Use BlueXP copy and sync

BlueXP copy and sync

Table of Contents

Use BlueXP copy and sync

Sync data between a source and target

Create sync relationships

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

Before you can create some types of sync relationships, you'll first need to create a working environment in BlueXP.

Create sync relationships for specific types of working environments

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

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

Steps

1. Create or discover the working environment.
   - Create an Amazon FSx for ONTAP working environment
   - 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
2. Select **Canvas**.
3. Select a working environment that matches any of the types listed above.
4. Select the action menu next to Sync.
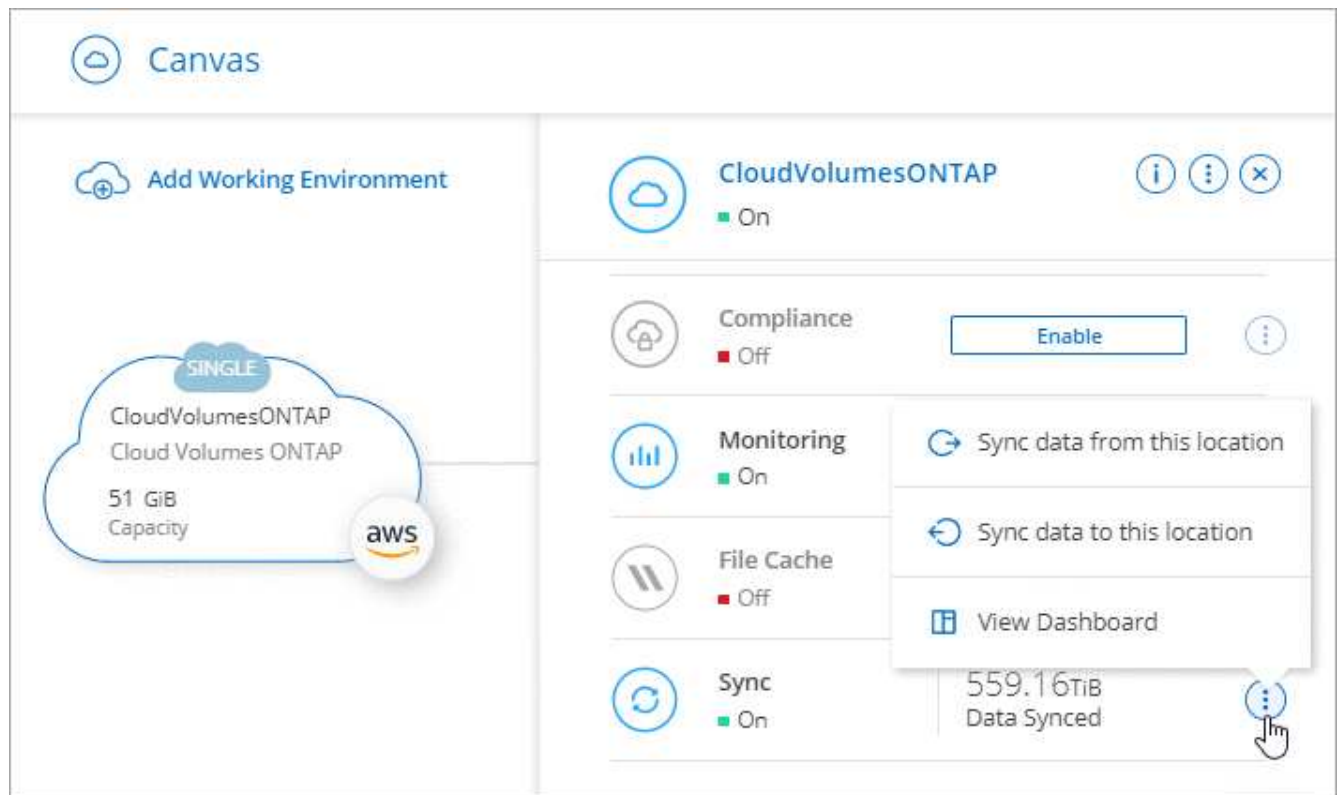
5. Select **Sync data from this location** or **Sync data to this location** and follow the prompts to set up the sync relationship.

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 BlueXP, 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.
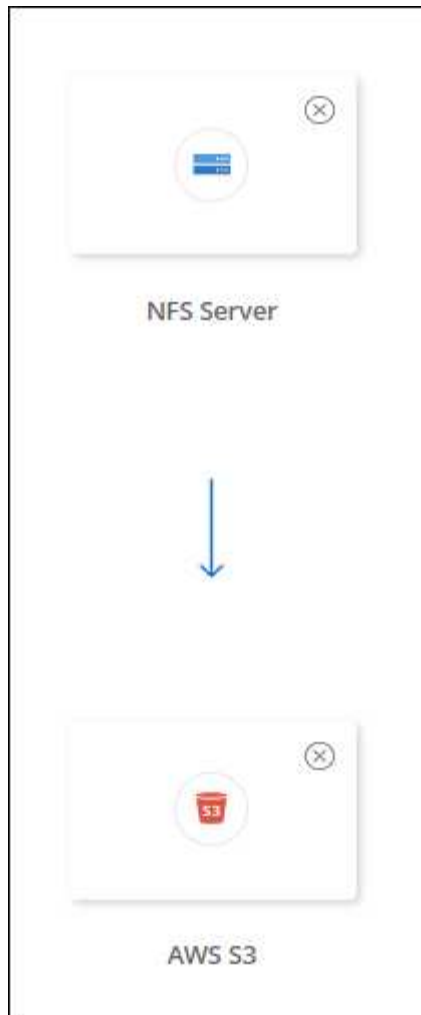
3. On the **NFS Server** page, enter the IP address or fully qualified domain name of the NFS server that you want to sync to AWS.

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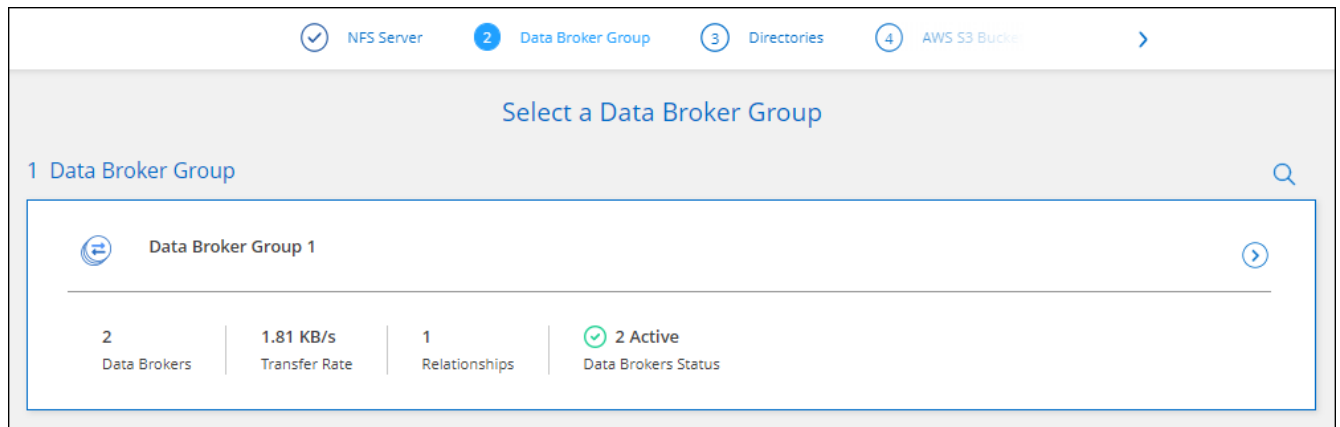
6. On the **Directories** page, select a top-level directory or subdirectory.

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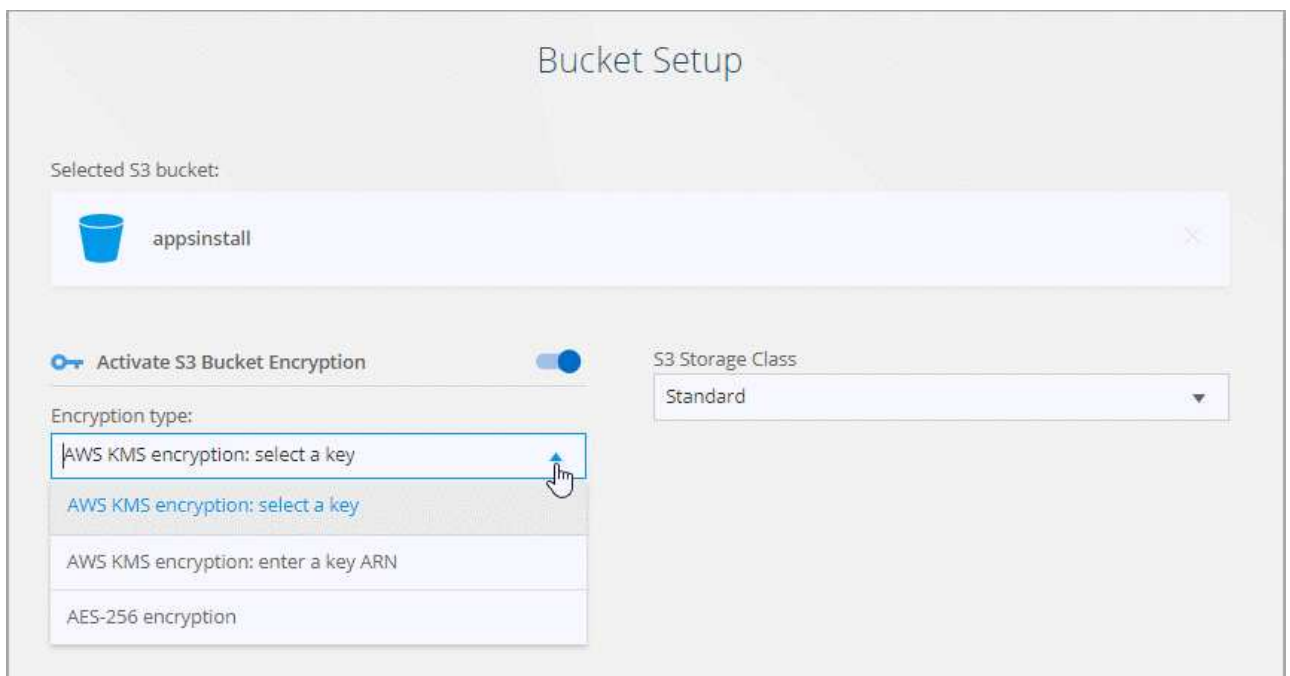
9. On the **Settings** page, define how source files and folders are synced and maintained in the target location:

   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 BlueXP 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 BlueXP copy and sync notifications in BlueXP'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 BlueXP copy and sync should retry to sync a file before skipping it.

   Continuous Sync

   After the initial data sync, BlueXP 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.

> (i) 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 BlueXP 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, BlueXP 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 BlueXP 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, BlueXP 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), BlueXP 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 BlueXP 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.

> (i) 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:

► https://docs.netapp.com/us-en/bluexp-copy-sync//media/video_file_extensions.mp4 *(video)*

> (i) 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.

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.

ⓠ 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

BlueXP copy and sync starts syncing data between the source and target.

Create sync relationships from BlueXP classification

BlueXP copy and sync is integrated with BlueXP classification. From within BlueXP classification, you can select the source files that you'd like to sync to a target location using BlueXP copy and sync.

After you initiate a data sync from BlueXP 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.

[Learn how to start a sync relationship from BlueXP classification](#).

Copying ACLs from SMB shares

BlueXP 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 BlueXP copy and sync to automatically copy ACLs](#)
- [Manually copy the ACLs between SMB shares](#)

Set up BlueXP 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. From BlueXP copy and sync, select **Create New Sync**.

2. 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**.

3. 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**.

4. 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.

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**.

Result

When syncing data, BlueXP 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.

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 BlueXP 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>

Syncing NFS data using data-in-flight encryption

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

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

How data-in-flight encryption works

Data-in-flight encryption encrypts NFS data when it's sent over the network between two data brokers. The following image shows a relationship between two NFS servers and two data brokers:
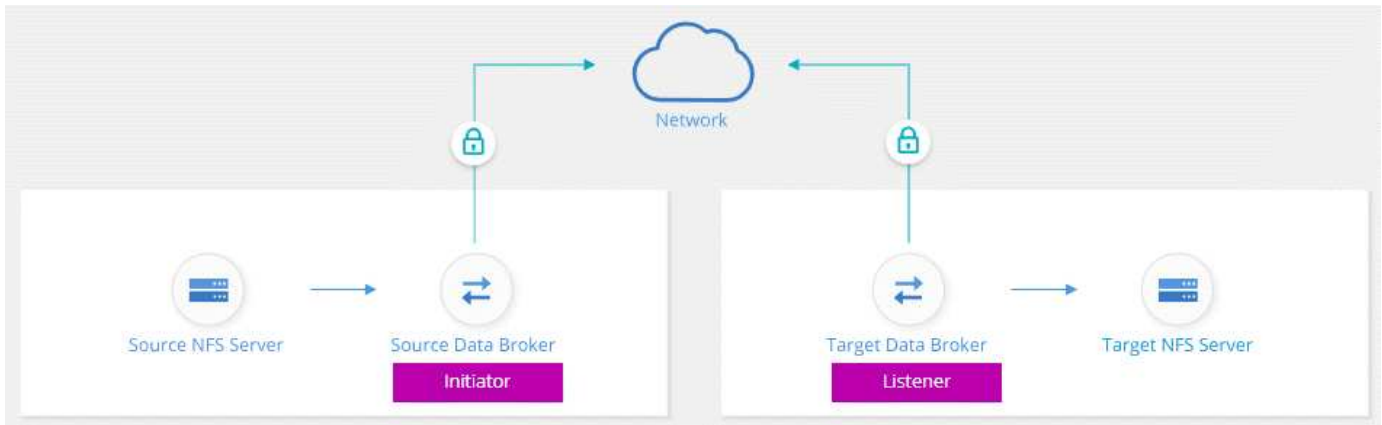
One data broker functions as the *initiator*. When it's time to sync data, it sends a connection request to the other data broker, which is the *listener*. That data broker listens for requests on port 443. You can use a different port, if needed, but be sure to check that the port is not in use by another service.

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

Here's how in-flight encryption works:

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

   If you prefer to sync data more frequently, you can change the schedule after you create the relationship.

Supported NFS versions

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

Proxy server limitation

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

What you'll need to get started

Be sure to have the following:

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

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

If you 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

Syncing NFS data using data-in-flight encryption

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

Steps

1. Select **Create New Sync**.

2. 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.

3. 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, BlueXP 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.
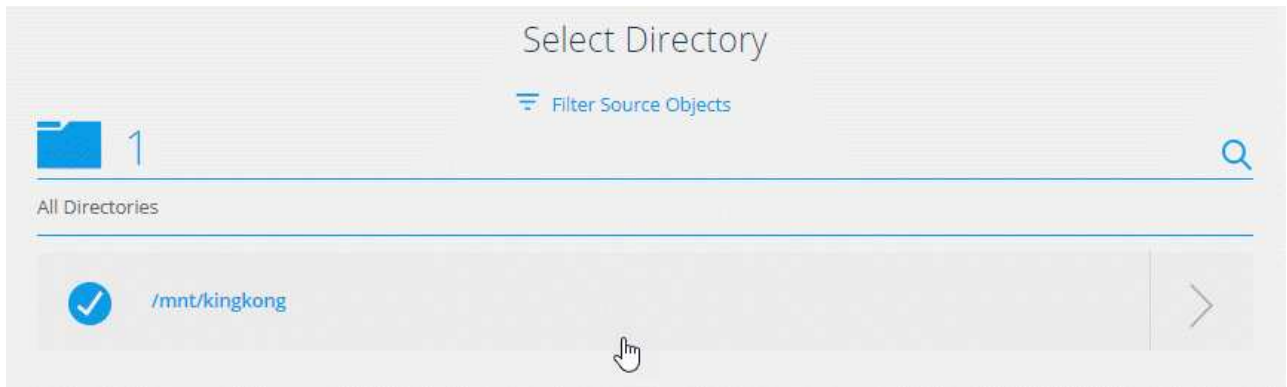
e. **Target NFS Server/Target Azure NetApp Files**: Choose the NFS version and then enter a new NFS target or select an existing server.

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.

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**.

Review Sync Relationship

Network

Source NFS Server → Source Data Broker *(Initiator)*

Target Data Broker → Target NFS Server *(Listener)*

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

BlueXP copy and sync starts creating the new sync relationship. When it's done, select **View in Dashboard** to view details about the new relationship.

Setting up a data broker group to use an external HashiCorp Vault

When you create a sync relationship that requires Amazon S3, Azure, or Google Cloud credentials, you need to specify those credentials through the BlueXP 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 BlueXP 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.

Preparing the vault

You'll need to provide BlueXP 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.

Preparing 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"
        }
    }
  }
}
```

Setting 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 BlueXP 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 BlueXP 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"
    }
    }
}
```

Paying for sync relationships after your free trial ends

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

You can 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 BlueXP copy and sync or the applicable website and **not** through a digital wallet.

Subscribing from AWS

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

Steps to pay-as-you-go

1. Select **Sync > Licensing**.

2. Select **AWS**

3. Select **Subscribe** and then select **Continue**.

4. Subscribe from the AWS Marketplace, and then log back in to the BlueXP copy and sync service to complete the registration.

    The following video shows the process:

    ► https://docs.netapp.com/us-en/bluexp-copy-sync//media/video_cloud_sync_registering.mp4 *(video)*

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**.

Subscribing from Azure

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

What you'll need

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

Steps

1. Select **Sync > Licensing**.

2. Select **Azure**.

3. Select **Subscribe** and then select **Continue**.

4. 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.

5. When deployment is complete, select the name of the SaaS resource in the notification pop-up.

6. Select **Configure Account** to return to BlueXP copy and sync.

    The following video shows the process:

    ► https://docs.netapp.com/us-en/bluexp-copy-sync//media/video_cloud_sync_registering_azure.mp4

(video)

Purchasing licenses from NetApp and adding them to BlueXP copy and sync

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

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. In BlueXP, select **Sync > 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:

      ▪ If your account was already added to BlueXP, 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**.

Updating a license

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

Steps

1. In BlueXP, select **Sync > Licensing**.

2. 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

BlueXP copy and sync updates the existing license with the new expiration date.

Managing sync relationships

You can manage sync relationships 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 press a button to immediately sync data between the source and target.

Steps

1. From the **Dashboard**, navigate to the sync relationship and select ⋮

2. Select **Sync Now** and then select **Sync** to confirm.

Result

BlueXP 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

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

Updatecredentials

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 BlueXP 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**.

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 BlueXP copy and sync will keep using the original domain that you provided.

Result

BlueXP 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 BlueXP copy and sync notifications in BlueXP's Notification Center. You can enable notifications for successful data syncs, failed data syncs, and canceled data syncs.

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 BlueXP copy and sync notifications in BlueXP'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.

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 BlueXP 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 BlueXP copy and sync notifications in BlueXP'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 BlueXP copy and sync should retry to sync a file before skipping it.

Compare By

Choose whether BlueXP 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, BlueXP 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 BlueXP 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 BlueXP 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.

> ℹ️ 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:

► https://docs.netapp.com/us-en/bluexp-copy-sync//media/video_file_extensions.mp4 *(video)*

> ℹ️ 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.

> ℹ️ 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

BlueXP 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

BlueXP 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

BlueXP copy and sync deletes the sync relationships.

Manage data broker groups

A data broker group 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. Select **Sync > Manage Data Brokers**.

2. Select **Add New Data Broker**.

3. 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. Select **Sync > Manage Data Brokers**.

2. Select the action menu and select **Add Data Broker**.

3. 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. Select **Sync > Manage Data Brokers**.
2. Select the action menu and select **Edit Group Name**.

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

Result

BlueXP 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. Select **Manage Data Brokers**.
2. Select the Settings icon for a data broker group.

3. 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 BlueXP copy and sync interface.
   ◦ If a data broker isn't running, it's configuration won't change because BlueXP 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. Select **Sync > Manage Data Brokers**.
2. Select to expand the list of data brokers in a group.
3. Select the action menu for a data broker and select **Move Data Broker**.

4. Create a new data broker group or select an existing data broker group.

5. Select **Move**.

Result

BlueXP 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 BlueXP 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. Select **Sync > Manage Data Brokers**.

2. Select to expand the list of data brokers in a group.

3. Select the action menu for a data broker and select **Edit Proxy Configuration**.

4. Specify details about the proxy: host name, port number, user name, and password.

5. Select **Update**.

Result

BlueXP 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.

BlueXP copy and sync provides the following details about a data broker:
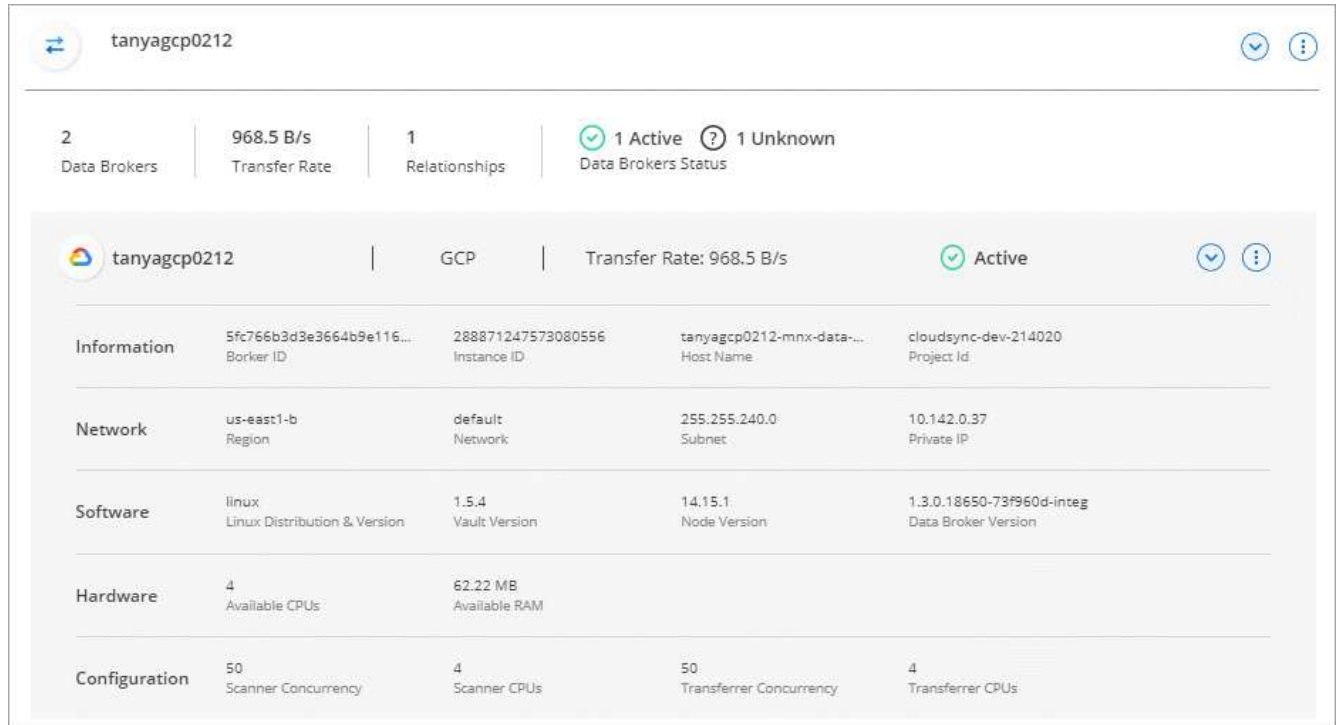
- Basic information: Instance ID, host name, etc.

- 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. Select **Sync > Manage Data Brokers**.

2. Select ⊗ to expand the list of data brokers in a group.

3. Select ⊗ to view details about a data broker.

Address issues with a data broker

BlueXP copy and sync displays a status for each data broker that can help you troubleshoot issues.

Steps

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

2. Hover over the ⓘ icon to see the failure reason.

3. Correct the issue.

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

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 BlueXP copy and sync's records. You'll need to manually delete the data broker and any additional cloud resources yourself.

Things you should know

- BlueXP 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. Select **Sync > Manage Data Brokers**.

2. Select ⊙ to expand the list of data brokers in a group.

3. Select the action menu for a data broker and select **Remove Data Broker**.

4. Select **Remove Data Broker**.

Result

BlueXP 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 BlueXP copy and sync.

Data brokers that BlueXP copy and sync removes are only deleted from BlueXP 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. Select **Sync > Manage Data Brokers**.

2. Select the action menu and select **Delete Group**.

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

Result

BlueXP copy and sync removes the data brokers and deletes the group.

Creating and viewing reports to tune your configuration

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

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

Creating reports

Each time that you create a report, BlueXP copy and sync scans the path and then compiles the details into a report.

Steps

1. Select **Sync > Reports**.

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

2. In the **Reports Actions** column, go to a specific path and select **Create**, or select the action menu and select **Create New**.

3. 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.

Downloading reports

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

Steps

1. Select **Sync > Reports**.
2. In the **Reports Actions** column, select the action menu and select **View**.
3. In the top right of the report, select the action menu and select **Download pdf**.
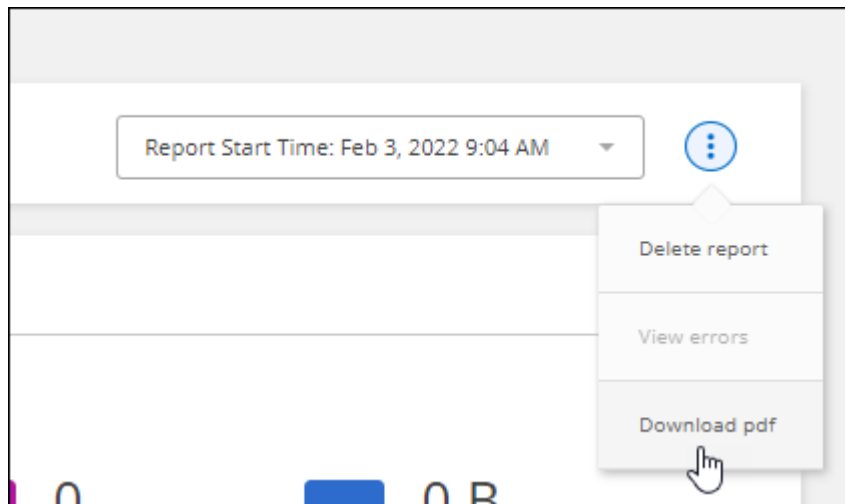
Viewing report errors

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

For example, a report might contain permission denied errors. This type of error can affect BlueXP 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. Select **Sync > Reports**.
2. In the **Errors** column, identify whether any errors are present in a report.
3. If errors are present, select the arrow next to the number of errors.

4. 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.

Deleting 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.

Uninstalling the data broker

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

Steps

1. Log in to the data broker host.

2. Change to the data broker directory: `/opt/netapp/databroker`

3. Run the following commands:

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

4. Press 'y' to confirm the uninstallation.