



Technical Report

Cloud Sync with NetApp E-Series Systems Solution Deployment

Kevin Wong and Austin Major, NetApp
January 2021 | TR-4880-DEPLOY

Abstract

This document describes a solution in which you can establish a NetApp® Cloud Sync relationship between two remote NetApp E-Series systems using a basic setup of an NFS server with Windows 2016 and Red Hat 7.9.

TABLE OF CONTENTS

Overview of Cloud Sync and E-Series systems	3
Use cases.....	3
Technology.....	3
Pricing	4
Install NFS	4
Requirements	4
NFS for RHEL installation.....	6
NFS for Windows installation.....	8
Create a sync relationship	14
Install the data broker.....	18
On-Prem Data Broker	18
On-Prem Data Broker with AWS	20
AWS Data Broker	21
Share the data broker	26
Uninstall the data broker	27
Conclusion	28
Where to find additional information	28
Version history.....	28

LIST OF TABLES

List 1) Hardware requirements.	5
-------------------------------------	---

LIST OF FIGURES

Figure 1) Cloud Sync overview.....	3
------------------------------------	---

Overview of Cloud Sync and E-Series systems

Cloud Sync is a data migration tool offered by NetApp that supports platforms such as AWS, Azure, Google Cloud, IBM Cloud Object Storage, NetApp® ONTAP®, NFS, and SMB. It can be accessed via a GUI, REST API, or the CLI.

There are a number of approaches to moving data between storage systems. Cloud Sync is one such approach that is supported by NetApp. It is an efficient, user-friendly tool that works with a variety of storage systems and storage system vendors.

To use Cloud Sync with NetApp E-Series systems, you must perform extra steps at the host level. This document describes how to get an NFS server running for both Windows 2016 and Red Hat 7.9 and then connect to NetApp E-Series storage systems.

Use cases

Although Cloud Sync is a generic solution that can be used in many different solutions, a solution may already exist for certain specific use cases. These specialized solutions may offer better tuning for a better experience.

With Cloud Sync, you can perform tasks such as the following:

- Moving data from a file system backed by an E-Series system to cloud backup, such as AWS Glacier.
- Moving data from a file system backed by an E-Series system to a NetApp StorageGRID® object store.
- Moving data from a cloud backup to an off-site E-Series system.

Note: NetApp E-Series has native support for Asynchronous Remote Volume Mirroring (ARVM), which may be better suited to this task.

Technology

Cloud Sync is a software-as-a-service (SaaS) solution that consists of a data broker, a cloud-based platform, a source, and a target. The data broker syncs the data from the source to the target, acting as a middleman for the transactions. The data broker software can be installed in AWS, Azure, or Google Cloud, or on a Linux system on the premises.

Figure 1 shows the relationship between Cloud Sync components.

Figure 1) Cloud Sync overview.


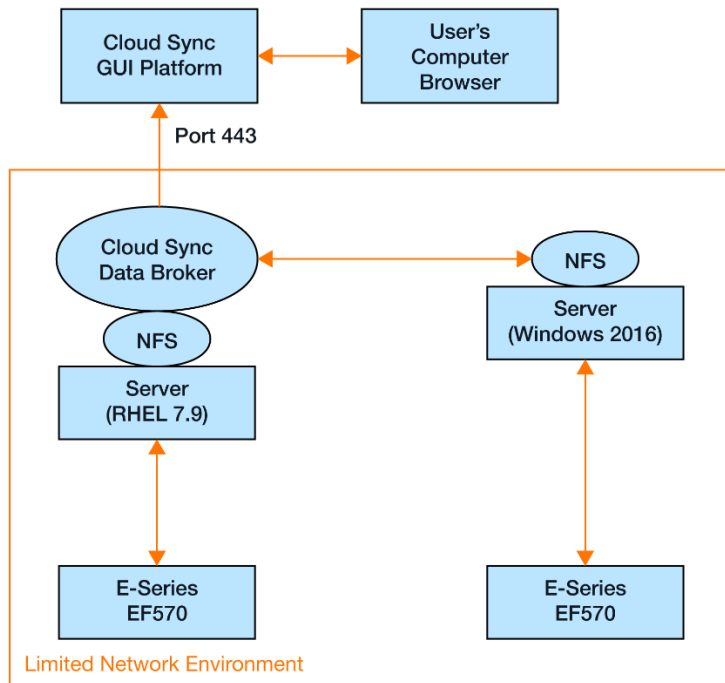


Figure 2 shows one possible configuration, with NetApp E-Series storage systems serving as both the source and the target. Because the data broker must be able to connect to both the target and the source, you might want to optionally install Cloud Sync onto the same system as the NFS server. The Cloud Sync Data Broker service should be installed on the Linux server inside the limited network environment, and an outbound connection should be allowed on port 443. This makes it possible to use the GUI outside of the environment while interacting with the connections inside the limited network environment.

Figure 2) Solution overview.



Pricing

You can select one of two pricing models for the Cloud Sync services:

- Hourly rate that is based on a maximum number of host relationships
- Annual payment

For current pricing information and information about how to acquire a license, go to [How Cloud Sync licenses work](#).

Install NFS

This solution demonstrates how to establish a simple synchronization relation between an NFS server on Windows 2016 and an NFS server running on Red Hat 7.9. This is a generic solution to get Cloud Sync working with NetApp E-Series systems.

Requirements

This section covers the technology requirements for the primary use case.

Network

You must have outbound internet access over port 443 to communicate with the Cloud Sync service for the data broker. The browser for initiating the synchronization relationship must also have port 443 open.

This solution describes a limited-network environment with internet access, but with limited inbound services. Therefore, a specific type of data broker is used to alleviate this pain point.

For full information about endpoints, go to [Endpoints that are required for Cloud Sync](#).

Hardware

This solution uses two hosts and two storage systems. The models of the storage systems and hosts do not have to be the ones addressed in Table 1. You simply need two systems that are attached to the E-Series system and that have a network path to communicate. Although this particular solution can be used on two Dell R720s or two NetApp HCI H615Cs, we opted to use both systems in the relationship to ensure correct functionality.

The following list shows the hardware components that are required to implement the use case described in this document.

List 1) Hardware requirements.

NetApp HCI H615C (Model H615-75031): Quantity 1

- Two Intel Xeon Gold 6242 processors @ 2.8GHz
- 512GB RAM
- iSCSI protocol used the base networking 2x 10/25Gbe (Mellanox Connect-X 4)

Dell R720: Quantity 1

- Two Intel Xeon CPU E5-2620 0 @ 2.00GHz
- 32GB RAM
- iSCSI protocol used the base networking 2x 10/25Gbe (Mellanox Connect-X 4)

NetApp EF280 all-flash array: Quantity 2

- 25Gb iSCSI HICs for the iSCSI solution
 - Built-in 16Gb FC baseboard ports for the FC solution
 - 24 MZILS15THMLS-0G4 - 15TB SSD drives
 - One disk pool spanning all drives
 - 8 volumes presented to host
-

Software

This solution can be used on different operating systems, but the steps may vary. Guides are available to create an NFS server for most operating systems.

List 2 contains the list of software components that are required to implement the exact use case described in this document.

List 2) Hardware requirements.

NetApp Cloud Sync License: Quantity 1

Windows 2016 OS License: Quantity 1

- Installed on the NetApp HCI VM
-

VMware ESXi 6.5.0 Update 3: Quantity 1

- Installed on the NetApp HCI VM
-

Red Hat Enterprise Linux 7.9: Quantity 1

- Installed on Dell 720
-

NFS for RHEL installation

The following instructions assume that the NetApp E-Series system has already been cabled and connected.

Note: The cabling and configuration of the NetApp E-Series system to the server is out of scope of this document. Refer to the [appropriate guide for configuration](#).

To install NFS for RHEL, follow these steps.

1. Locate the block device to use.
 - a. To find the specific volume ID, open NetApp SANtricity® System Manager. Go to Storage > Volumes. Click the desired volume and then select View/Edit settings. The WWID and EUI appear under Identifiers.

Volume Settings

Basic

Advanced

Name

Cloud_Sync_vol1

Capacities

Reported capacity (GiB): 2048.00

Allocated capacity (GiB): 2048.00

Pool/volume group

Name: Pool Disk_Pool

RAID level: Pool

Secure-capable: No

Secure-enabled: No

Data Assurance (DA) capable: Yes

Host

Assigned to

Host baremetal-rhel

LUN

1

Identifiers

World-wide identifier (WWID): 60:0A:09:80:00:A8:C8:9E:00:00:DD:40:5F:B6:B9:E4

Extended unique identifier (EUI):

Subsystem identifier (SSID): 0

Save

Cancel

b. You can then match that ID on the Red Hat server, as shown in the following example.

```
[root@localhost ~]# multipath -ll
...
3600a098000a8c89e0000dd405fb6b9e4 dm-27 NETAPP ,INF-01-00
size=2.0T features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='1
alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
   `- 10:0:0:2 sdh 8:112 active ready running
...
```

2. Create a file system on the block device. You can use EXT4 or XFS. You must make a directory for it, mount that directory to the file system, and then give it higher permissions.

```
# mkfs.xfs /dev/mapper/<BLOCK_DEVICE>
mkfs.xfs /dev/mapper/3600a098000a8c89e0000dd405fb6b9e4
#mkdir /mnt/<DIRECTORY_NAME>
mkdir /mnt/data

#mount /dev/mapper/<BLOCK_DEVICE> /mnt/<DIRECTORY_NAME>
mount /dev/mapper/3600a098000a8c89e0000dd405fb6b9e4 /mnt/data
```

```
#chmod 777 /mnt/<DIRECTORY_NAME>
chmod 777 /mnt/data/
```

3. Make sure that the file system is mounted in case the system were to lose power or be rebooted. The `_netdev` attribute tells the system to wait for the network to establish before it is remounted.

WARNING: Make sure that the connection is created and established before `fstab` is called (for example, iSCSI sessions have been established).

```
# echo '/dev/mapper/<BLOCK_DEVICE> /mnt/<DIRECTORY_NAME> <FILESYSTEM> _netdev 0 0' >> /etc/fstab
echo '/dev/mapper/3600a098000a8c89e0000dd405fb6b9e4 /mnt/data xfs _netdev 0 0' >> /etc/fstab
```

4. Install `nfs-utils`.

```
yum install nfs-utils
```

5. Set up the NFS export.

Note: If you have other file systems, you may need to change your `fsid` to an unused number. You should also understand the implications of the `sync` option, which ensures that data is written to persistent memory before being acknowledged. This option leads to a performance impact at the cost of ensuring that files are secure in various unexpected events such as a power outage. For this reason, NetApp strongly recommends using `sync`.

```
#bash -c 'echo "/mnt/<DIRECTORY_NAME> *(rw, sync, fsid=1, no_subtree_check, no_root_squash)" >> /etc/exports'
bash -c 'echo "/mnt/data *(rw, sync, fsid=1, no_subtree_check, no_root_squash)" >> /etc/exports'
exportfs -a
```

6. Enable the services to come up on reboot.

```
systemctl enable --now rpcbind
systemctl enable --now nfs-server
```

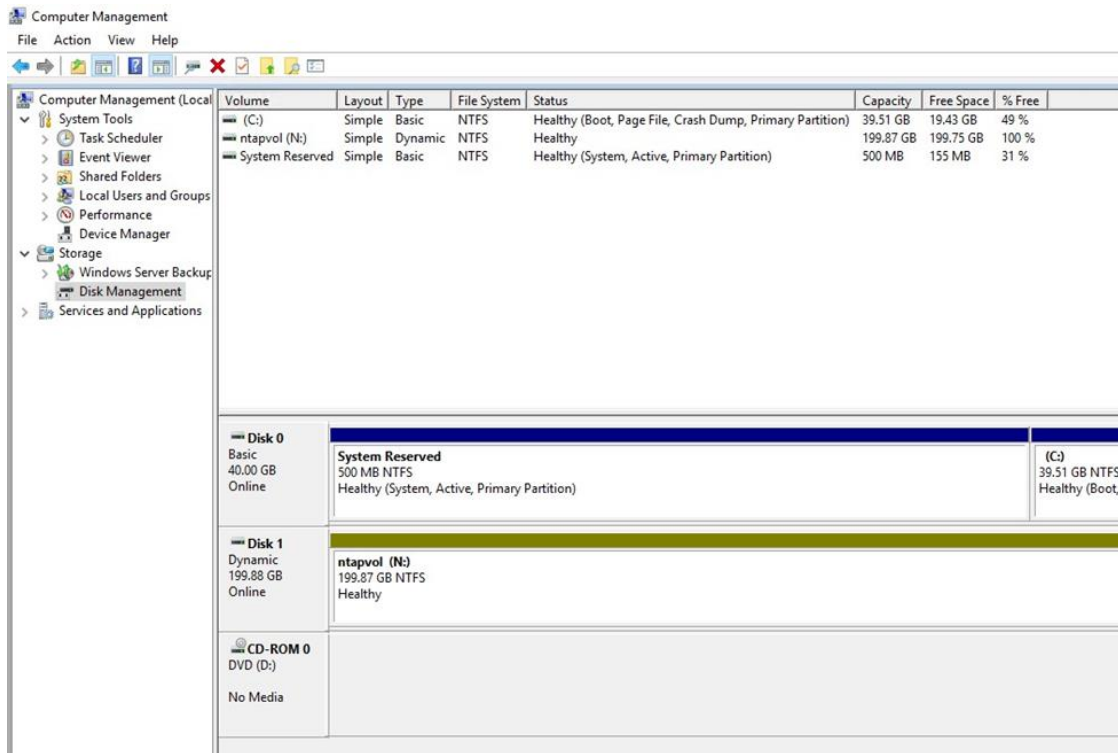
7. Test whether the mount can be written to. If the following can be executed without errors, you are finished.

```
mkdir /mnt/data_via_nfs
chmod 777 /mnt/data_via_nfs
mount -t nfs -overs=4.1 localhost:/mnt/data /mnt/data_via_nfs/
touch /mnt/data_via_nfs/test_file
```

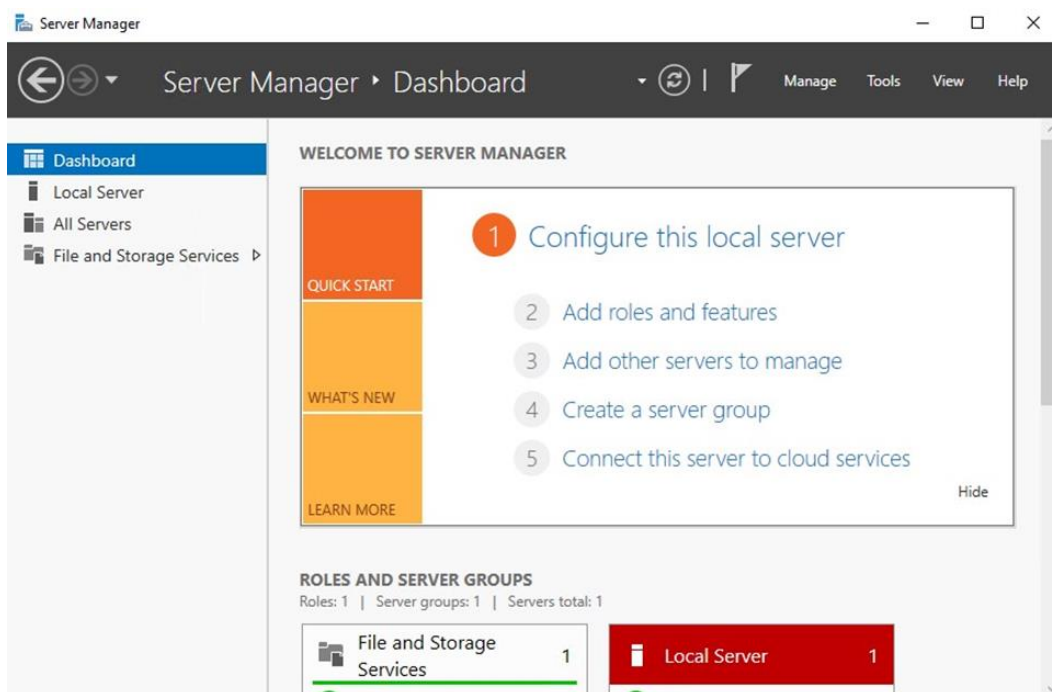
NFS for Windows installation

To install NFS for Windows, follow these steps.

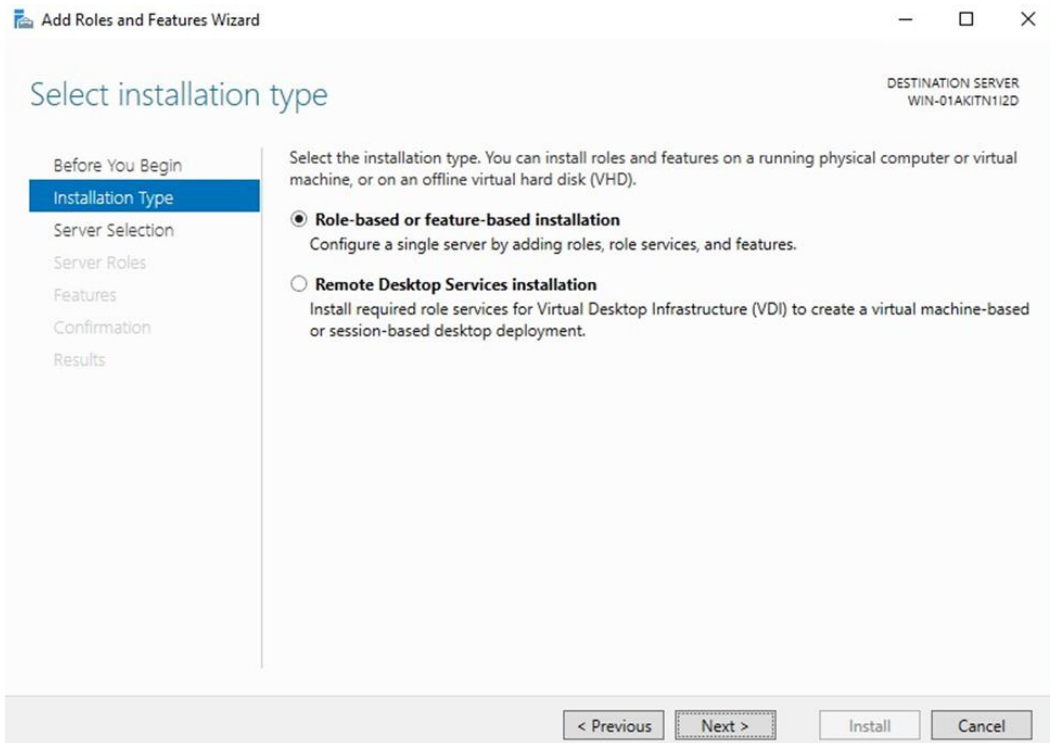
1. In Windows Disk Management, locate the disk to use. Right-click the disk and format it into an NTFS file system.



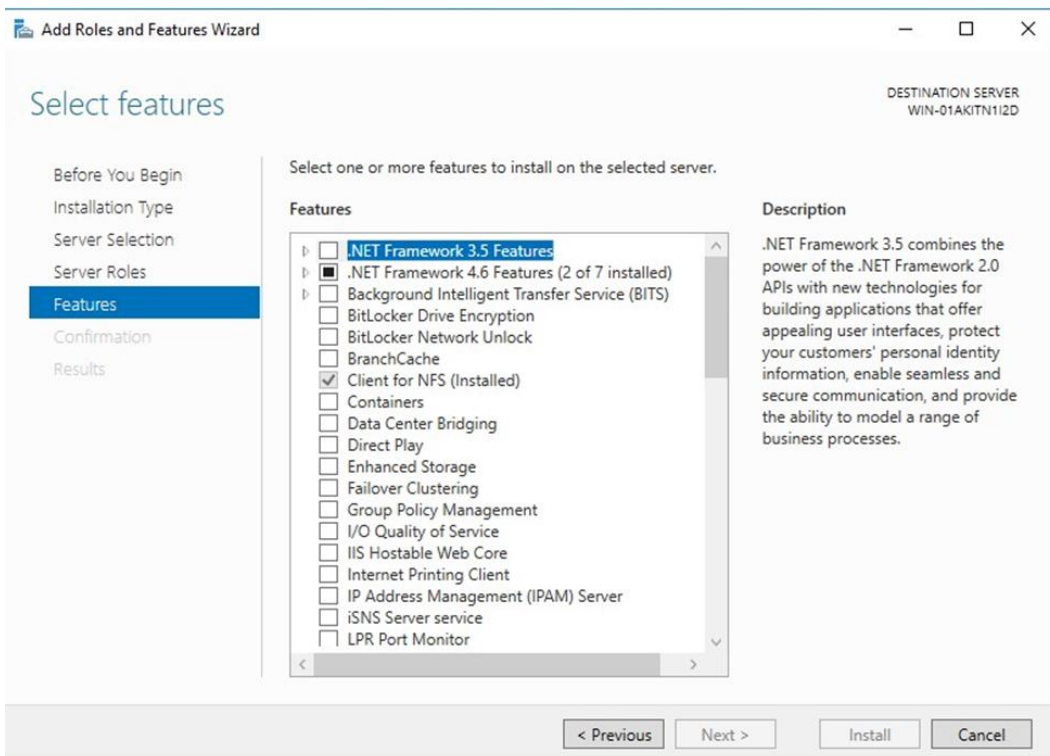
2. To install the NFS Client, go to the Windows Server Manager and click Add Roles and Features.



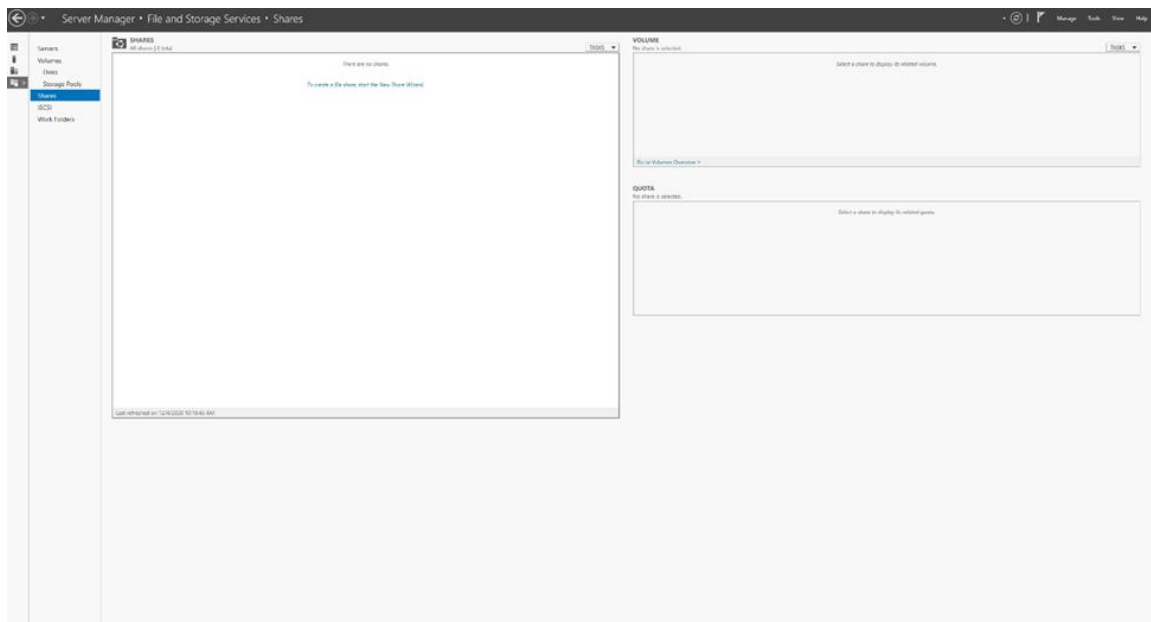
3. In the wizard, select Role-Based or Feature-Based Installation.



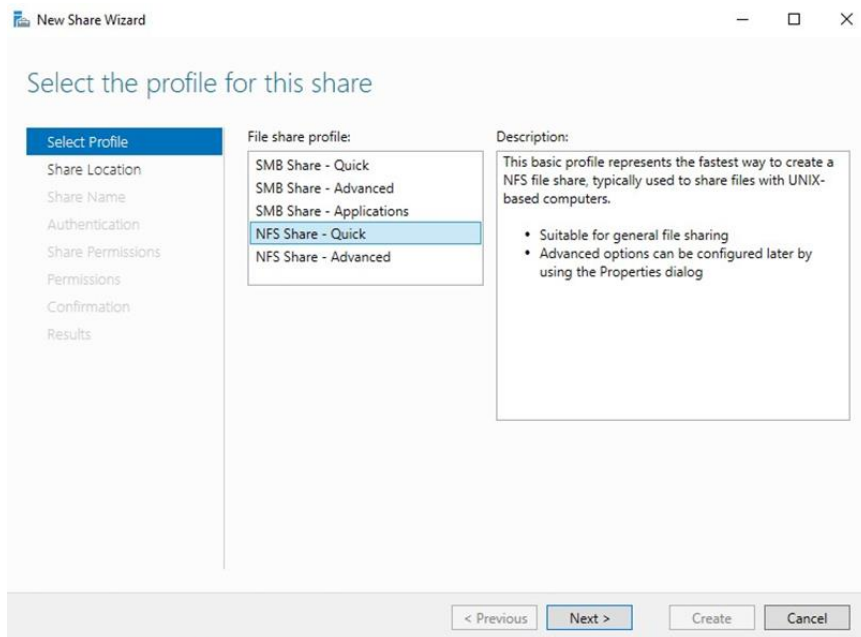
4. Skip Server Roles and move on to features. Check Client for NFS (Installed), click Next, and then click Install.



5. After installing the necessary tools, return to Server Manager, select File and Storage Service, and then select Shares.



6. To start the file share wizard, open the Tasks drop-down menu or click the link "To create a file share, start the New Share Wizard."
7. Use NFS Share - Quick profile. These settings are generally fine for most setups; advanced settings can be done through the Share properties.



8. Locate and highlight your server and the appropriate file system. In this case, put the file system that represents the NetApp E-Series volume in volume N:.

New Share Wizard

Select the server and path for this share

Select Profile
Share Location
Share Name
Authentication
Share Permissions
Permissions
Confirmation
Results

Server:

Server Name	Status	Cluster Role	Owner Node
WIN-01AKITN112D	Online	Not Clustered	

i The list is filtered to show only servers that have Server for NFS installed.

Share location:

☒ Select by volume:

Volume	Free Space	Capacity	File System
C:	19.5 GB	39.5 GB	NTFS
N:	200 GB	200 GB	NTFS

The location of the file share will be a new folder in the \Shares directory on the selected volume.

☐ Type a custom path:

9. Give the share a name.

New Share Wizard

Specify share name

Select Profile
Share Location
Share Name
Authentication
Share Permissions
Permissions
Confirmation
Results

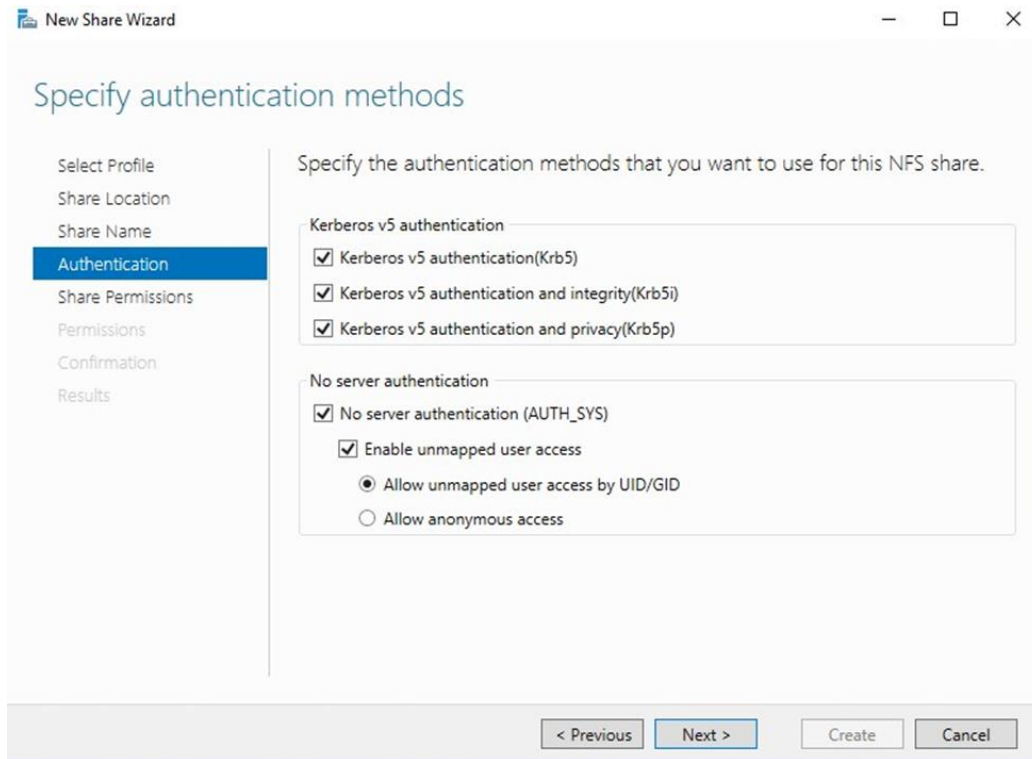
Share name:

Local path to share:

i If the folder does not exist, the folder is created.

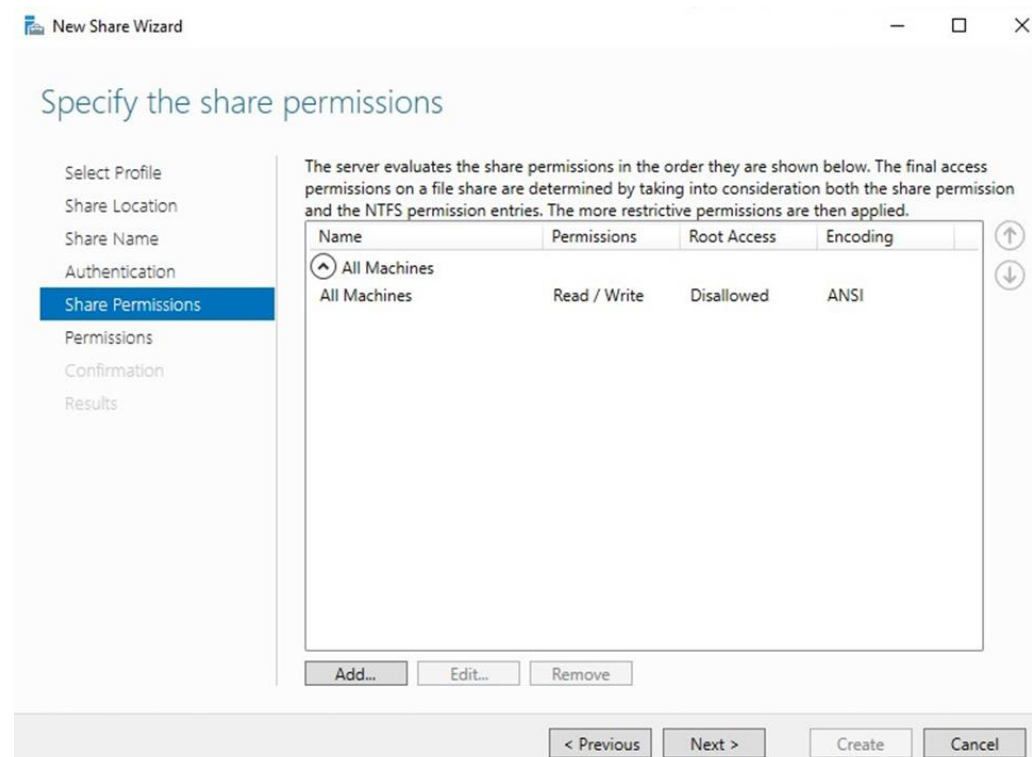
Remote path to share:

10. Change the next settings as desired. In this case, you can allow unmapped users to authenticate to the NFS server for ease of use.



The screenshot shows the 'Specify authentication methods' step of the 'New Share Wizard'. The left sidebar lists steps: Select Profile, Share Location, Share Name, Authentication (selected), Share Permissions, Permissions, Confirmation, and Results. The main area has the title 'Specify authentication methods' and a subtitle 'Specify the authentication methods that you want to use for this NFS share.' There are two sections: 'Kerberos v5 authentication' with three checked options: 'Kerberos v5 authentication(Krb5)', 'Kerberos v5 authentication and integrity(Krb5i)', and 'Kerberos v5 authentication and privacy(Krb5p)'. Below it is 'No server authentication' with 'No server authentication (AUTH_SYS)' checked, and 'Enable unmapped user access' checked. Under 'Enable unmapped user access', 'Allow unmapped user access by UID/GID' is selected with a radio button, and 'Allow anonymous access' is unselected. At the bottom are buttons: '< Previous', 'Next >', 'Create', and 'Cancel'.

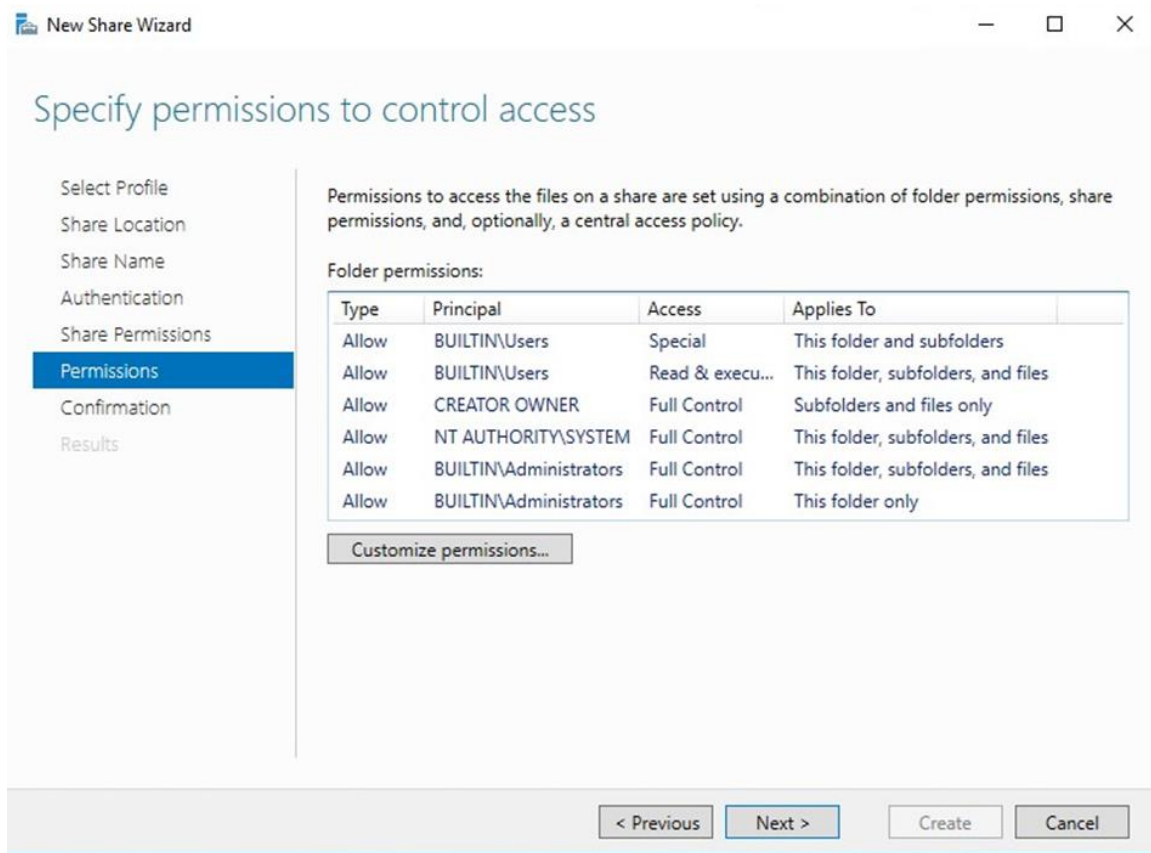
11. Allow all machines Read/Write access for this example. This access can vary based on how the permissions are structured in the environment.




The screenshot shows the 'Specify the share permissions' step of the 'New Share Wizard'. The left sidebar lists steps: Select Profile, Share Location, Share Name, Authentication, Share Permissions (selected), Permissions, Confirmation, and Results. The main area has the title 'Specify the share permissions' and a subtitle 'The server evaluates the share permissions in the order they are shown below. The final access permissions on a file share are determined by taking into consideration both the share permission and the NTFS permission entries. The more restrictive permissions are then applied.' Below this is a table with columns: Name, Permissions, Root Access, and Encoding. The table has one entry: 'All Machines' (with a caret icon) in the Name column, 'Read / Write' in the Permissions column, 'Disallowed' in the Root Access column, and 'ANSI' in the Encoding column. There are up and down arrow buttons to the right of the table. Below the table are buttons: 'Add...', 'Edit...', and 'Remove'. At the bottom are buttons: '< Previous', 'Next >', 'Create', and 'Cancel'.

Name	Permissions	Root Access	Encoding
^ All Machines	Read / Write	Disallowed	ANSI

12. Use the default permissions; these can differ based on the environment.



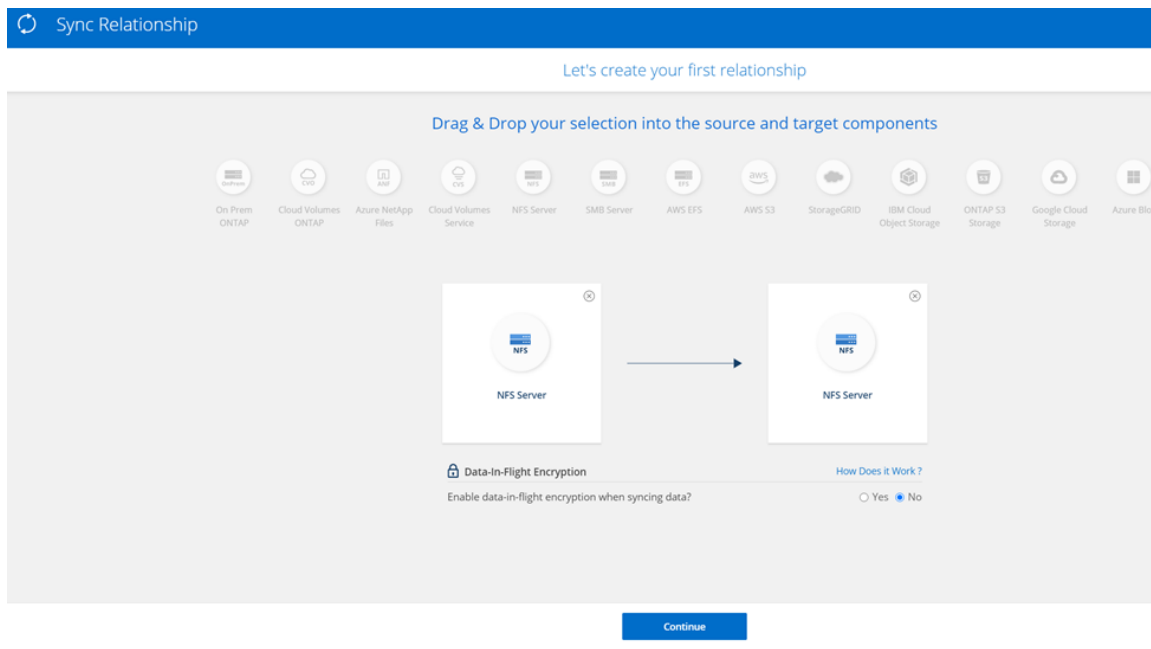
13. (Optional) To test whether files can be easily written, go to Share and make an empty file. The Share will be in the following directory:



Create a sync relationship

After creating the NFS source and NFS targets, follow these steps to create the relationship.

1. Log into <https://www.cloudsync.netapp.com>.
2. Drag the NFS Server button onto Drag Source Here and Drag Target Here. For this example, do not enable in-flight encryption.




- Specify which system is the NFS source and select the NFS version. In virtually all cases, you should specify the latest version that can be supported. In this case, up to NFS 4.1 is supported.

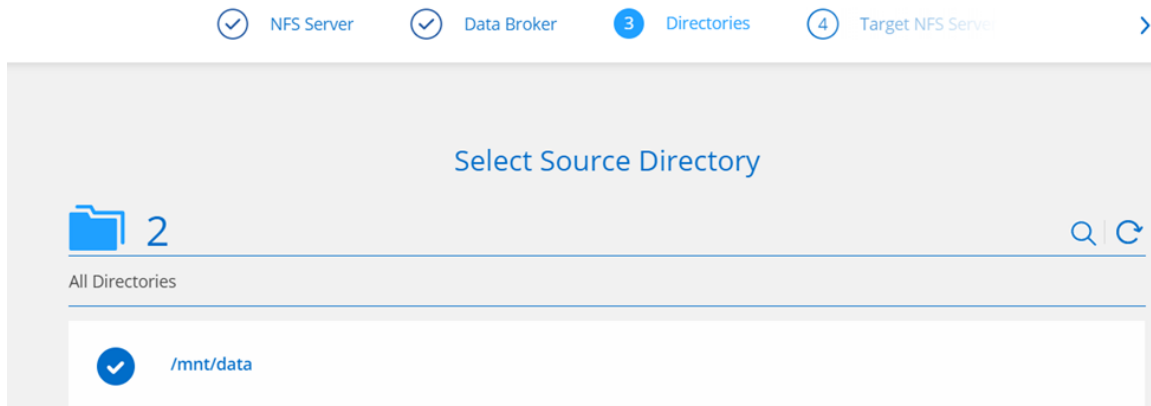


- Select a data broker.

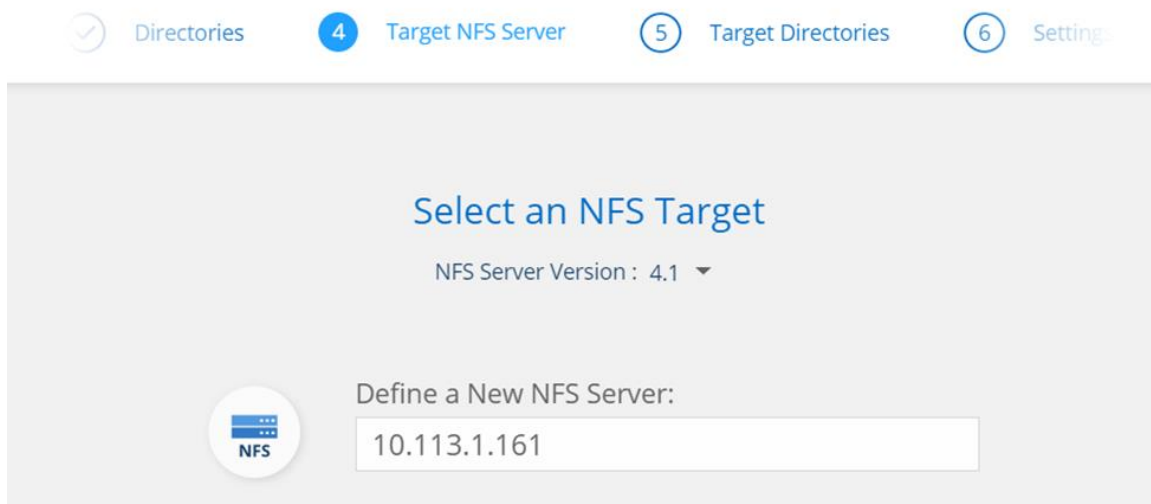
For this example, there are two NetApp E-Series systems in a limited-network environment. The data broker will be installed on the Linux server in the environment so that the data broker can communicate to both servers in the limited-network environment. Refer to “Install the Data Broker,” later in this document, to understand and install the data broker that best fits your needs.

3 Data Brokers		
<div>  redhat-prem-broker Active </div>		
localhost.localdomain Hostname	linux Platform	5fcf03abac1e7e1308ffb2e8 Broker ID
10.113.1.130 Private IPs		

- Select the directory for the sync.



- Select the target NFS server. In this case, use an IP address instead for the latest NFS version that can be supported.



- Select the target directory.
- Select the settings to edit. For example, we will set this operation to sync now and schedule a sync every day at 10:00 PM. In addition, we will look to never delete files on the source or target if a particular file is missing.

Directories

Target NFS Server

5 Target Directories

6 Settings

Target NFS Server

Target Directories

6 Settings

7 Review

General

Schedule

☒ Start sync now ☐ Future sync ☐ One time copy

The first sync will start after the wizard is completed.

The next sync will automatically start at 22:00

and will repeat every 1 Days

Retries

Retry 3 times before skipping file

Files and Directories

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

[Reset to defaults](#)

Continue

On the dashboard, you should see Synced Successfully.

2 Syncs

Source
nfs://localhost/mnt/data

Target
nfs://10.113.1.161/nfs_cloud_...

On-Prem Data Broker
redhat-prem-broker

Schedule [ON] | No Tags

Synced Successfully

Initial Copy | Duration: 2 minutes | 5 minutes ago

Scan

Succeeded	1	0	0
	Directories	Files	B
Failed	0		
	Directories		
Marked for Copy	0	0	0
	Directories	Files	B
Marked for Delete	0	0	0
	Directories	Files	B

Copy

Succeeded	0	0	0
	Directories	Files	B
Failed	0	0	0
	Directories	Files	B
Deleted	0	0	0
	Directories	Files	B

Download Data Broker Logs

9. (Optional) You can do some testing by creating a large file, for example with the following command, and rerunning the sync. Click the three dots and then click Sync Now.

Note: There is a setting to exclude files that are modified up to 30 seconds before the scheduled sync.

```
truncate -s 100G /mnt/data/test_file.txt

# Create 100MB file with non-human readable data
dd if=/dev/urandom of=file.txt bs=1048576 count=100
```

Install the data broker

On-Prem Data Broker

An On-Prem Data Broker is useful in environments with limited network access and if the user does not have or won't need an Amazon Web Services account or service.

Requirements

Make sure that your environment meets the following requirements:

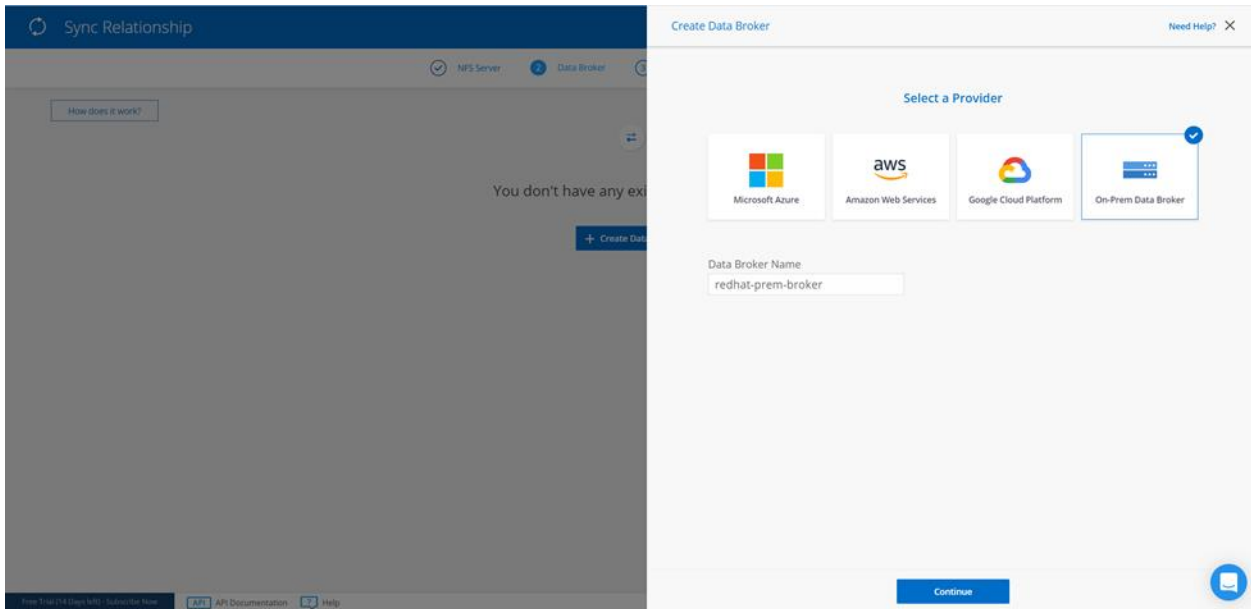
- Centos 7.0, 7.7, 8.0
- Red Hat Enterprise Linux 7 and 8.0
- Ubuntu Server 20.04 LTS
- SUSE 15 SP1
- Internet access with outbound port 443 open
- 16GB RAM, 4 CPUs
- OpenSSL is installed
- The data broker must have access to both the source and the target

For additional information, see [Installing the data broker on a Linux host](#).

Installation

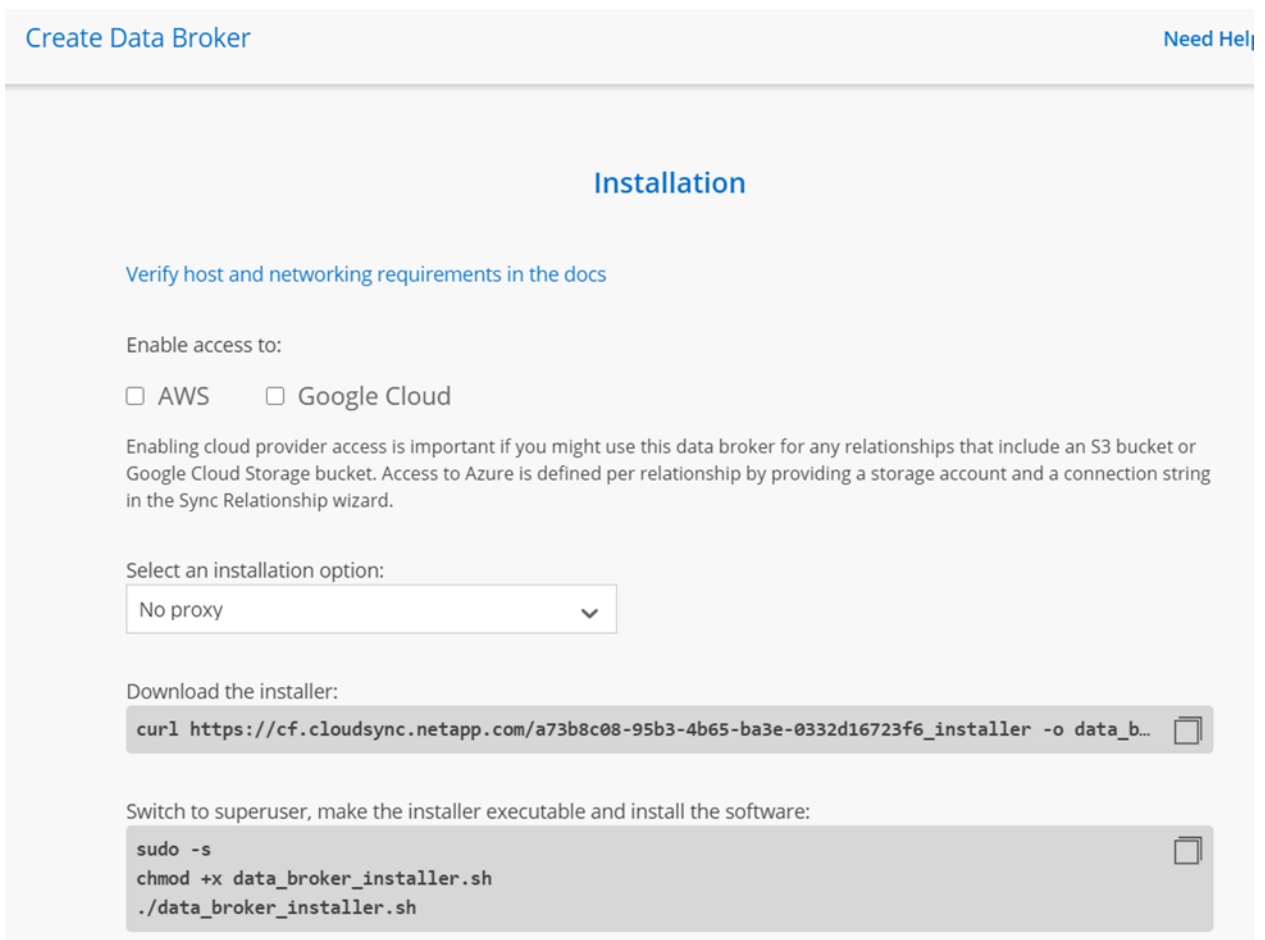
To install the data broker, follow these steps.

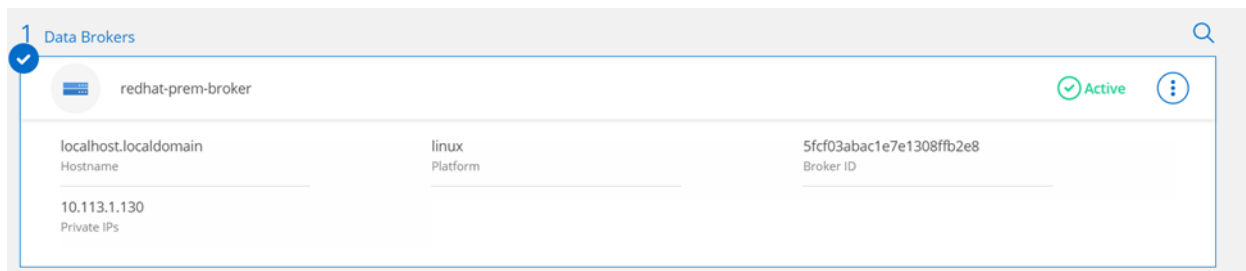
1. Create a data broker. This example shows a data broker on the Red Hat server, which is needed for its closed environment. In this closed environment, you need 443 open as an outbound port.



2. Click Create Data Broker and then click Continue. These steps create a personalized installer and instructions to set up the broker. When you have completed the steps, click Continue.

You should see the following screen.





Note: If you see issues, make sure that `pm2-root` is started.

```
[root@localhost]# systemctl status pm2-root
```

- pm2-root.service - PM2 process manager
 - Loaded: loaded (/etc/systemd/system/pm2-root.service; enabled; vendor preset: disabled)
 - Active: active (running) since Mon 2020-12-07 22:46:20 CST; 1s ago
 - Docs: <https://pm2.keymetrics.io/>
 - Process: 22726 ExecStart=/usr/lib/node_modules/pm2/bin/pm2 resurrect (code=exited, status=0/SUCCESS)
 - Main PID: 34315 (node)

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

On-Prem Data Broker with AWS

If you have a limited networking environment and want to have access to an external AWS data source or target, follow these instructions.

Requirements

To enable access to AWS, see [Enabling access to AWS](#).

Installation

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

4. Enter a name for the data broker and click Continue.
5. On the instructions page:
 - a. Enable access to AWS.
 - b. Specify No Proxy.
 - c. Execute the following commands:

```
#download the custom installer script
curl https://cf.cloudsync.netapp.com/<Unique_Key_Installer> -o data_broker_installer.sh
#ignore if root
sudo -s
#assign executable permissions to installer script
chmod +x data_broker_installer.sh
#run installer script
./data_broker_installer.sh -a <aws_access_key> -s <aws_secret_key>
```

- d. Complete the pages in the wizard to create the new sync relationship.

AWS Data Broker

If you have an open network environment, follow these instructions. You are required to set up your Amazon Web Services account. This is a convenient data broker, which is easy to set up and maintain, and is generally the recommended method for an open network environment.

Requirements

For supported AWS regions, networking requirements, and permissions, see [Installing the data broker in AWS](#).

Installation

1. On the Data Broker page, click Create Data Broker and then select Amazon Web Services.

Create Data Broker [Need Help?](#) ✕

Select a Provider

Microsoft Azure

aws
Amazon Web Services

Google Cloud Platform

On-Prem Data Broker

Data Broker Name
aws-data-broker

Continue

2. Enter a name for the data broker and click Continue.
3. Use one of the following two methods for deploying a data broker to your AWS account.
First method for deploying a data broker to your AWS account:
 - a. Enter an AWS access key so that Cloud Sync can create the data broker in AWS.
 - b. Select a location for the instance, select a key pair, choose whether to enable a public IP address, and then select an existing IAM role. Or leave the field blank so that Cloud Sync creates the role for you.
 - c. If a proxy is required for internet access in the VPC, specify a proxy configuration.

Create Data Broker in AWS
Need Help? X

1 Credentials
2 Settings
3 Proxy

Previous Step

Basic Settings

Location

VPC
vpc-04c98d05a99fd4402 - 172.30.0.0/24


Subnet
172.30.0.0/26

Connectivity

Key Pair
parallel-cluster-key

Enable Public IP?
☒ Enable ☐ Disable

IAM Role (optional) ⓘ

Continue


- d. When the data broker is available, click Continue in Cloud Sync.
- e. Complete the pages in the wizard to create the new sync relationship.

Second method for deploying a data broker to your AWS account:

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

Create stack

Step 1: Specify template

Prerequisite - Prepare template

Prepare template
Every stack is based on a template. A template is a JSON or YAML file that contains configuration information about the AWS resources you want to include in the stack.

☒ Template is ready ☐ Use a sample template ☐ Create template in Designer

Specify template

A template is a JSON or YAML file that describes your stack's resources and properties.

Template source
Selecting a template generates an Amazon S3 URL, where it will be stored.

☒ Amazon S3 URL ☐ Upload a template file

Amazon S3 URL
https://s3.amazonaws.com/metadata.datafabric.io/data-mover-cf%2Fa28e3c1d-b2be-413e-8ed7-f033cadbf5d
Amazon S3 template URL

S3 URL: https://s3.amazonaws.com/metadata.datafabric.io/data-mover-cf%2Fa28e3c1d-b2be-413e-8ed7-f033cadbf5d [View in Designer](#)

[Cancel](#) [Next](#)

- The necessary information has already been filled in for you. Click Next.
- Select a VPC, subnet, and EC2 KeyPair. You do not need to configure anything under Assign a Public IP Address. Click Next.

Specify stack details

Stack name
aws-data-broker
Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).

Parameters

Parameters are defined in your template and allow you to input custom values when you create or update a stack.

NetApp Data Broker Placement

Which VPC should this be deployed to?
The list of your Virtual Private Cloud (VPC)

vpc-04c98d05a9964402 (172.30.0.0/24) (parallel-cluster-vpc)

Which subnet should this be deployed to?
The list of subnet IDs in your Virtual Private Cloud (VPC)

subnet-01b4070ff149dc4ab (172.30.0.0/26) (parallel-cluster-subnet-az-a)

NetApp Data Broker Security

EC2 KeyPair
Name of an existing EC2 KeyPair to enable SSH access to the NetApp Data Broker

parallel-cluster-key

Assign a public IP address?
Indicates whether the NetApp Data Broker instance should receive a public IP address

True

IAM role name (Optional)
The name of an existing IAM role to use with NetApp Data Broker

[Next](#)

- You do not need to configure anything under Configure Stack Options. Click Next.

Configure stack options

Tags
You can specify tags (key-value pairs) to apply to resources in your stack. You can add up to 50 unique tags for each stack. [Learn more](#)

Key: Value:

Permissions
Choose an IAM role to explicitly define how CloudFormation can create, modify, or delete resources in the stack. If you don't choose a role, CloudFormation uses permissions based on your user credentials. [Learn more](#)

IAM role - optional
Choose the IAM role for CloudFormation to use for all operations performed on the stack.

IAM role name:

Advanced options
You can set additional options for your stack, like notification options and a stack policy. [Learn more](#)

Stack policy
Defines the resources that you want to protect from unintentional updates during a stack update.

Rollback configuration
Specify alarms for CloudFormation to monitor when creating and updating the stack. If the operation breaches an alarm threshold, CloudFormation rolls it back. [Learn more](#)

d. You do not need to configure anything under Review.

Review aws

Step 1: Specify template

Template

Template URL
https://s3.amazonaws.com/metadata.datafabric.io/data-mover-cf%2F08dcfb5-98c5-42ee-8119-181bd25642a7

Stack description
Launch NetApp Data Broker for Cloud Sync

[Estimate cost](#)

Step 2: Specify stack details

Parameters (9)

Key	Value
AssociatePublicAddress	True
KeyPair	parallel-cluster-key
ProxyHost	-
ProxyPassword	-

e. Select the option that specifies that AWS CloudFormation might create IAM resources.

f. Click Create Stack.

Capabilities

The following resource(s) require capabilities: [AWS::IAM::InstanceProfile]

This template contains Identity and Access Management (IAM) resources that might provide entities access to make changes to your AWS account. Check that you want to create each of these resources and that they have the minimum required permissions. [Learn more](#)

☒ I acknowledge that AWS CloudFormation might create IAM resources.

- g. When the data broker is available, click Continue in Cloud Sync.
- h. Complete the pages in the wizard to create the new sync relationship.

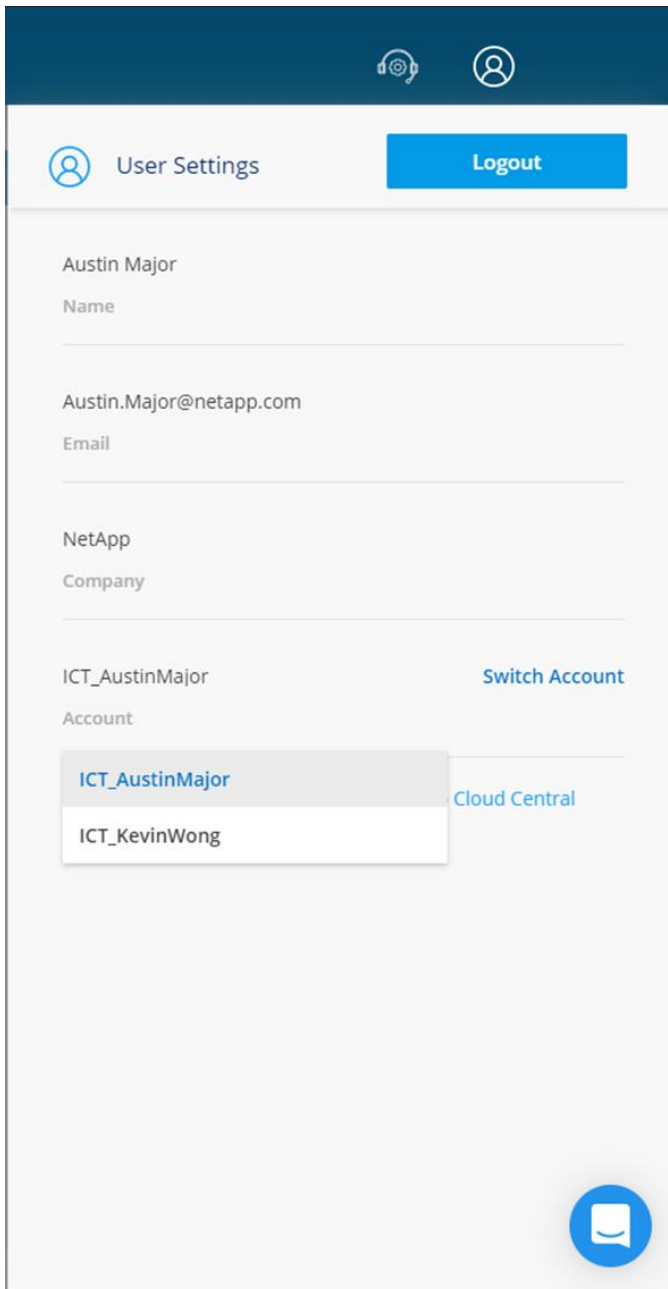
Share the data broker

If you have multiple users and want to share an instance of cloud sync with others, you might need to perform the following steps.

1. Go to <https://cloudmanager.netapp.com/>.
2. Go to Accounts > Manage Account > Users > Associate User > Enter User's Email > Select Role > Select Workspaces > Associate User.
 - a. To share NetApp cloud services such as Cloud Sync, provide only a Workspace Admin role.

The screenshot shows the 'Associate User' form in the NetApp Cloud Manager interface. The breadcrumb trail at the top reads: Manage Account: ICT_AustinMajor > Overview > Users > Workspaces > Service Connector > Subscriptions. The form title is 'Associate User'. A message states: 'To add a user to your NetApp Cloud Account, that user must already have signed up at NetApp Cloud Central. Enter the email address that they used when signing up with Cloud Central.' The form contains three main input fields: 'User's Email' with the value 'kevin.wong@netapp.com', 'Role' with a dropdown menu showing 'Workspace Admin', and 'Associate User to Workspaces' with a dropdown menu showing 'Workspace-1'. At the bottom are 'Cancel' and 'Associate User' buttons.

- b. Go to <https://cloudsync.netapp.com/>.
- c. On the top right, the “person” icon is the profile button. Click the icon, and then click Switch Account.



Uninstall the data broker

You can have only one data broker per Linux environment. It is directly tied to your cloud-based platform and cannot be managed by another user's Cloud Sync platform. The Installer is not idempotent and must be uninstalled.

1. Find a bash script in `/opt/netapp/uninstall.sh`.
2. Provide `chmod +x` permissions and execute the bash script.
3. When completed, check `/opt/netapp/*` and delete any remaining directories, if desired.
4. For more information, see [Uninstalling the data broker](#).

Conclusion

Using a basic setup of an NFS server with Windows 2016 and Red Hat 7.9, it is relatively easy to establish a Cloud Sync relationship between two remote NetApp E-Series systems. With Cloud Sync, you can now establish data mobility between NetApp E-Series systems and a number of data targets such as AWS S3, Azure, Google Cloud, NetApp ONTAP, and more.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- E-Series and SANtricity 11 Documentation Center
<https://docs.netapp.com/ess-11/index.jsp>
- E-Series and SANtricity documentation resources
<https://www.netapp.com/documentation/eseries-santricity/>
- NetApp Product Documentation
<https://docs.netapp.com/us-en/cloudsync/>

Version history

Version	Date	Document Version History
Version 1.0	January 2021	Initial release.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2021 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

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.

TR-4880-DEPLOY