



## Release notes

### Cloud Sync

NetApp  
May 08, 2024

# Table of Contents

- Release notes ..... 1
- What's new in Cloud Sync ..... 1
- Known limitations ..... 20

# Release notes

## What's new in Cloud Sync

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

### 12 Mar 2021

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

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

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

### 9 Mar 2021

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

[View supported sync relationships.](#)

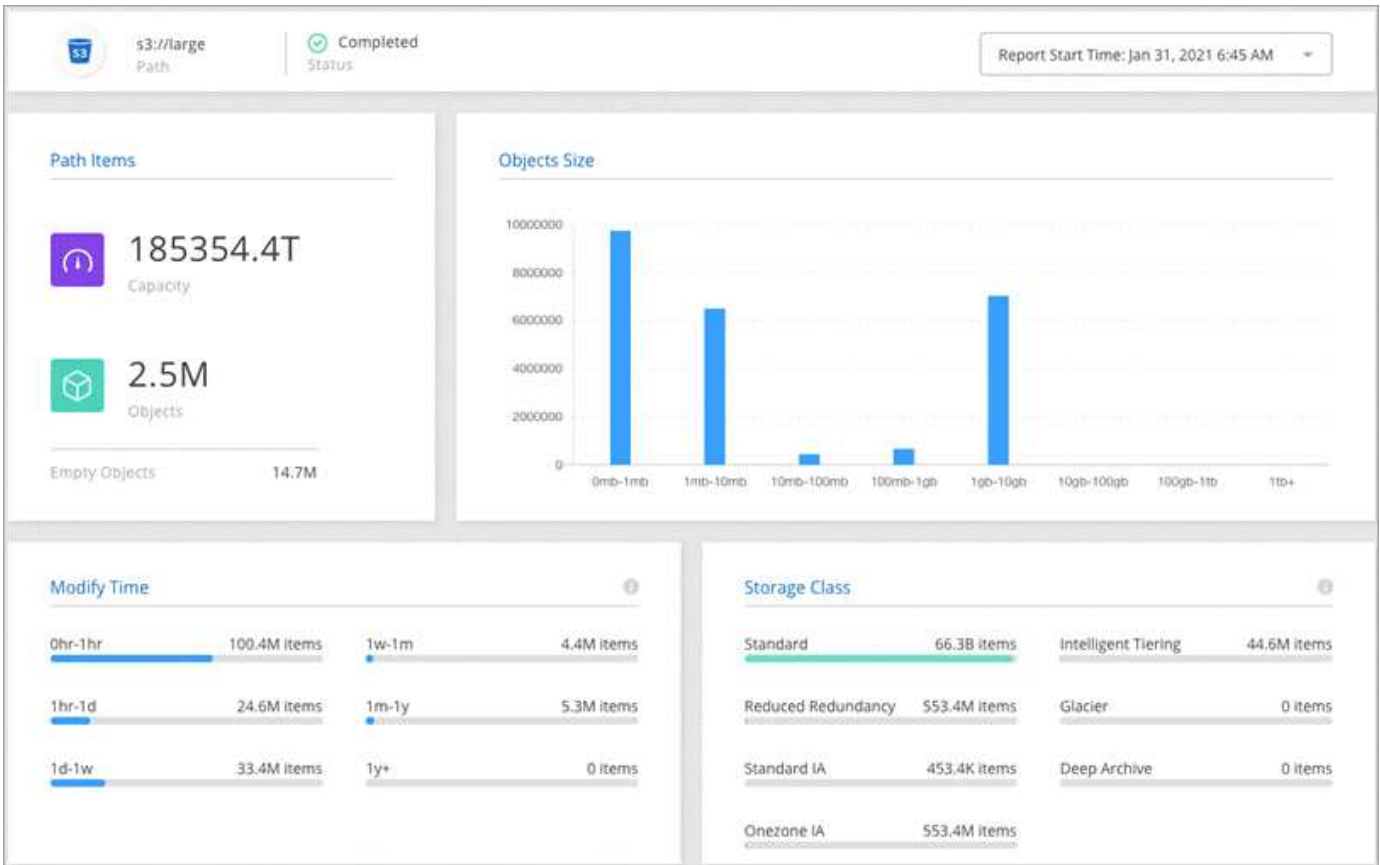
- Cloud Sync now enables you to unify a data broker group's configuration directly from the user interface.

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

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

### 10 Feb 2021

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



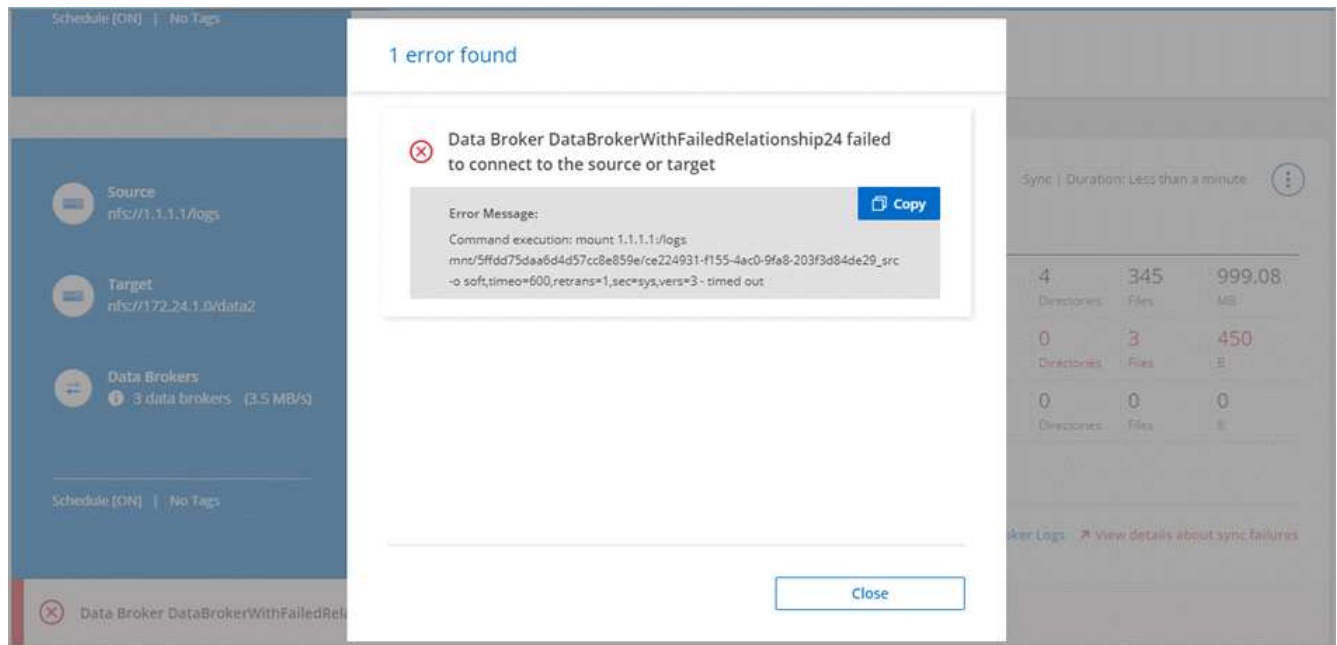
### 13 Jan 2021

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

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

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

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



## 7 Dec 2020

- You can now manage data broker groups.

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

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

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

[View supported sync relationships.](#)

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

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

## 1 Nov 2020

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

**Proxy Configuration (optional)**

Host ⓘ

Port ⓘ

Define credentials for this proxy ▲

User Name ⓘ

Password ⓘ

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

## 6 Sept 2020

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

## 2 Aug 2020

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

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

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

## 6 July 2020

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

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

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

The screenshot displays the Sync Relationships dashboard. On the left, a blue sidebar contains three items: 'Source' with a file icon and path 'nfs://.../mnt/kingk...', 'Target' with a folder icon and path 'nfs-anf://172.16.14.4/vol3/test1', and 'On-Prem Data Broker' with a bidirectional arrow icon. Below these is a 'Schedule [ON]' button. The main area shows a green checkmark icon and the text 'Synced Successfully'. Below this is a 'Scan' table:

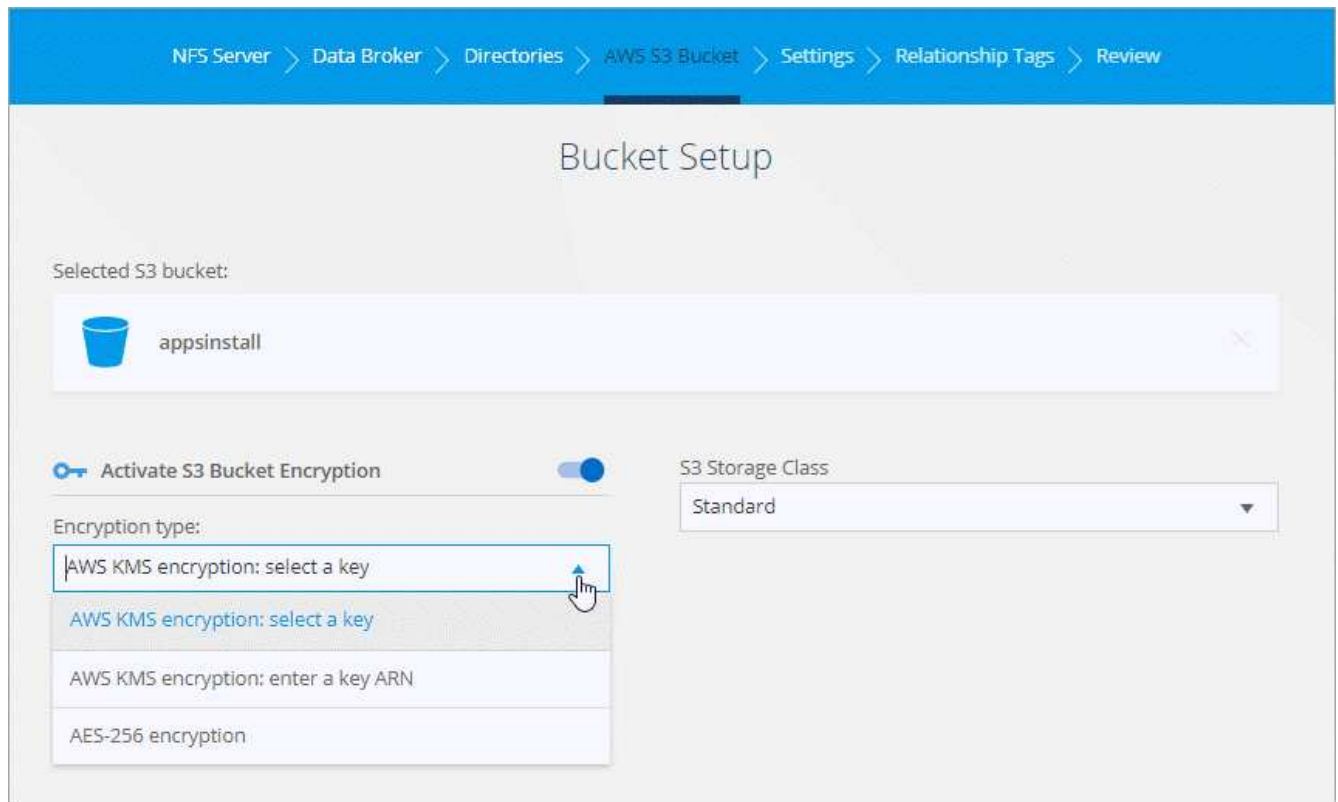
Scan			
Succeeded	683	2.7K	457.42
	Directories	Files	MB
Failed	0		
	Directories		

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

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

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

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



- SMB version 3.11 is now supported.



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



**Source**  
nfs://[redacted]/mnt/kingkong...

**Target**  
s3://[redacted]/testush

Selected Storage Class for target bucket: "Standard"

Schedule [ON] | No Tags

**Synced Successfully**

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

## 4 June 2020

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

NFS Server > Data Broker > Directories > Target Cloud Volumes Service > Target Directories > Settings > Review

Select Directory Refresh Directories

1 🔍

All Directories

- ✓ /MyExportName

## 5 May 2020

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

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

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

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

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

- Two new settings are available for sync relationships:

- Delete files on source

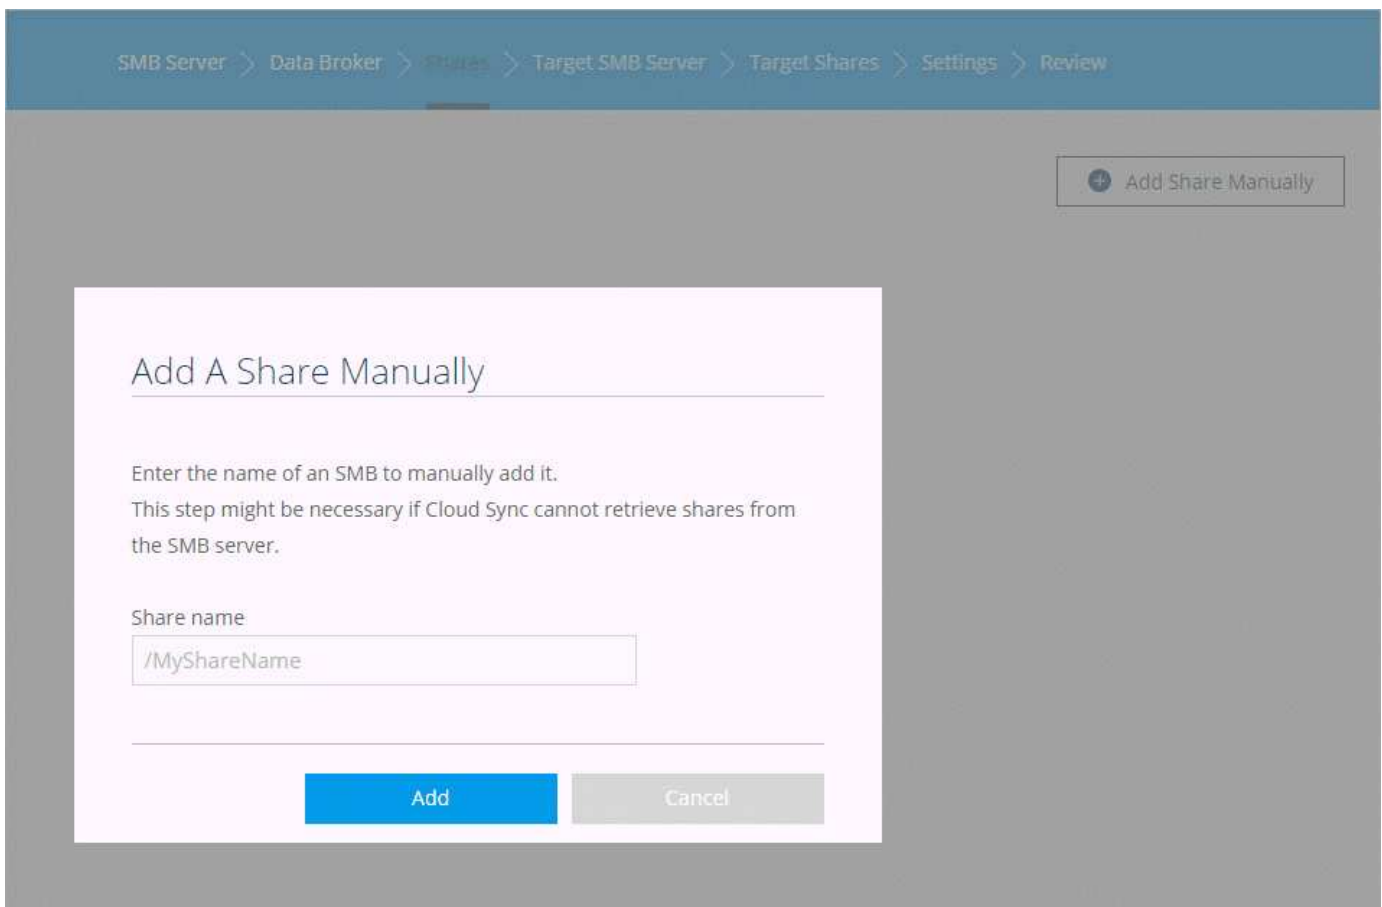
When you enable this setting, Cloud Sync deletes files from the source location after it copies the files to the target location. [Learn more about this setting.](#)

- Object tagging

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

## 5 Apr 2020

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



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

## 24 Mar 2020

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

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

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

## 8 Mar 2020

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

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

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

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

Target: s3://vadim-service-2test-empt...

AWS Data Broker: vadimBroker1

Schedule [ON] | No Tags

Synced Successfully | Sync | Duration: 2 minutes | 5 days ago

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

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

Download Data Broker Logs

## 23 Feb 2020

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

[Learn more about accounts.](#)

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

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

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

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

[View supported sync relationships](#)

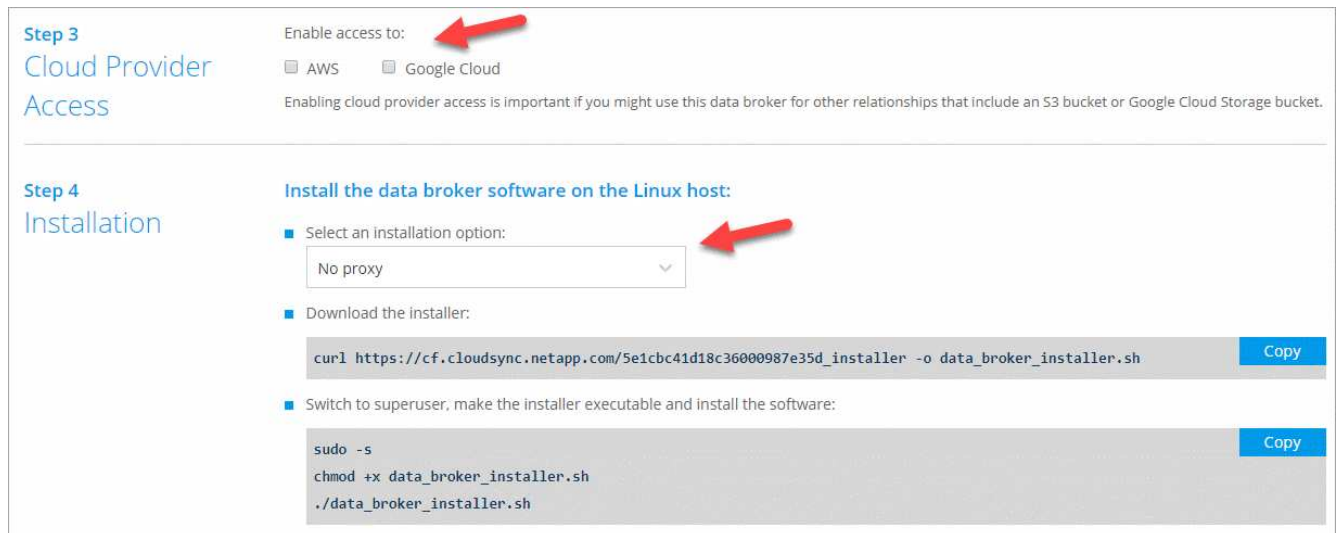
## 14 Jan 2020

- NFS 4.1 is now supported with Azure NetApp Files.



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

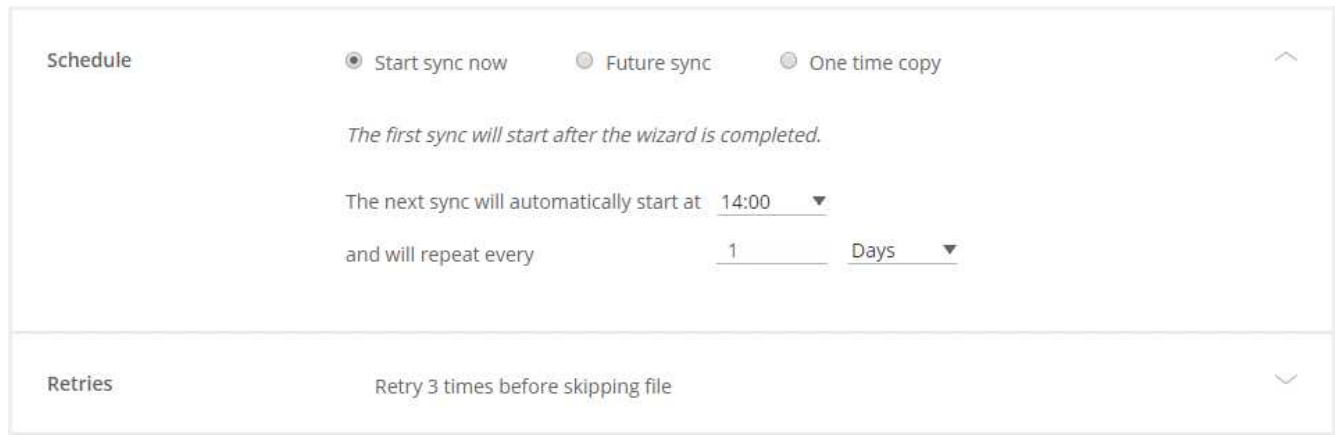
You now choose whether the data broker should have AWS access, GCP access, or both, and the type of proxy configuration that you have. Cloud Sync then shows you the right commands to use.



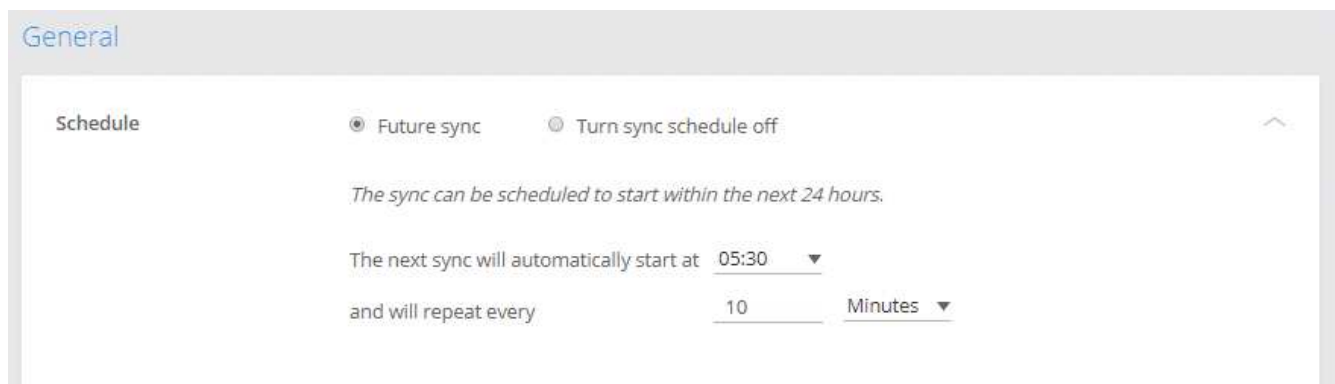
## 9 Dec 2019

- We improved how to schedule data syncs.

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



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



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


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

Here's an example:

Scan				Copy			
Succeeded	8.2K Directories	18.3K Files	4.66 GB	Succeeded	6.9K Directories	18.2K Files	4.56 GB
Failed	1.8K Directories			Failed	1 Directories	104 Files	95.37 MB
Marked for Copy	6.9K Directories	18.3K Files	4.66 GB	Deleted	0 Directories	0 Files	0 B
Marked for Delete	0 Directories	0 Files	0 B	<a href="#">View details about sync failures ↗</a>			

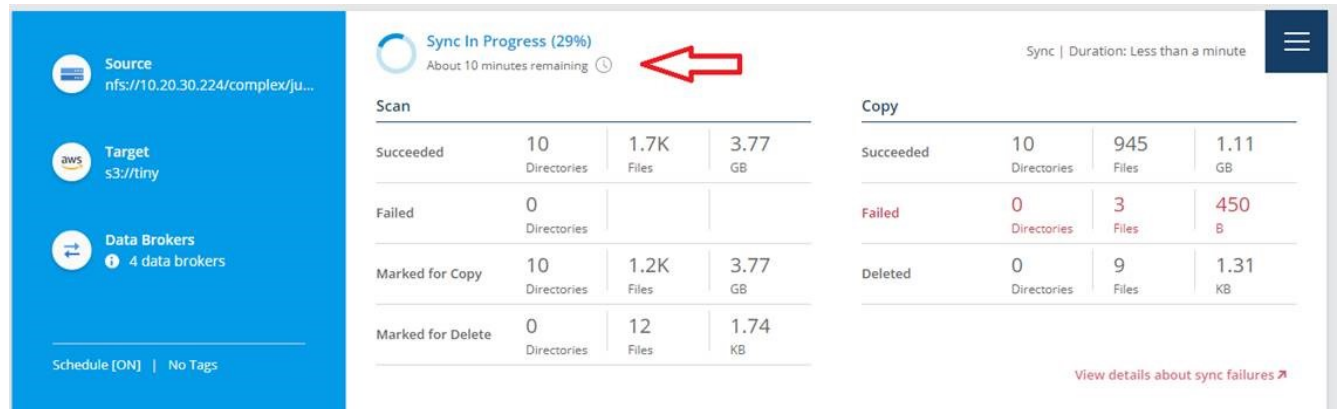
- Cloud Sync now shows you details about failures that occurred in the most recent data sync. If errors occur, you can use the failure reason to correct the issue.

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



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

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



## 17 Nov 2019

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

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

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

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

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

## 4 Nov 2019

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



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

## 8 Sept 2019

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

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

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

Date Modified

Sync files modified between a time range

From: 04/01/2019

To: 09/08/2019

## 15 July 2019

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

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

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

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

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

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

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

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

## 20 June 2019

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

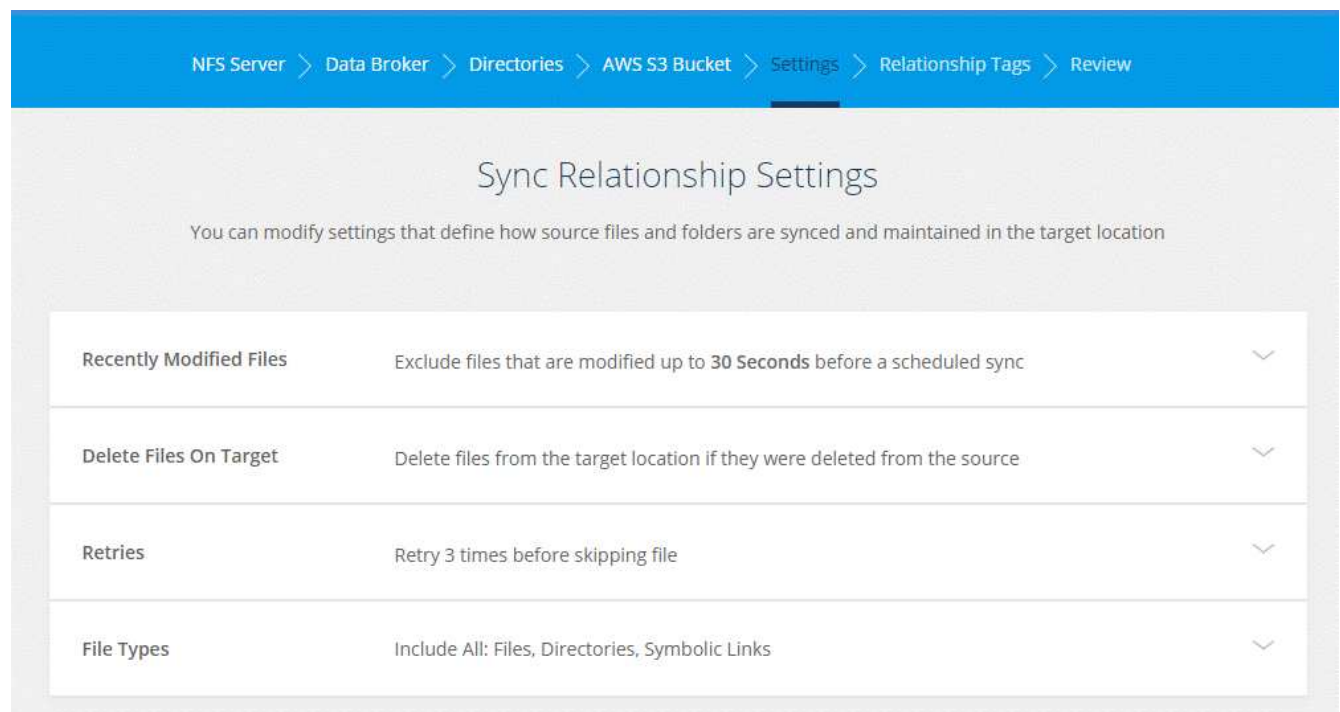
[View supported sync relationships](#)

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

[Learn about S3 storage classes](#)

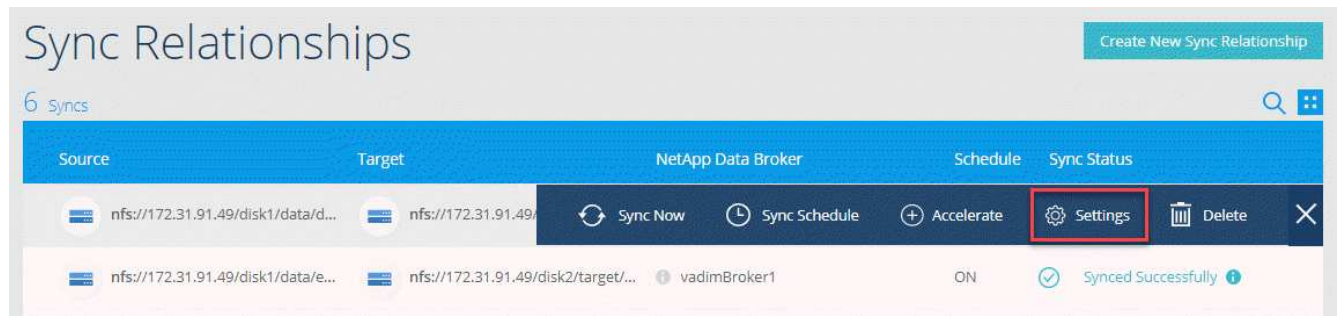
- New settings enable you to define the number of retries and file types for a sync relationship.
- Sync relationship settings were moved to a new page when setting up a relationship and when editing a relationship.

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



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





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

## 16 May 2019

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

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

## 21 Mar 2019

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

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

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

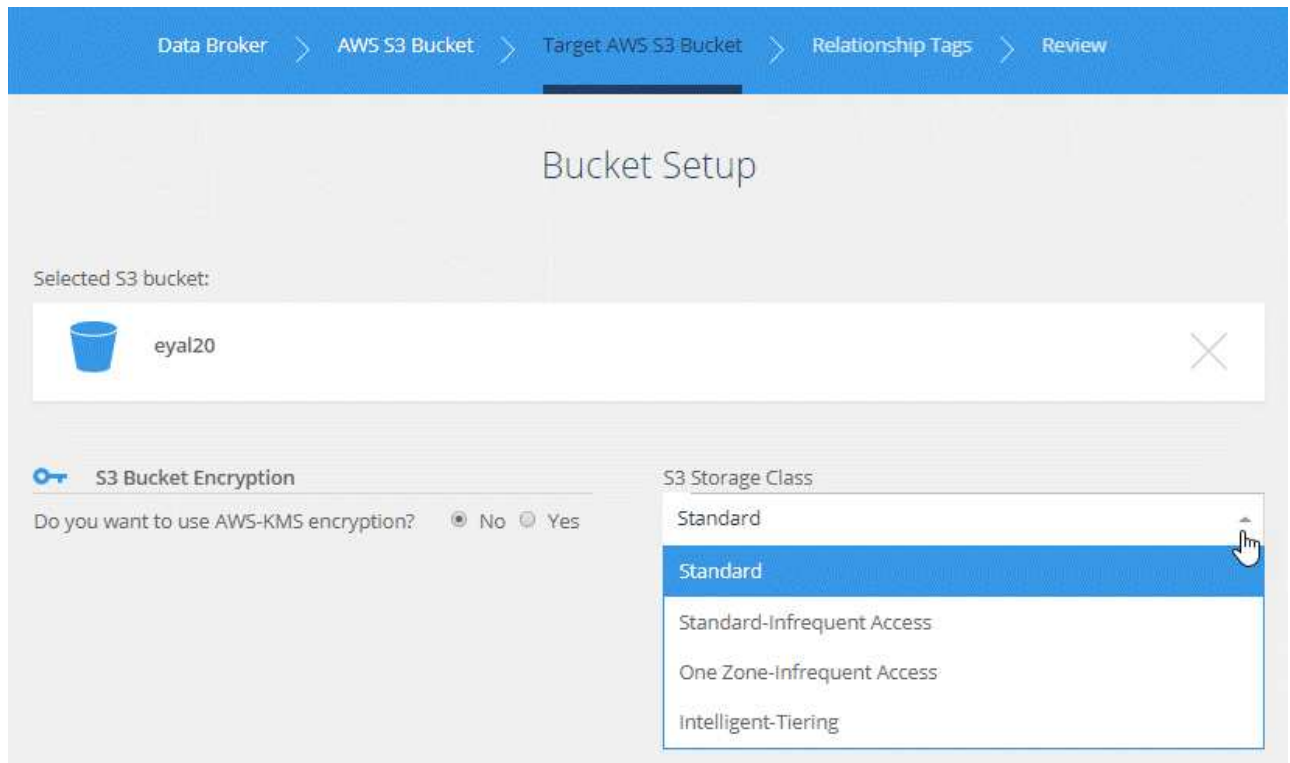
## 25 Feb 2019

Two new sync relationships are supported:

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

## 2 Dec 2018

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



## 8 Nov 2018

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

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

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

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

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

[View requirements for Azure Blob storage](#)

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

[View requirements for Azure Blob storage](#)

## 11 Oct 2018

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

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

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

or for deletion:

The screenshot shows the 'Sync Relationships' page in the Cloud Sync interface. The page title is 'Sync Relationships' and it shows '3 Syncs'. A 'Create New Sync Relationship' button is visible in the top right. The main content area displays a sync relationship with the following details:

- Source:** `s3://my-aws-bucket/10x/adam`
- Target:** `nfs://10.60.13.13/share/target/10x/adam`
- AWS Data Broker:** `hiliAWSDev`
- Schedule:** [ON] | **No Tags** | **No Custom Settings**

The sync status is 'Synced Successfully' with a duration of 'a minute' and it was completed '3 days ago'. A table shows the scan results:

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

Below the scan table, a 'Copy' section shows:

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

At the bottom of the page, there is a 'FREE TRIAL [11 Days left]' banner and an 'API API Documentation' link.

## 12 Sept 2018

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

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


The screenshot shows the 'Add a NetApp Data Broker' dialog box. It features four buttons for selecting a data broker:

- AWS Data Broker** (with the AWS logo)
- Azure Data Broker** (with the Azure logo)
- Google Cloud Data Broker** (with the Google Cloud logo)
- On-Prem Data Broker** (with a server icon)

A 'Cancel' button is located at the bottom right of the dialog.

## 21 Aug 2018

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



## 14 Aug 2018

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

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

## 17 July 2018

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

## 10 July 2018

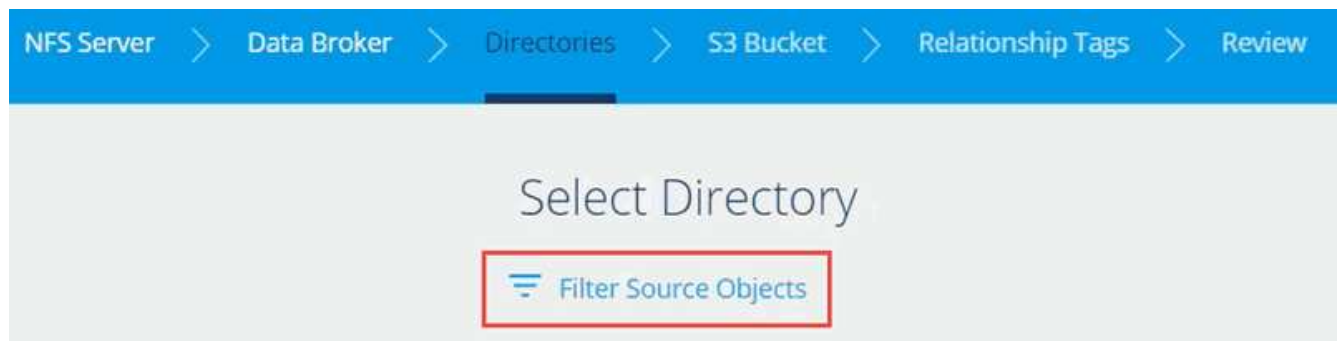
- An Azure Blob container is now supported as the source or target in a sync relationship with an NFS or SMB server.

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

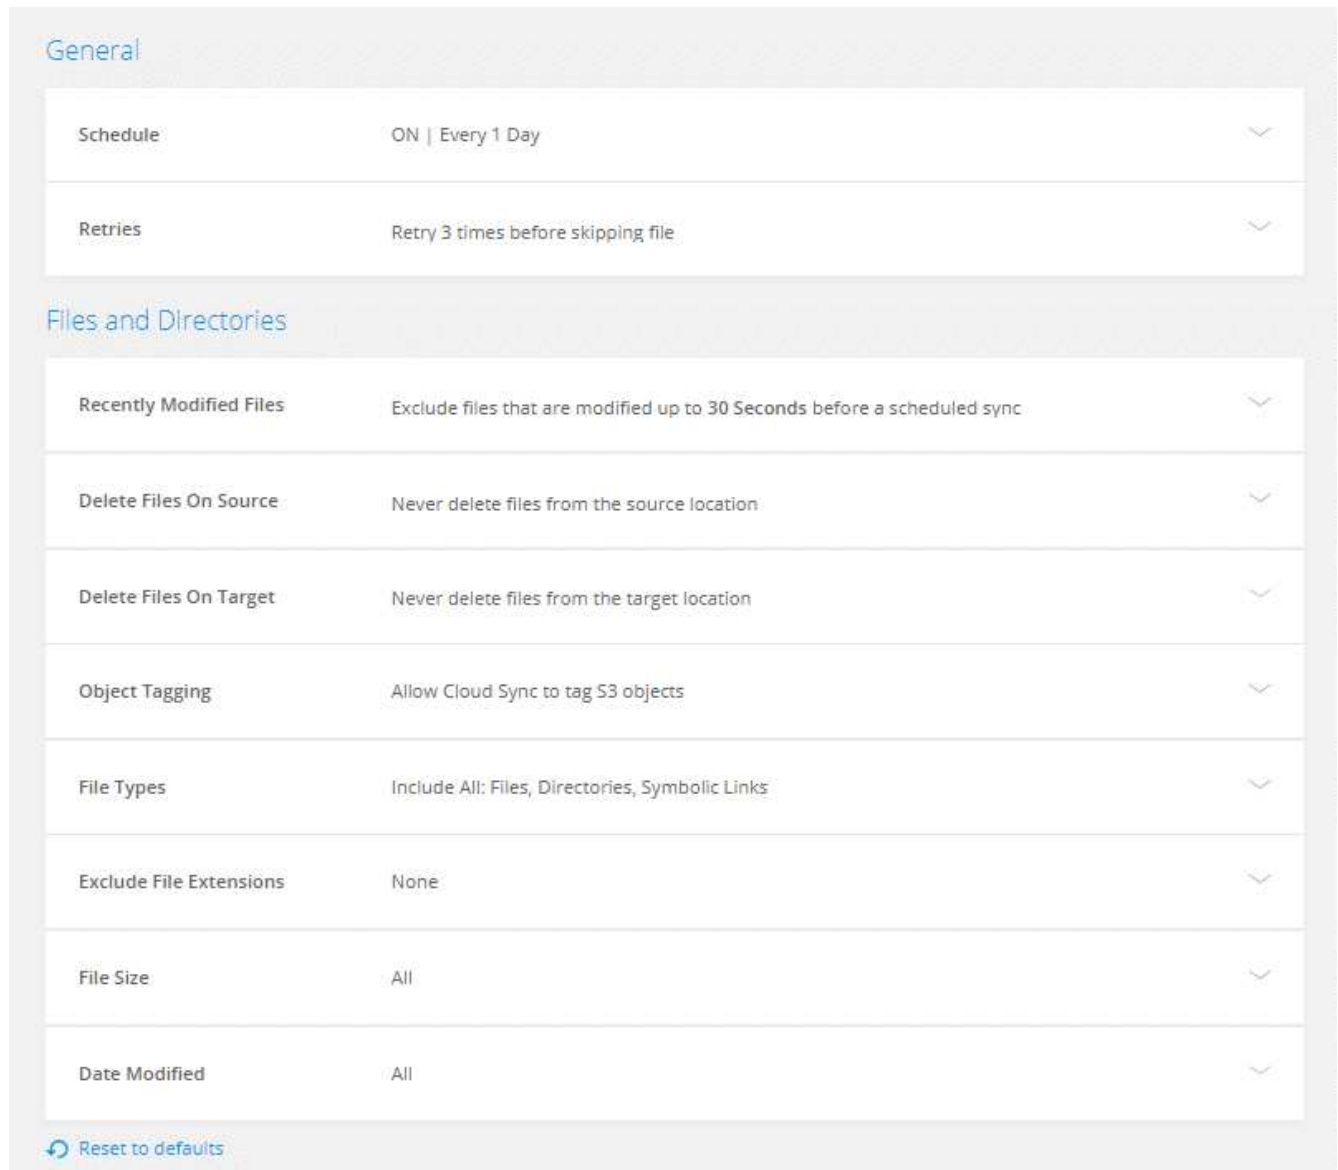
- You must now select an NFS version or SMB version when setting up a new sync relationship for an NFS or SMB server.
  - For NFS, you can select version 3, 4.1, or 4.2.
  - For SMB, you can select version 1.0, 2.0, 2.1, or 3.0.
- You can now filter source objects when setting up a new sync relationship.

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

You can access the option when selecting a directory:



The following options are available when filtering source objects:



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

## April 2018

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

## February 2018

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

## January 2018

- You can now abort an in-progress sync.

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

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

## December 2017

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

## November 2017

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

## October 2017

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

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

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

## Known limitations

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

### Unsupported regions

- Cloud Sync is not supported in China.

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

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

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

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

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

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

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

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

## Copyright information

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

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

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

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

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

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

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

## Trademark information

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