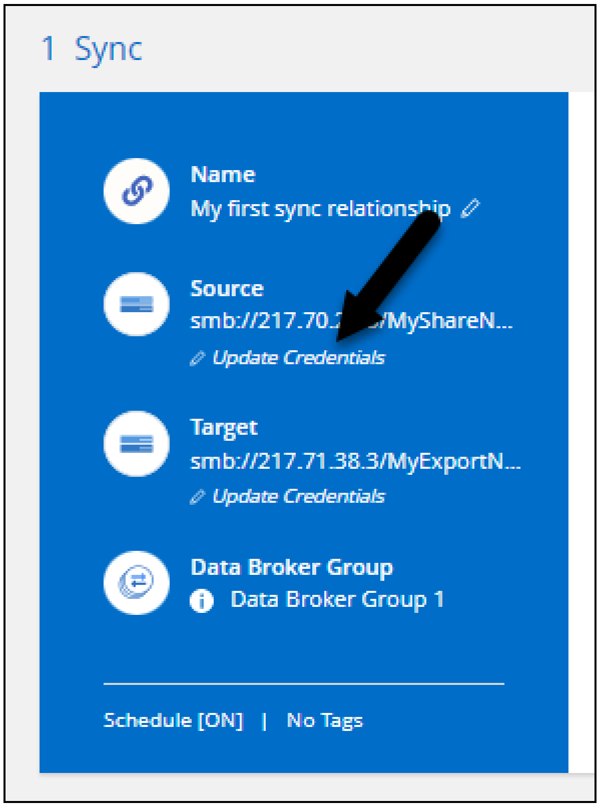

1 Sync

Name
My first sync relationship

Source
smb://217.70.2 /MyShareN...
Update Credentials

Target
smb://217.71.38.3/MyExportN...
Update Credentials

Data Broker Group
Data Broker Group 1

Schedule [ON] | No Tags

2. Enter the credentials and select **Update**.

   A note about SMB servers: if the domain is new, then you'll need to specify it when you update the credentials. If the domain hasn't changed, then you don't need to enter it again.

   If you entered a domain when you created the sync relationship, but you don't enter a new domain when you update the credentials, then Copy and Sync will keep using the original domain that you provided.

Result

Copy and Sync updates the credentials on the data broker. It can take up 10 minutes until the data broker starts using the updated credentials for data syncs.

Set up notifications

A **Notifications** setting for each sync relationship enables you to choose whether to receive Copy and Sync notifications in the NetApp Console’s Notification Center. You can enable notifications for successful data syncs, failed data syncs, and canceled data syncs.

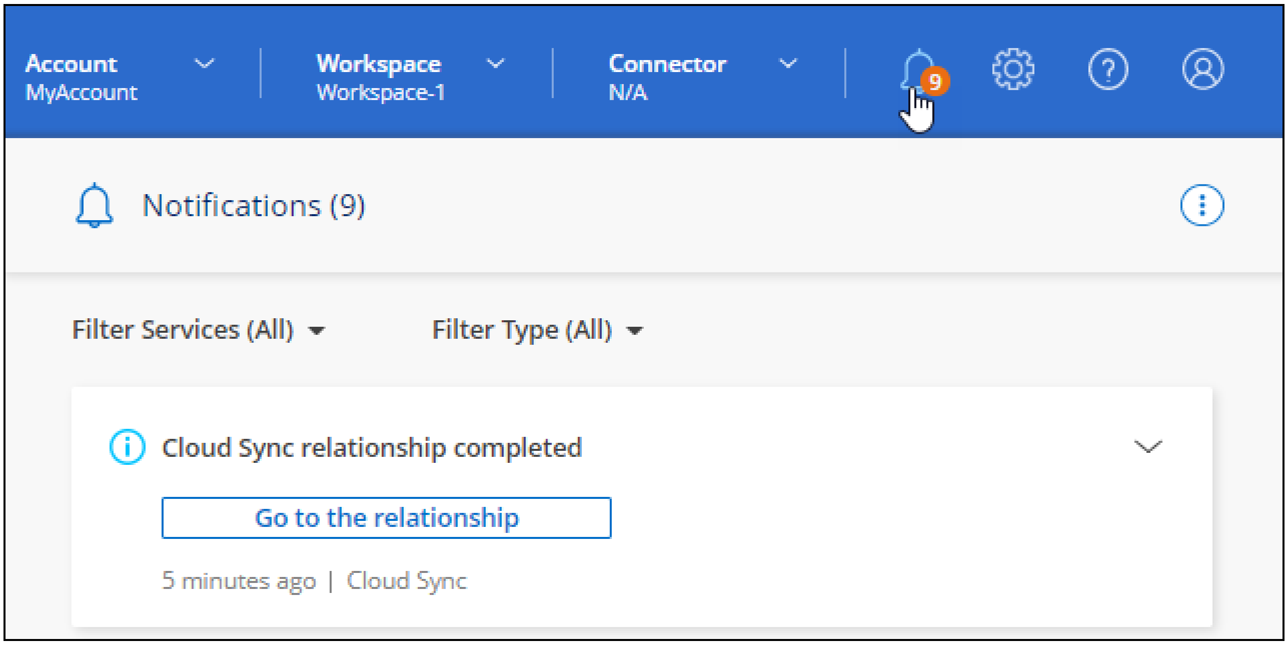

Cloud Sync relationship completed

5 minutes ago | Cloud Sync

In addition, you can also receive notifications by email.

Steps

1. Modify the settings for a sync relationship:
   a. From the **Dashboard**, navigate to the sync relationship and select ⋮
   b. Select **Settings**.
   c. Enable **Notifications**.
   d. Select **Save Settings**.
2. If you want to receive notifications by email, configure alert and notifications settings:
   a. Select **Settings > Alerts and Notifications Settings**.
   b. Select a user or multiple users and choose the **Info** notification type.
   c. Select **Apply**.

Result

You’ll now receive Copy and Sync notifications in the NetApp Console’s Notification Center, with a few notifications arriving by email, if you configured that option.

Change the settings for a sync relationship

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

1. From the **Dashboard**, navigate to the sync relationship and select ⋮
2. Select **Settings**.
3. Modify any of the settings.

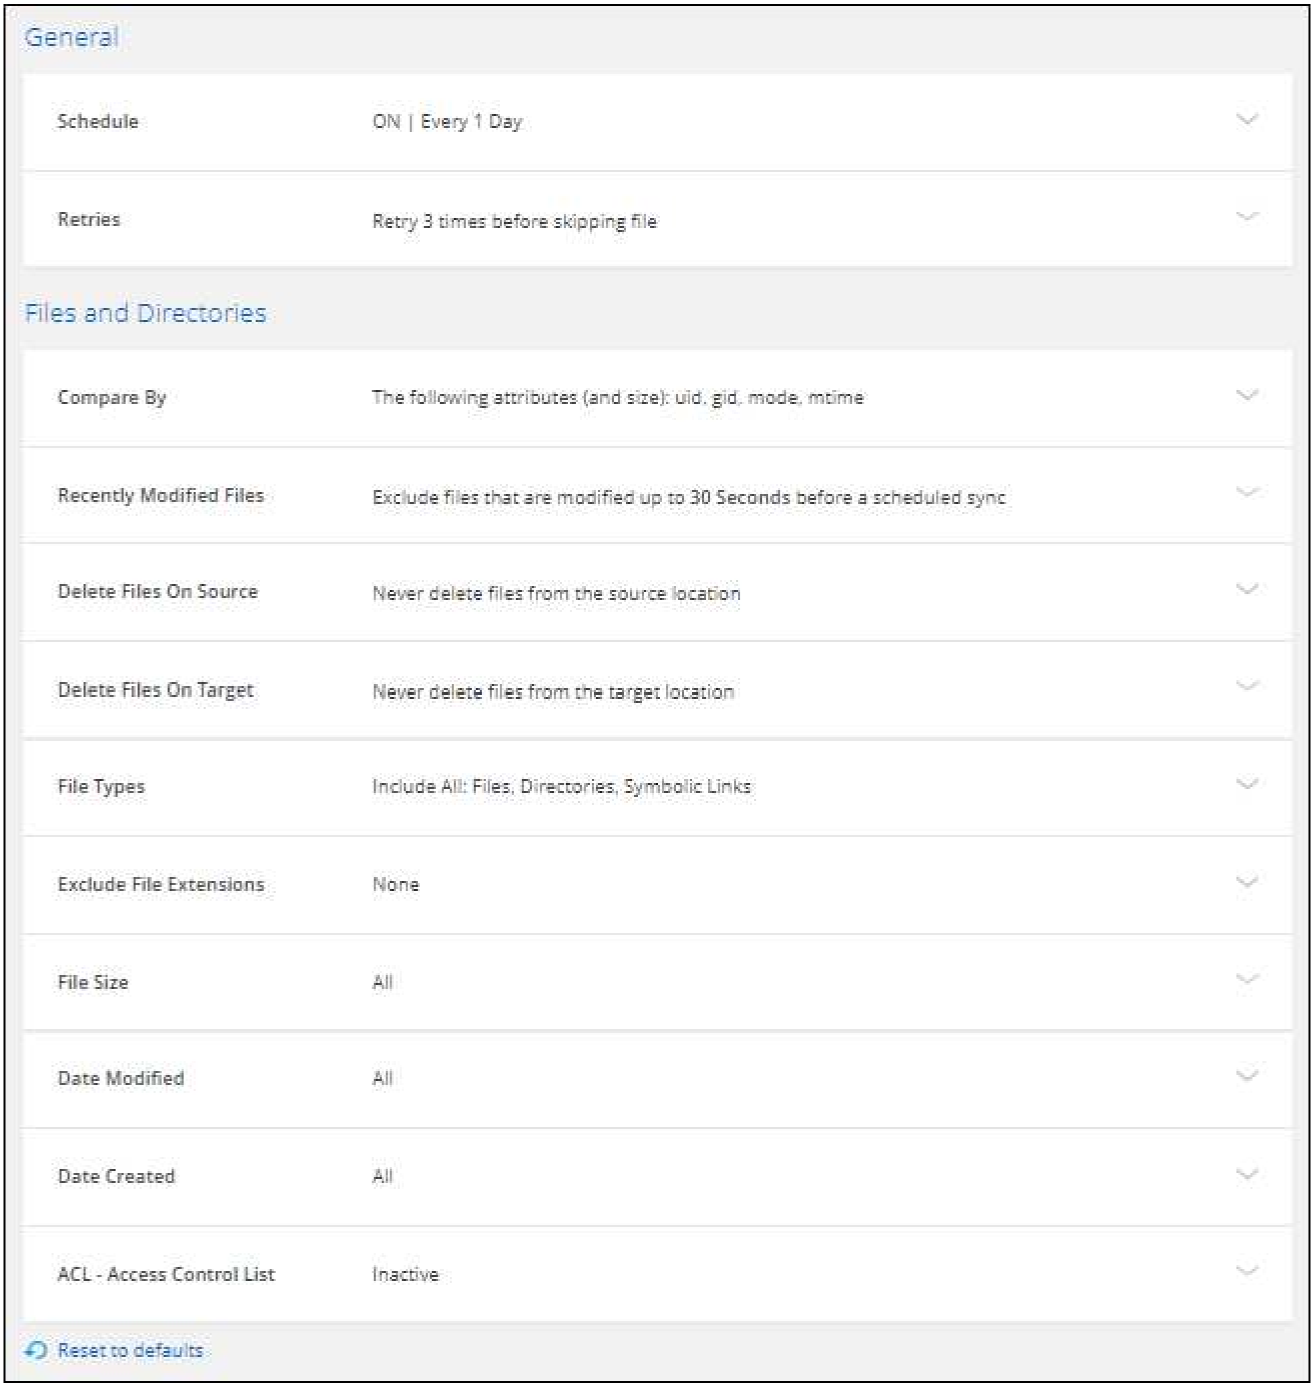

General

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

Files and Directories

Compare By	The following attributes (and size): uid, gid, mode, mtime
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
File Types	Include All: Files, Directories, Symbolic Links
Exclude File Extensions	None
File Size	All
Date Modified	All
Date Created	All
ACL - Access Control List	Inactive

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.

Sync Timeout

Define whether Copy and Sync should cancel a data sync if the sync hasn’t completed in the specified number of minutes, hours, or days.

Notifications

Enables you to choose whether to receive Copy and Sync notifications in the NetApp Console’s Notification Center. You can enable notifications for successful data syncs, failed data syncs, and canceled data syncs.

If you want to receive notifications for

Retries

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

Compare By

Choose whether Copy and Sync should compare certain attributes when determining whether a file or directory has changed and should be synced again.

Even if you uncheck these attributes, Copy and Sync still compares the source to the target by checking the paths, file sizes, and file names. If there are any changes, then it syncs those files and directories.

You can choose to enable or disable Copy and Sync from comparing the following attributes:

- **mtime**: The last modified time for a file. This attribute isn’t valid for directories.
- **uid**, **gid**, and **mode**: Permission flags for Linux.

Copy for Objects

You can’t edit this option after you create the relationship.

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 Copy and 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 update it as follows:

```
{
"workers":{
"transferrer":{
"delete-on-source": true
}
}
}
```

After updating the local.json file, you should do a restart: `pm2 restart all`.

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.

File Types

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

> Hard links are only available for unsecured NFS to NFS relationships. Users will be limited to one scanner process and one scanner concurrency, and scans must be run from a root directory.

Exclude File Extensions

Specify the regex or 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:

Exclude file extensions for a sync relationship

> Regex, or regular expressions, differ from wildcards or glob expressions. This feature **only** works with regex.

Exclude Directories

Specify a maximum of 15 regex or directories to exclude from the sync by typing their name or directory full path and pressing **Enter**. The .copy-offload, .snapshot, ~snapshot directories are excluded by default.

> Regex, or regular expressions, differ from wildcards or glob expressions. This feature **only** works with regex.

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.

Date Created

When an SMB server is the source, this setting enables you to sync files that were created after a specific date, before a specific date, or between a specific time range.

ACL - Access Control List

Copy ACLs only, files only, or ACLs and files from an SMB server by enabling a setting when you create a relationship or after you create a relationship.

4. Select **Save Settings**.

Result

Copy and Sync modifies the sync relationship with the new settings.

Delete relationships

You can delete a sync relationship, if you no longer need to sync data between the source and target. This action doesn’t delete the data broker group (or the individual data broker instances) and it does not delete data from the target.

Option 1: Delete a single sync relationship

Steps

1. From the **Dashboard**, navigate to the sync relationship and select ⋮
2. Select **Delete** and then select **Delete** again to confirm.

Result

Copy and Sync deletes the sync relationship.

Option 2: Delete multiple sync relationships

Steps

1. From the **Dashboard**, navigate to the "Create New Sync" button and select ⋮
2. Select the sync relationships you want to delete, select **Delete** and then select **Delete** again to confirm.

Result

Copy and Sync deletes the sync relationships.

Manage data broker groups in NetApp Copy and Sync

A data broker group in NetApp Copy and Sync syncs data from a source location to a target location. At least one data broker is required in a group for each sync relationship that you create. Manage data broker groups by adding a new data broker to a group, by viewing information about groups, and more.

How data broker groups work

A data broker group can include one or more data brokers. Grouping data brokers together can help improve the performance of sync relationships.

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 is managed by data broker group A
- Relationship 2 is managed by data broker group B
- Relationship 3 is managed by data broker group A

You want to accelerate the performance of relationship 1 so you add a new data broker to data broker group A. Because group A also manages sync relationship 3, the sync performance for relationship is automatically accelerated as well.

Number of data brokers in a group

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.

Security recommendations

To ensure the security of your data broker machine, NetApp recommends the following:

- SSH should not permit X11 Forwarding
- SSH should not permit TCP connection forwarding
- SSH should not permit tunnels
- SSH should not accept client environment variables

These security recommendations can help prevent unauthorized connections to the data broker machine.

Add a new data broker to a group

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 selecting **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

Before you get started

- 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 Google Cloud data broker because they aren’t the same data broker type.

Steps to create a data broker in a new group

1. Log in to Copy and Sync.
2. Select **Sync > Manage Data Brokers**.
3. Select **Add New Data Broker**.
4. Follow the prompts to create the data broker.

   For help, refer to the following pages:

   - Create a data broker in AWS
   - Create a data broker in Azure
   - Create a data broker in Google Cloud
   - Installing the data broker on a Linux host

Steps to create a data broker in an existing group

1. Log in to Copy and Sync.
2. Select **Sync > Manage Data Brokers**.
3. Select the action menu and select **Add Data Broker**.

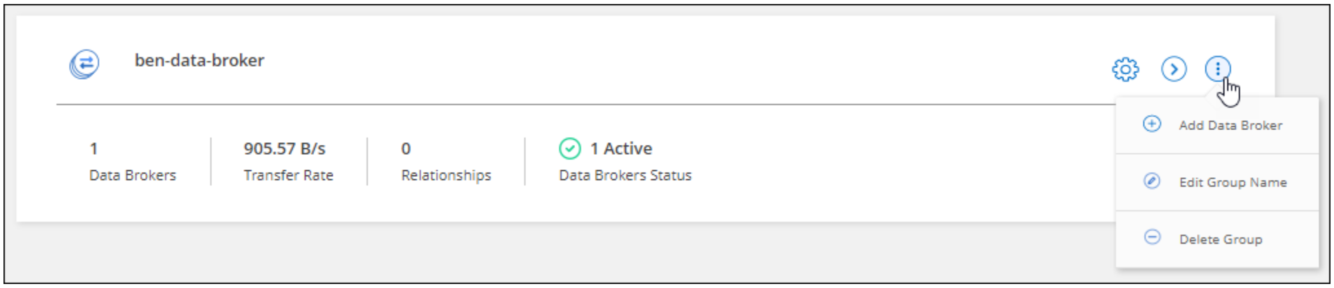

4. Follow the prompts to create the data broker in the group.

   For help, refer to the following pages:

   - Create a data broker in AWS
   - Create a data broker in Azure
   - Create a data broker in Google Cloud
   - Installing the data broker on a Linux host

Edit a group’s name

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

Steps

1. Log in to Copy and Sync.
2. Select **Sync > Manage Data Brokers**.
3. Select the action menu and select **Edit Group Name**.

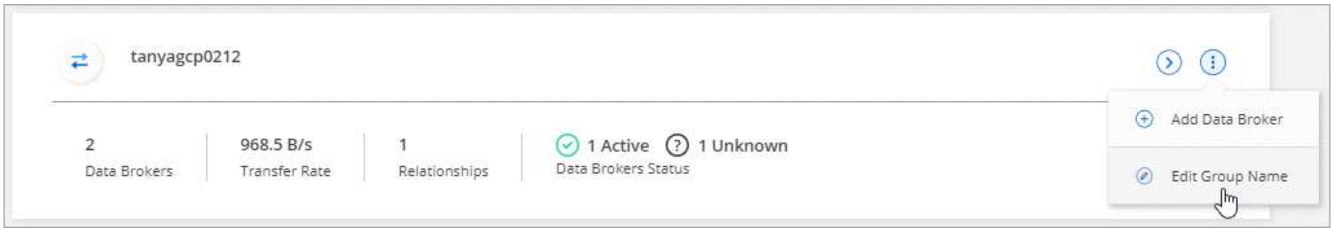

4. Enter a new name and select **Save**.

Result

Copy and Sync updates the name of the data broker group.

Set up a unified configuration

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. Log in to Copy and Sync.
2. Select **Manage Data Brokers**.
3. Select the Settings icon for a data broker group.

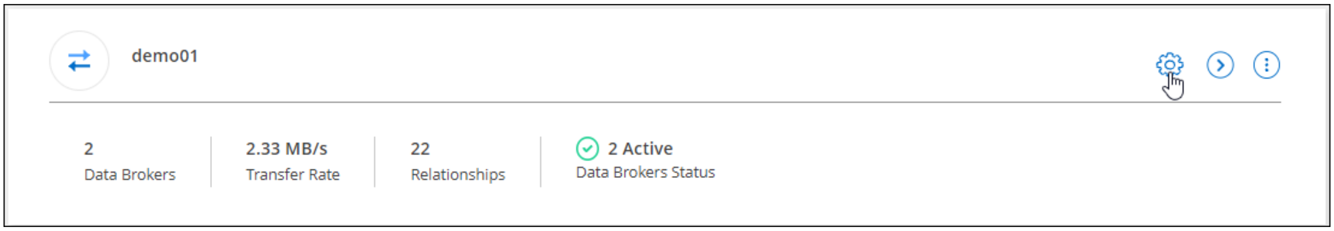

4. Change the settings as needed and then select **Unify Configuration**.

   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 Copy and Sync interface.
   - If a data broker isn’t running, it’s configuration won’t change because Copy and 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.

Move data brokers between groups

Move a data broker from one group to another group if you need to accelerate the performance of the target data broker group.

For example, if a data broker is no longer managing a sync relationship, you can easily move it to another group that is managing sync relationships.

Limitations

- If a data broker group is managing a sync relationship and there's only one data broker in the group, then you can't move that data broker to another group.
- You can't move a data broker to or from a group that manages encrypted sync relationships.
- You can't move a data broker that is currently being deployed.

Steps

1. Log in to Copy and Sync.
2. Select **Sync > Manage Data Brokers**.
3. Select › to expand the list of data brokers in a group.
4. Select the action menu for a data broker and select **Move Data Broker**.

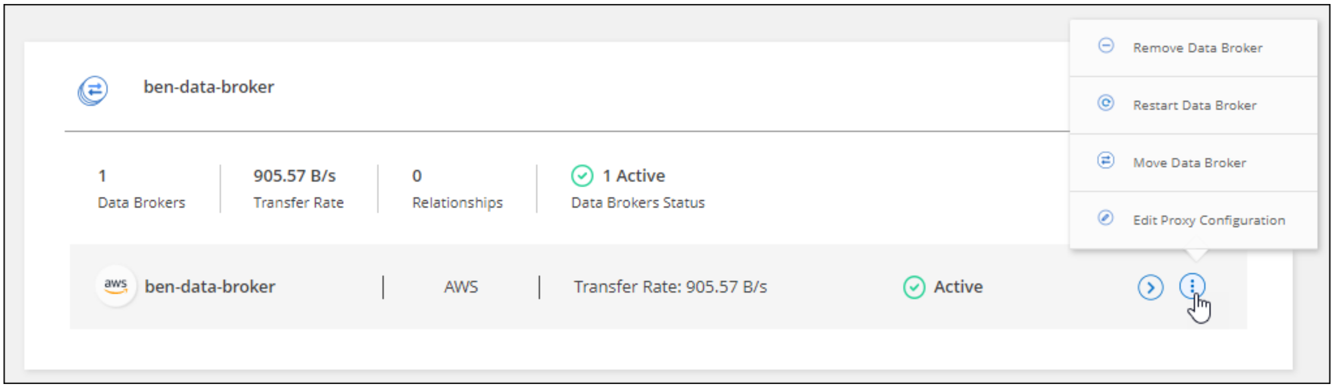

5. Create a new data broker group or select an existing data broker group.
6. Select **Move**.

Result

Copy and Sync moves the data broker to a new or existing data broker group. If there are no other data brokers in the previous group, then Copy and Sync deletes it.

Update proxy configuration

Update the proxy configuration for a data broker by adding details about a new proxy configuration or by editing the existing proxy configuration.

Steps

1. Log in to Copy and Sync.
2. Select **Sync > Manage Data Brokers**.
3. Select › to expand the list of data brokers in a group.
4. Select the action menu for a data broker and select **Edit Proxy Configuration**.
5. Specify details about the proxy: host name, port number, user name, and password.
6. Select **Update**.

Result

Copy and Sync updates the data broker to use the proxy configuration for internet access.

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.

Copy and 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. Log in to Copy and Sync.
2. Select **Sync > Manage Data Brokers**.
3. Select ❯ to expand the list of data brokers in a group.
4. Select ❯ to view details about a data broker.

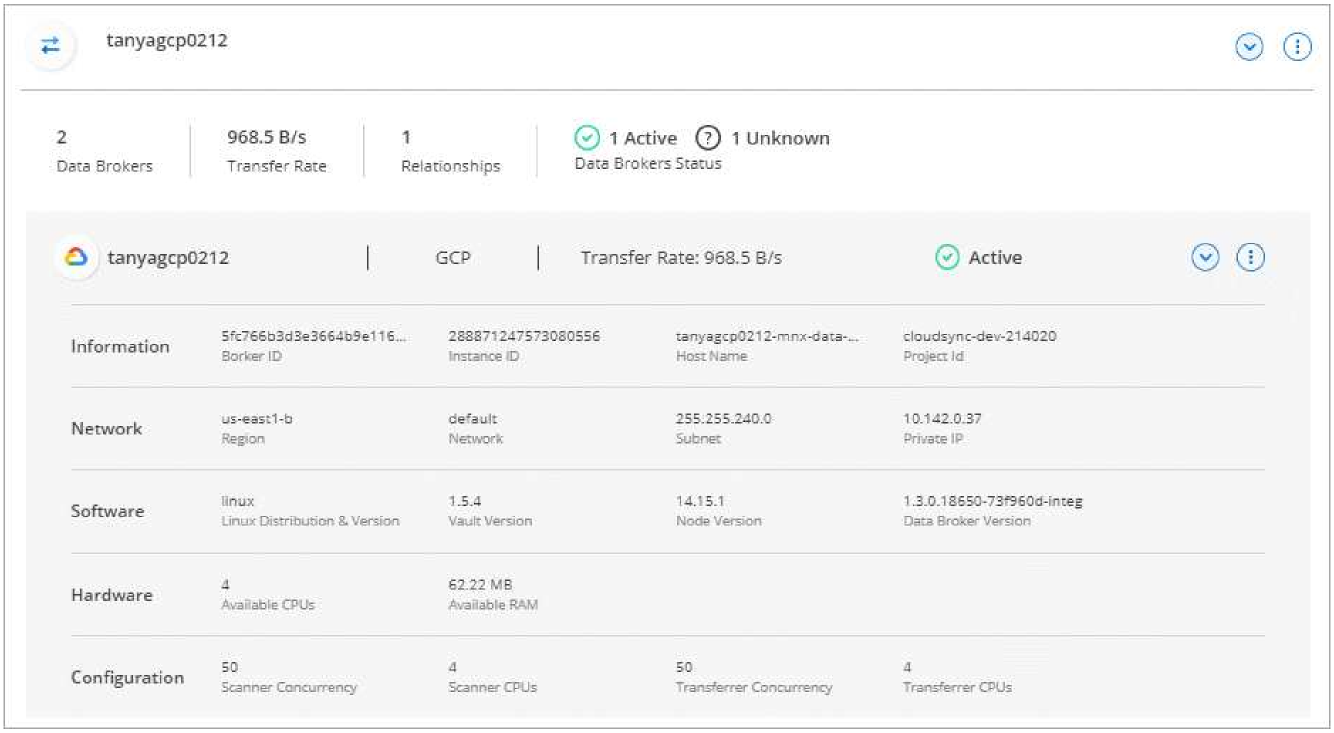

tanyagcp0212

2	968.5 B/s	1	✓ 1 Active ? 1 Unknown
Data Brokers	Transfer Rate	Relationships	Data Brokers Status

tanyagcp0212 | GCP | Transfer Rate: 968.5 B/s ✓ Active

Information	5fc766b3d3e3664b9e116... Borker 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

Address issues with a data broker

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

Steps

1. Log in to Copy and Sync.
2. Identify any data brokers that have a status of "Unknown" or "Failed."

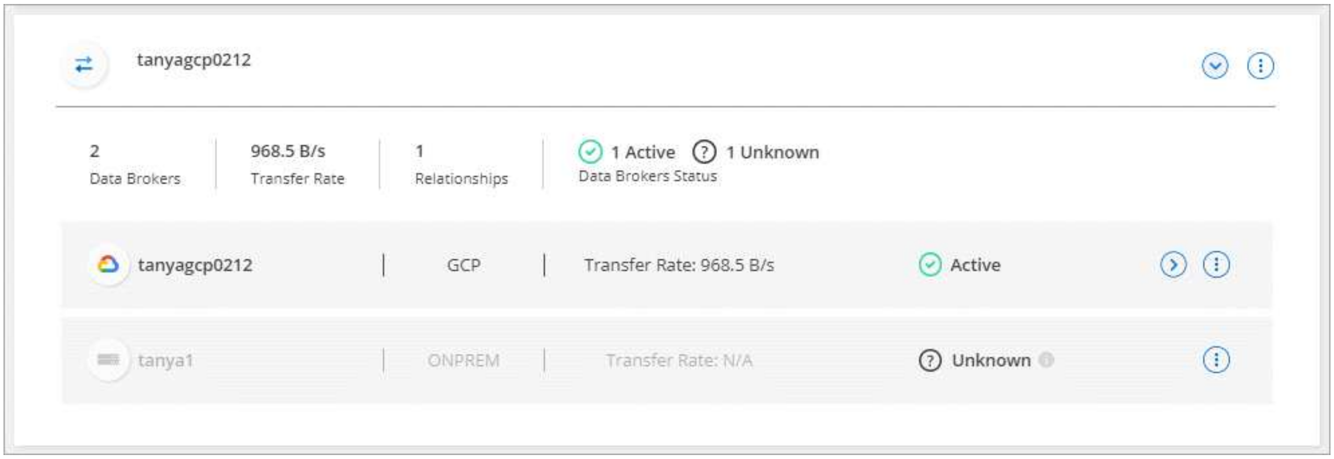

tanyagcp0212

2	968.5 B/s	1	✓ 1 Active ? 1 Unknown
Data Brokers	Transfer Rate	Relationships	Data Brokers Status

tanyagcp0212	GCP	Transfer Rate: 968.5 B/s	✓ Active
tanya1	ONPREM	Transfer Rate: N/A	? Unknown ⓘ

3. Hover over the ⓘ icon to see the failure reason.
4. 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.

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 Copy and Sync's records. You'll need to manually delete the data broker and any additional cloud resources yourself.

Things you should know

- Copy and 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. Log in to Copy and Sync.
2. Select **Sync > Manage Data Brokers**.
3. Select › to expand the list of data brokers in a group.
4. Select the action menu for a data broker and select **Remove Data Broker**.

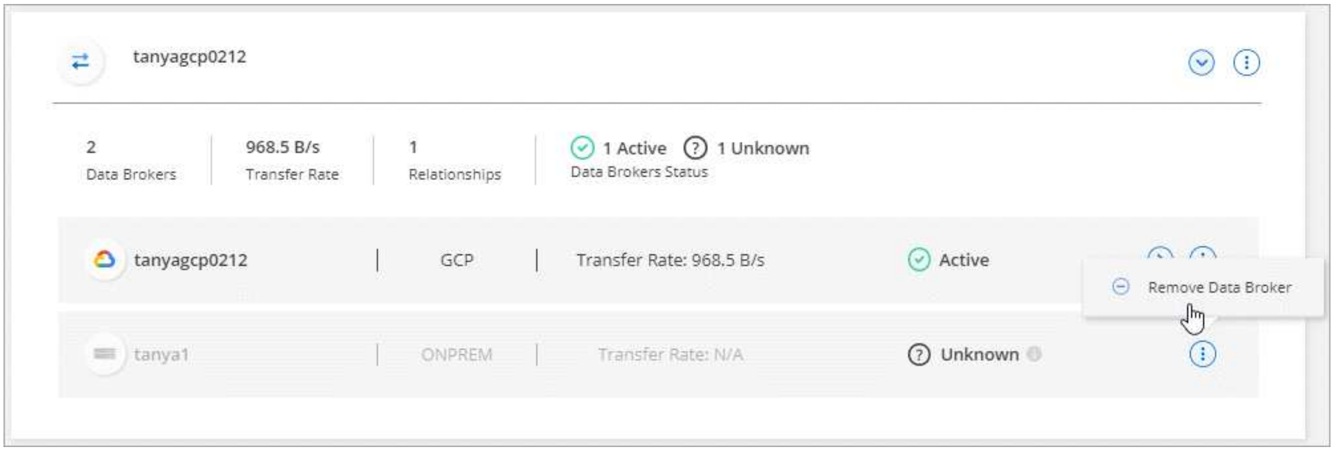

5. Select **Remove Data Broker**.

Result

Copy and Sync removes the data broker from the group.

Delete a data broker group

If a data broker group no longer manages any sync relationships, you can delete the group, which removes all of the data brokers from Copy and Sync.

Data brokers that Copy and Sync removes are only deleted from Copy and Sync's records. You'll need to manually delete the data broker instance from your cloud provider and any additional cloud resources.

Steps

1. Log in to Copy and Sync.
2. Select **Sync > Manage Data Brokers**.
3. Select the action menu and select **Delete Group**.

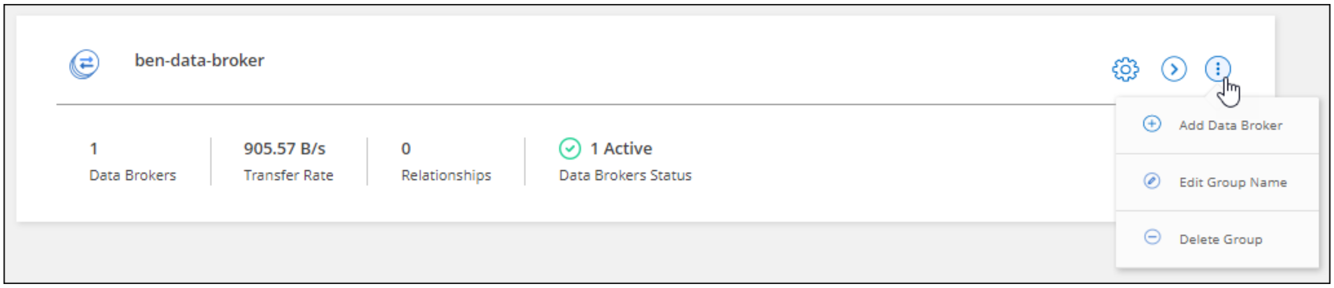

4. To confirm, enter the name of the group and select **Delete Group**.

Result

Copy and Sync removes the data brokers and deletes the group.

Create and view reports to tune your configuration in NetApp Copy and Sync

Create and view reports in NetApp Copy and Sync 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. It will include how many directories, files, and symbolic links there are, the distribution of file size, how deep and wide the directories are, modify time, and access time. This differs from sync statics, which are available from the dashboard after successfully creating and completing a sync.

Create reports

Each time that you create a report, Copy and Sync scans the path and then compiles the details into a report.

Steps

1. Log in to Copy and Sync.
2. Select **Sync > Reports**.

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

3. In the **Reports Actions** column, go to a specific path and select **Create**, or select the action menu and select **Create New**.
4. When the report is ready, select the action menu and select **View**.

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

   And here’s a sample report for object storage.

Download reports

You can download a report in PDF so that you can view it offline or share it.

Steps

1. Log in to Copy and Sync.
2. Select **Sync > Reports**.
3. In the **Reports Actions** column, select the action menu and select **View**.
4. In the top right of the report, select the action menu and select **Download pdf**.

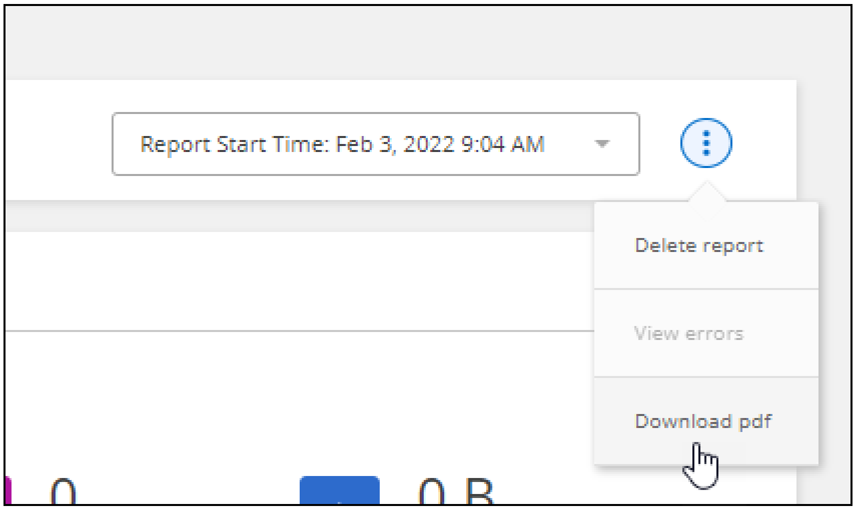

Report Start Time: Feb 3, 2022 9:04 AM

View report errors

The Paths table identifies whether any errors are present in the most recent report. An error identifies an issue that Copy and Sync faced when scanning the path.

For example, a report might contain permission denied errors. This type of error can affect Copy and Sync's ability to scan the entire set of files and directories.

After you view the list of errors, you can then address the issues and run the report again.

Steps

1. Log in to Copy and Sync.
2. Select **Sync > Reports**.
3. In the **Errors** column, identify whether any errors are present in a report.
4. If errors are present, select the arrow next to the number of errors.

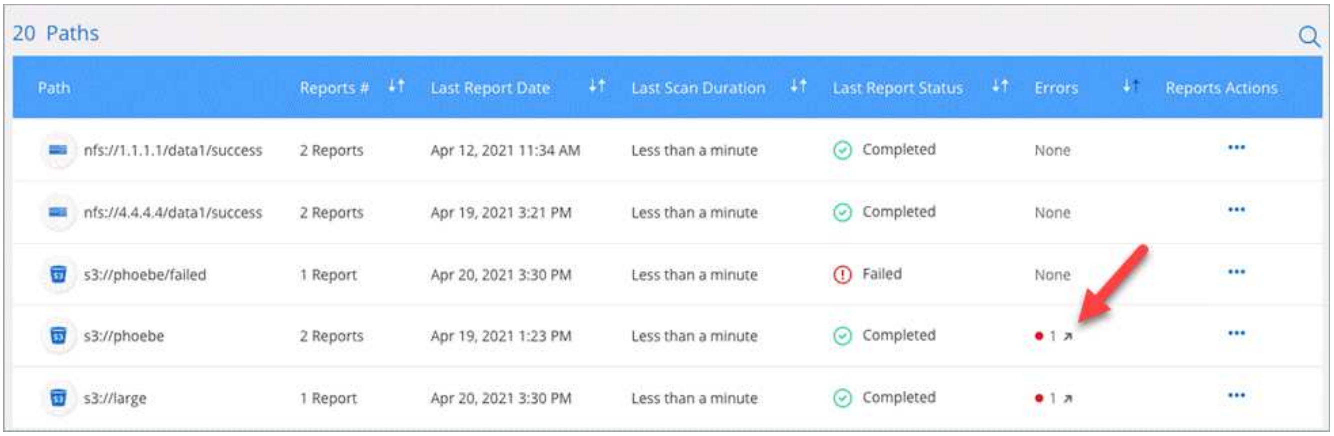

20 Paths

Path	Reports #	Last Report Date	Last Scan Duration	Last Report Status	Errors	Reports Actions
nfs://1.1.1.1/data1/success	2 Reports	Apr 12, 2021 11:34 AM	Less than a minute	✓ Completed	None	•••
nfs://4.4.4.4/data1/success	2 Reports	Apr 19, 2021 3:21 PM	Less than a minute	✓ Completed	None	•••
s3://phoebe/failed	1 Report	Apr 20, 2021 3:30 PM	Less than a minute	! Failed	None	•••
s3://phoebe	2 Reports	Apr 19, 2021 1:23 PM	Less than a minute	✓ Completed	● 1 ↗	•••
s3://large	1 Report	Apr 20, 2021 3:30 PM	Less than a minute	✓ Completed	● 1 ↗	•••

5. Use the information in the error to correct the issue.

   After you resolve the issue, the error shouldn't appear the next time that you run the report.

Delete reports

You might delete a report of it contained an error that you fixed, or if the report is related to a sync relationship that you removed.

Steps

1. Select **Sync > Reports**.
2. In the **Reports Actions** column, select the action menu for a path and select **Delete last report** or **Delete all reports**.
3. Confirm that you want to delete the report or reports.

Uninstall the data broker for NetApp Copy and Sync

If needed, run an uninstall script to remove the data broker and the packages and directories that were created for NetApp Copy and Sync 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.

NetApp Copy and Sync APIs

The NetApp Copy and Sync capabilities that are available through the web UI are also available through the RESTful API.

Get started

To get started with the Copy and Sync API, you need to obtain a user token and your NetApp Console 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 Console.

```
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>"
}
```

ⓘ If you are using a personal email account with no client ID, you can use the default client ID "QC3AgHk6qdbmC7Yyr82ApBwaaJLwRrNO."

2. Obtain your NetApp Console 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's next?

The user token from NetApp Console 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.

Use 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"
    }
]
```

API reference

Documentation for each Copy and Sync API is available from https://api.cloudsync.netapp.com/docs.

Concepts

Licensing overview for NetApp Copy and Sync

There are two ways to pay for NetApp Copy and 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.

Licenses should be managed through NetApp Copy and Sync or the applicable website and **not** through the NetApp Console.

Marketplace subscription

Subscribing to Copy and Sync 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.

> ⓘ Copy and Sync supports Marketplace subscriptions from **AWS** and **Azure** only. Google Cloud Marketplace subscriptions are not supported for Copy and Sync.

Hourly subscription

With an hourly pay-as-you-go subscription, you’re charged hourly based on the number of sync relationships that you create.

- View pricing in Azure
- View pay-as-you-go pricing in AWS

Annual subscription

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 AWS, you pay for the additional relationships by the hour.

View annual pricing in AWS

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 NetApp Copy and Sync.

License terms

Customers who purchase a Bring Your Own License (BYOL) to Copy and Sync 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.

Data privacy in NetApp Copy and Sync

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

Depending on the configuration that you choose, Copy and 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.

NetApp Copy and 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 NetApp Copy and Sync.

How does NetApp Copy and Sync work?

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

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

After the initial copy, Copy and Sync 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 Copy and Sync. You’re not subject to NetApp charges for Copy and Sync relationships you create for 14 days. However, all resource charges for any data brokers that you deploy still applies.

How much does Copy and Sync cost?

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

Service charges

For pay-as-you-go pricing, Copy and 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

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

Learn more about licenses.

ⓘ Copy and Sync relationships are free for Azure NetApp Files.

Resource charges

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

How is Copy and Sync billed and how do I manage my subscription?

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. In each case, your subscription will be managed through your provider marketplace and **not** via the Copy and Sync user interface.

Can I use Copy and Sync outside the cloud?

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

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

- A data broker group needs an internet connection to communicate with Copy and Sync.
- If you don’t purchase a license directly from NetApp, you will need an AWS or Azure account for PAYGO Copy and Sync billing.

How do I access Copy and Sync?

Copy and Sync is available from the the NetApp Console. From the Console left navigation, select **Mobility** > **Copy and Sync**.

What’s a data broker group?

Each data broker belongs to a data broker group. Grouping data brokers together helps improve the performance of sync relationships.

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 Copy and Sync support?

Copy and Sync supports many different types of sync relationships. View the entire list.

What versions of NFS and SMB does Copy and Sync support?

Copy and 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 Deep Archive
- Glacier Flexible Retrieval
- Glacier Instant Retrieval

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

Do you support Google Cloud storage tiers?

Yes, you can choose a specific storage class when a Google Cloud Storage bucket is the target:

- Standard
- Nearline
- Coldline
- Archive

Networking

The following questions relate to networking requirements for Copy and Sync.

What are the networking requirements for Copy and Sync?

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

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

For more details, review networking requirements.

Can I use a proxy server with the data broker?

Yes.

Copy and 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 group retry when a file fails to sync? Or does it timeout?

A data broker group doesn't timeout when a single file fails to transfer. Instead, the data broker group 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 that we can help you configure the data broker group to handle the payload. We might need to add additional memory to the data broker group.

Note that there's no limit to the total number of files in the mount point. The extra memory is required for large directories with 600,000 files or more, regardless of their level in the hierarchy (top directory or subdirectory).

Security

The following questions related to security.

Is Copy and Sync secure?

Yes. All Copy and Sync networking connectivity is done using Amazon Simple Queue Service (SQS).

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

If you're using Copy and 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 Copy and Sync data broker group.

Is data encrypted by Copy and Sync?

- Copy and Sync supports data-in-flight encryption between source and target NFS servers. Learn more.
- For SMB, Copy and Sync supports SMB 3.0 and 3.11 data that you've encrypted on the server side. Copy and Sync copies the encrypted data from the source to the target where the data remains encrypted.

  Copy and Sync cannot encrypt SMB data itself.
- 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.
- When a Google Storage bucket is the target in a sync relationship, you can choose whether to use the default, Google-managed encryption key or your own KMS key.

Permissions

The following questions relate to data permissions.

Are SMB data permissions synced to the target location?

You can set up Copy and Sync to preserve access control lists (ACLs) between a source SMB share and a target SMB share, and from a source SMB share to object storage (except for ONTAP S3).

> ⓘ Copy and Sync doesn't support copying ACLs from object storage to SMB shares.

Learn how to copy ACLs between SMB shares.

> ⓘ Copy Sync copies SMB ACLs (permissions), but it does not copy file or folder ownership. The ownership attribute is not included in the SMB ACL copy operation. If you need to preserve ownership when copying data between SMB shares, use `robocopy` to manually copy the security information. For example, the `/copyall` flag copies ACLs, owner, and audit data.

Are NFS data permissions synced to the target location?

Copy and Sync automatically copies NFS permissions between NFS servers as follows:

- NFS version 3: Copy and Sync copies the permissions and the user group owner.
- NFS version 4: Copy and Sync copies the ACLs.

Object storage metadata

What kinds of sync relationships preserve object storage metadata?

Copy and Sync copies object storage metadata from the source to the target for the following types of sync relationships:

- Amazon S3 → Amazon S3 [1]
- Amazon S3 → StorageGRID
- StorageGRID → Amazon S3
- StorageGRID → StorageGRID
- StorageGRID → Google Cloud Storage
- Google Cloud Storage → StorageGRID [1]
- Google Cloud Storage → IBM Cloud Object Storage [1]
- Google Cloud Storage → Amazon S3 [1]
- Amazon S3 → Google Cloud Storage
- IBM Cloud Object Storage → Google Cloud Storage
- StorageGRID → IBM Cloud Object Storage
- IBM Cloud Object Storage → StorageGRID
- IBM Cloud Object Storage → IBM Cloud Object Storage

[1] For these sync relationships, you need to enable the Copy for Objects setting when you create the sync relationship.

What kinds of metadata are replicated during syncs where NFS or SMB are the source?

Metadata such as user ID, modification time, access time, and GID are replicated by default. Users may opt into replicating ACL from CIFs by marking it as required when creating a sync relationship.

Performance

The following questions relate to Copy and Sync performance.

What does the progress indicator for a sync relationship represent?

The sync relationship shows the throughput of the data broker group’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?

If you have very large files (multiple TiBs 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.

How many data brokers are required in a group?

When you create a new relationship, you start with a single data broker in a group (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 to the group. 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 group can reach 100 MB/s, while disk throughput on the target might only allow 64 MB/s. As a result, the data broker group keeps trying to copy the data, but the target can’t meet the performance of the data broker group.

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 brokers to a group 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 Copy and 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 Copy and 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, Copy and Sync compares the source to the target, identifies that the item is missing, and Copy and Sync copies it again from the source to the target.

Troubleshooting

NetApp Knowledgebase: Copy and 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.
- Copy and 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 group to increase transfer speed and add high availability. There is service resiliency if one data broker fails.

Knowledge and support

Register for support

Support registration is required to receive technical support specific to the NetApp Console and its storage solutions and data services. Support registration is also required to enable key workflows for Cloud Volumes ONTAP systems.

Registering for support does not enable NetApp support for a cloud provider file service. For technical support related to a cloud provider file service, its infrastructure, or any solution using the service, refer to "Getting help" in the documentation for that product.

- Amazon FSx for ONTAP
- Azure NetApp Files
- Google Cloud NetApp Volumes

Support registration overview

There are two forms of registration to activate support entitlement:

- Registering your NetApp Console account serial number (your 20 digit 960xxxxxxxxx serial number located on the Support Resources page in the Console).

  This serves as your single support subscription ID for any service within the Console. Each Console account must be registered.

- Registering the Cloud Volumes ONTAP serial numbers associated with a subscription in your cloud provider's marketplace (these are 20 digit 909201xxxxxxxx serial numbers).

  These serial numbers are commonly referred to as *PAYGO serial numbers* and get generated by the NetApp Console at the time of Cloud Volumes ONTAP deployment.

Registering both types of serial numbers enables capabilities like opening support tickets and automatic case generation. Registration is completed by adding NetApp Support Site (NSS) accounts to the Console as described below.

Register NetApp Console for NetApp support

To register for support and activate support entitlement, one user in your NetApp Console account must associate a NetApp Support Site account with their Console login. How you register for NetApp support depends on whether you already have a NetApp Support Site (NSS) account.

Existing customer with an NSS account

If you're a NetApp customer with an NSS account, you simply need to register for support through the Console.

Steps

1. Select **Administration** > **Credentials**.
2. Select **User Credentials**.

3. Select **Add NSS credentials** and follow the NetApp Support Site (NSS) authentication prompt.
4. To confirm that the registration process was successful, select the Help icon, and select **Support**.

   The **Resources** page should show that your Console account is registered for support.

   Note that other Console users will not see this same support registration status if they have not associated a NetApp Support Site account with their login. However, that doesn't mean that your account is not registered for support. As long as one user in the organization has followed these steps, then your account has been registered.

Existing customer but no NSS account

If you're an existing NetApp customer with existing licenses and serial numbers but *no* NSS account, you need to create an NSS account and associate it with your Console login.

Steps

1. Create a NetApp Support Site account by completing the NetApp Support Site User Registration form
   a. Be sure to select the appropriate User Level, which is typically **NetApp Customer/End User**.
   b. Be sure to copy the Console account serial number (960xxxx) used above for the serial number field. This will speed up the account processing.
2. Associate your new NSS account with your Console login by completing the steps under Existing customer with an NSS account.

Brand new to NetApp

If you are brand new to NetApp and you don't have an NSS account, follow each step below.

Steps

1. In the upper right of the Console, select the Help icon, and select **Support**.
2. Locate your account ID serial number from the Support Registration page.

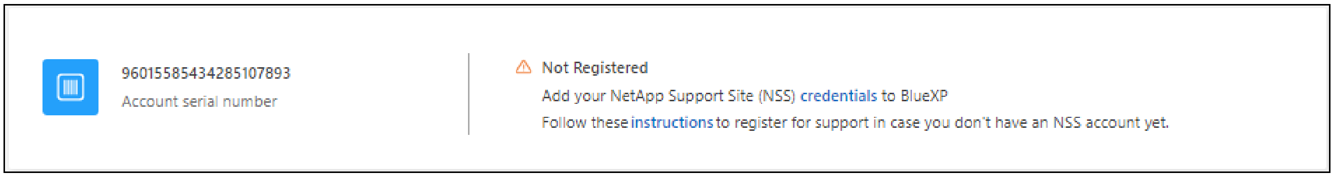

960155854342851078893

Account serial number

⚠ Not Registered

Add your NetApp Support Site (NSS) credentials to BlueXP

Follow these instructions to register for support in case you don't have an NSS account yet.

3. Navigate to NetApp's support registration site and select **I am not a registered NetApp Customer**.
4. Fill out the mandatory fields (those with red asterisks).
5. In the **Product Line** field, select **Cloud Manager** and then select your applicable billing provider.
6. Copy your account serial number from step 2 above, complete the security check, and then confirm that you read NetApp's Global Data Privacy Policy.

   An email is immediately sent to the mailbox provided to finalize this secure transaction. Be sure to check your spam folders if the validation email doesn't arrive in few minutes.

7. Confirm the action from within the email.

   Confirming submits your request to NetApp and recommends that you create a NetApp Support Site account.

8. Create a NetApp Support Site account by completing the NetApp Support Site User Registration form
   a. Be sure to select the appropriate User Level, which is typically **NetApp Customer/End User**.
   b. Be sure to copy the account serial number (960xxxx) used above for the serial number field. This will speed up processing.

After you finish

NetApp should reach out to you during this process. This is a one-time onboarding exercise for new users.

Once you have your NetApp Support Site account, associate the account with your Console login by completing the steps under Existing customer with an NSS account.

Associate NSS credentials for Cloud Volumes ONTAP support

Associating NetApp Support Site credentials with your Console account is required to enable the following key workflows for Cloud Volumes ONTAP:

- Registering pay-as-you-go Cloud Volumes ONTAP systems for support

  Providing your NSS account is required to activate support for your system and to gain access to NetApp technical support resources.

- Deploying Cloud Volumes ONTAP when you bring your own license (BYOL)

  Providing your NSS account is required so that the Console can upload your license key and to enable the subscription for the term that you purchased. This includes automatic updates for term renewals.

- Upgrading Cloud Volumes ONTAP software to the latest release

Associating NSS credentials with your NetApp Console account is different than the NSS account that is associated with a Console user login.

These NSS credentials are associated with your specific Console account ID. Users who belong to the Console organization can access these credentials from **Support > NSS Management**.

- If you have a customer-level account, you can add one or more NSS accounts.
- If you have a partner or reseller account, you can add one or more NSS accounts, but they can’t be added alongside customer-level accounts.

Steps

1. In the upper right of the Console, select the Help icon, and select **Support**.

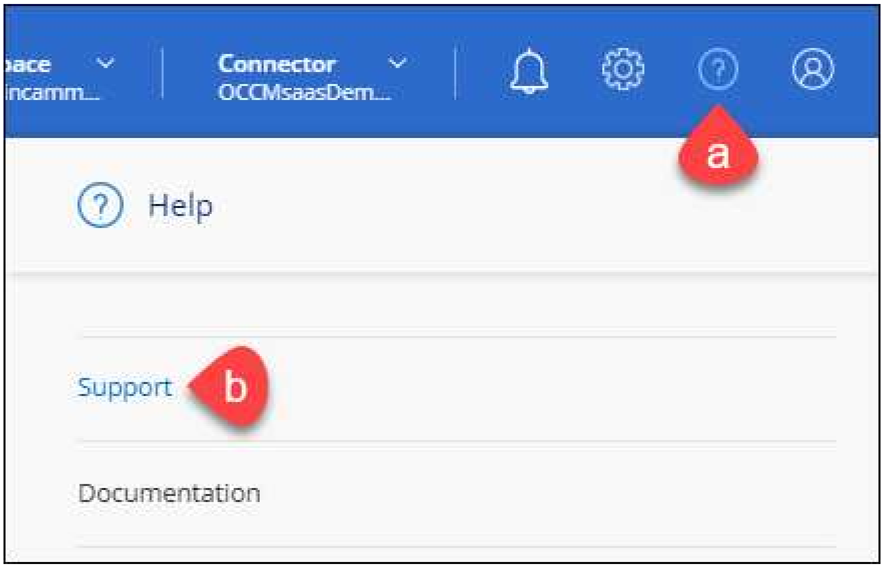

2. Select **NSS Management > Add NSS Account**.
3. When you’re prompted, select **Continue** to be redirected to a Microsoft login page.

   NetApp uses Microsoft Entra ID as the identity provider for authentication services specific to support and licensing.

4. At the login page, provide your NetApp Support Site registered email address and password to perform the authentication process.

   These actions enable the Console to use your NSS account for things like license downloads, software upgrade verification, and future support registrations.

   Note the following:

   - The NSS account must be a customer-level account (not a guest or temp account). You can have multiple customer-level NSS accounts.
   - There can be only one NSS account if that account is a partner-level account. If you try to add customer-level NSS accounts and a partner-level account exists, you’ll get the following error message:

     "The NSS customer type is not allowed for this account as there are already NSS Users of different type."

     The same is true if you have pre-existing customer-level NSS accounts and try to add a partner-level account.

   - Upon successful login, NetApp will store the NSS user name.

     This is a system-generated ID that maps to your email. On the **NSS Management** page, you can display your email from the ••• menu.

   - If you ever need to refresh your login credential tokens, there is also an **Update Credentials** option in the ••• menu.

     Using this option prompts you to log in again. Note that the token for these accounts expire after 90 days. A notification will be posted to alert you of this.

Get help

NetApp provides support for BlueXP and its cloud services in a variety of ways. Extensive free self-support options are available 24/7, such as knowledgebase (KB) articles and a community forum. Your support registration includes remote technical support via web ticketing.

Get support for a cloud provider file service

For technical support related to a cloud provider file service, its infrastructure, or any solution using the service, refer to "Getting help" in the BlueXP documentation for that product.

- Amazon FSx for ONTAP
- Azure NetApp Files
- Google Cloud NetApp Volumes

To receive technical support specific to BlueXP and its storage solutions and services, use the support options described below.

Use self-support options

These options are available for free, 24 hours a day, 7 days a week:

- Documentation

  The BlueXP documentation that you're currently viewing.

- Knowledge base

  Search through the BlueXP knowledge base to find helpful articles to troubleshoot issues.

- Communities

  Join the BlueXP community to follow ongoing discussions or create new ones.

Create a case with NetApp support

In addition to the self-support options above, you can work with a NetApp Support specialist to resolve any issues after you activate support.

Before you get started

- To use the **Create a Case** capability, you must first associate your NetApp Support Site credentials with your BlueXP login. Learn how to manage credentials associated with your BlueXP login.
- If you're opening a case for an ONTAP system that has a serial number, then your NSS account must be associated with the serial number for that system.

Steps

1. In BlueXP, select **Help > Support**.
2. On the **Resources** page, choose one of the available options under Technical Support:

a. Select **Call Us** if you'd like to speak with someone on the phone. You'll be directed to a page on netapp.com that lists the phone numbers that you can call.

b. Select **Create a Case** to open a ticket with a NetApp Support specialist:

   - **Service**: Select the service that the issue is associated with. For example, BlueXP when specific to a technical support issue with workflows or functionality within the service.
   - **Working Environment**: If applicable to storage, select **Cloud Volumes ONTAP** or **On-Prem** and then the associated working environment.

     The list of working environments are within scope of the BlueXP organization (or account), project (or workspace), and Connector you have selected in the top banner of the service.

   - **Case Priority**: Choose the priority for the case, which can be Low, Medium, High, or Critical.

     To learn more details about these priorities, hover your mouse over the information icon next to the field name.

   - **Issue Description**: Provide a detailed description of your problem, including any applicable error messages or troubleshooting steps that you performed.
   - **Additional Email Addresses**: Enter additional email addresses if you'd like to make someone else aware of this issue.
   - **Attachment (Optional)**: Upload up to five attachments, one at a time.

     Attachments are limited to 25 MB per file. The following file extensions are supported: txt, log, pdf, jpg/jpeg, rtf, doc/docx, xls/xlsx, and csv.

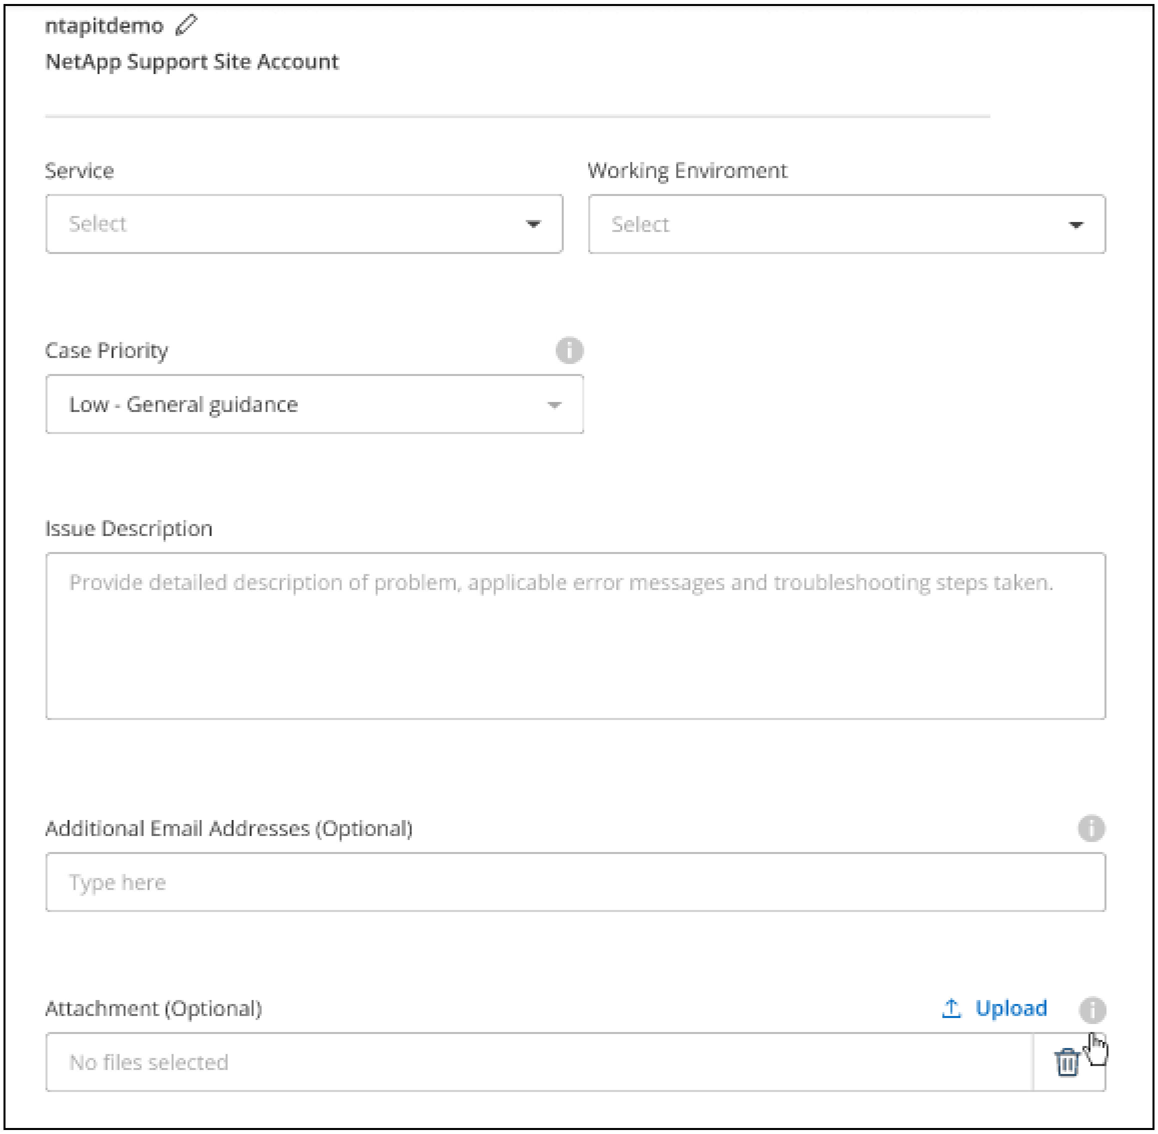

ntapitdemo

NetApp Support Site Account

Service: Select

Working Enviroment: Select

Case Priority: Low - General guidance

Issue Description: Provide detailed description of problem, applicable error messages and troubleshooting steps taken.

Additional Email Addresses (Optional): Type here

Attachment (Optional): No files selected

After you finish

A pop-up will appear with your support case number. A NetApp Support specialist will review your case and get back to you soon.

For a history of your support cases, you can select **Settings > Timeline** and look for actions named "create support case." A button to the far right lets you expand the action to see details.

It’s possible that you might encounter the following error message when trying to create a case:

"You are not authorized to Create a Case against the selected service"

This error could mean that the NSS account and the company of record it’s associated with is not the same company of record for the BlueXP account serial number (ie. 960xxxx) or the working environment serial number. You can seek assistance using one of the following options:

- Use the in-product chat
- Submit a non-technical case at https://mysupport.netapp.com/site/help

Manage your support cases (Preview)

You can view and manage active and resolved support cases directly from BlueXP. You can manage the cases associated with your NSS account and with your company.

Case management is available as a Preview. We plan to refine this experience and add enhancements in upcoming releases. Please send us feedback by using the in-product chat.

Note the following:

- The case management dashboard at the top of the page offers two views:
  - The view on the left shows the total cases opened in the past 3 months by the user NSS account you provided.
  - The view on the right shows the total cases opened in the past 3 months at your company level based on your user NSS account.

  The results in the table reflect the cases related to the view that you selected.

- You can add or remove columns of interest and you can filter the contents of columns like Priority and Status. Other columns provide just sorting capabilities.

  View the steps below for more details.

- At a per-case level, we offer the ability to update case notes or close a case that is not already in Closed or Pending Closed status.

Steps

1. In BlueXP, select **Help > Support**.
2. Select **Case Management** and if you’re prompted, add your NSS account to BlueXP.

   The **Case management** page shows open cases related to the NSS account that is associated with your BlueXP user account. This is the same NSS account that appears at the top of the **NSS management** page.

3. Optionally modify the information that displays in the table:
   - Under **Organization’s cases**, select **View** to view all cases associated with your company.
   - Modify the date range by choosing an exact date range or by choosing a different time frame.

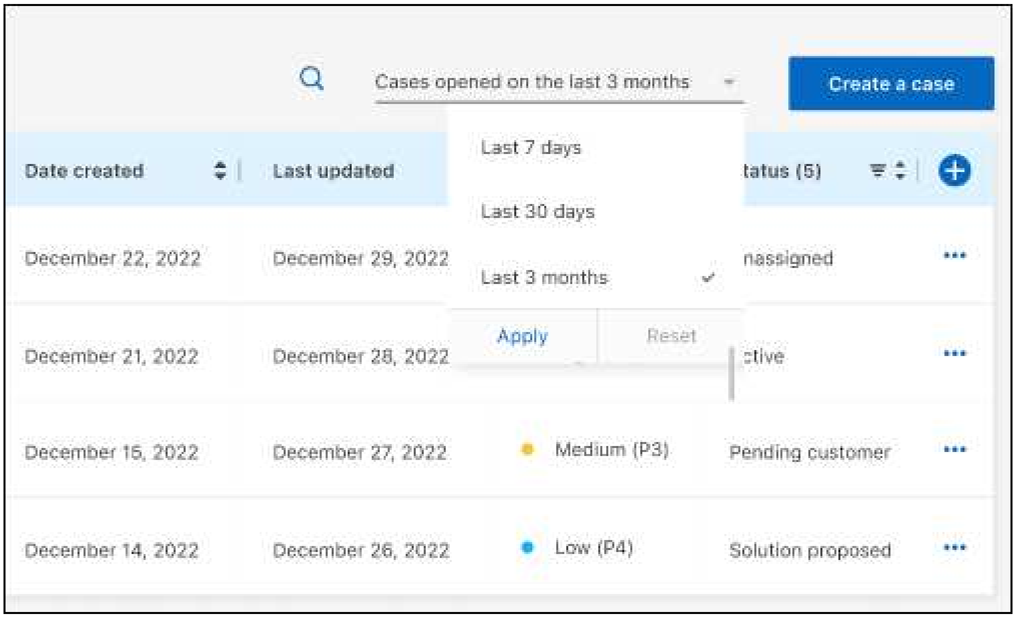

Date created	Last updated		tatus (5)
December 22, 2022	December 29, 2022		nassigned
December 21, 2022	December 28, 2022		ctive
December 15, 2022	December 27, 2022	● Medium (P3)	Pending customer
December 14, 2022	December 26, 2022	● Low (P4)	Solution proposed

- Filter the contents of the columns.

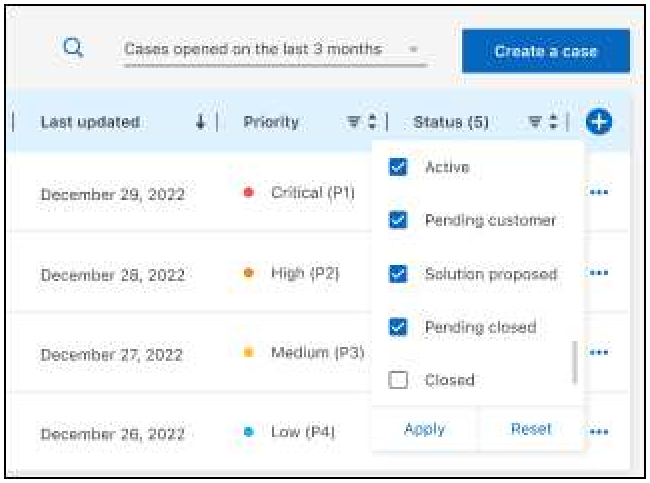

Last updated	Priority	Status (5)
December 29, 2022	● Critical (P1)	
December 28, 2022	● High (P2)	
December 27, 2022	● Medium (P3)	
December 26, 2022	● Low (P4)	

- Change the columns that appear in the table by selecting ⊕ and then choosing the columns that you'd like to display.

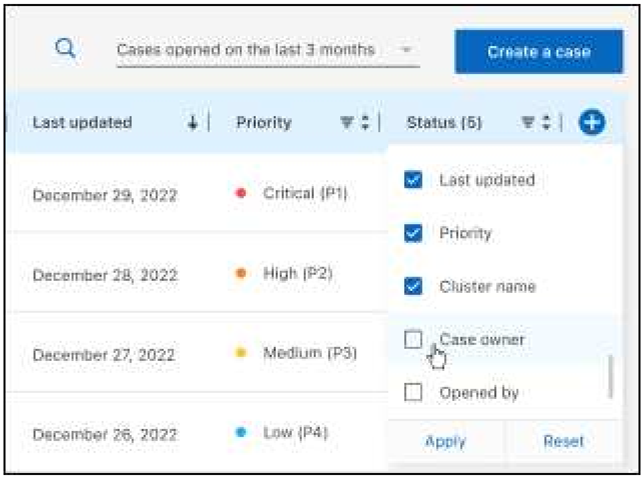

Last updated	Priority	Status (5)
December 29, 2022	● Critical (P1)	
December 28, 2022	● High (P2)	
December 27, 2022	● Medium (P3)	
December 26, 2022	● Low (P4)	

4. Manage an existing case by selecting ••• and selecting one of the available options:

   - **View case**: View full details about a specific case.
   - **Update case notes**: Provide additional details about your problem or select **Upload files** to attach up to a maximum of five files.

     Attachments are limited to 25 MB per file. The following file extensions are supported: txt, log, pdf, jpg/jpeg, rtf, doc/docx, xls/xlsx, and csv.

   - **Close case**: Provide details about why you're closing the case and select **Close case**.

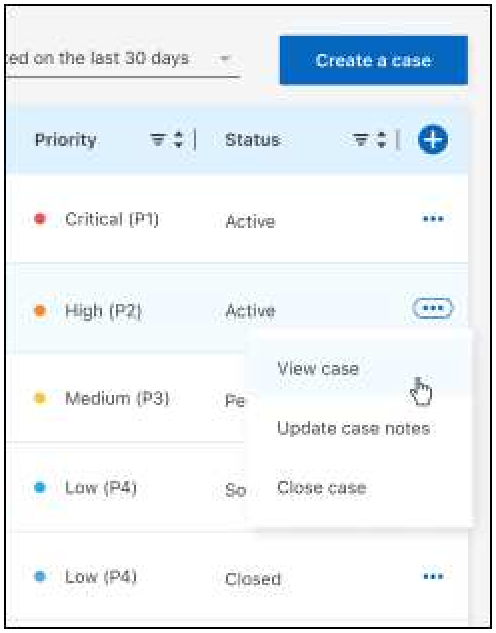

Priority	Status
● Critical (P1)	Active
● High (P2)	Active
● Medium (P3)	Pe
● Low (P4)	So
● Low (P4)	Closed

Legal notices

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

Copyright

https://www.netapp.com/company/legal/copyright/

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.

https://www.netapp.com/company/legal/trademarks/

Patents

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

https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf

Privacy policy

https://www.netapp.com/company/legal/privacy-policy/

Open source

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

Notice for NetApp Copy and Sync

Copyright information

Trademark information