

Dropbox to S3/ADLS

Date published: 2021-04-06

Date modified: 2024-06-03

CLOUdera

Legal Notice

© Cloudera Inc. 2024. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

ReadyFlow overview: Dropbox to S3/ADLS.....	4
Prerequisites.....	4
List of required configuration parameters for the Dropbox to S3/ADLS ReadyFlow.....	8

ReadyFlow overview: Dropbox to S3/ADLS

You can use the Dropbox to S3/ADLS ReadyFlow to ingest data from Dropbox and write it to a destination in Amazon S3 or Azure Data Lake Service (ADLS).

This ReadyFlow consumes files from a Dropbox folder and writes them to a Cloudera managed destination S3 or ADLS location. For the source, specify the Dropbox App credentials (app key, app secret, access token, refresh token) and the ID or path of the Dropbox folder. You can choose whether to include subfolder files. For the destination, specify the S3 or ADLS storage location and path.

The ReadyFlow polls the Dropbox folder for new files (it performs a listing periodically). Failed S3 or ADLS write operations are retried automatically to handle transient issues. Define a KPI on the failure_WriteToS3/ADLS connection to monitor failed write operations.



Note: This ReadyFlow leverages Cloudera Public Cloud's centralized access control for cloud storage access. Make sure to either set up Ranger policies or an IDBroker mapping allowing your workload user access to the target S3 or ADLS location.

Dropbox to S3/ADLS ReadyFlow details	
Source	Dropbox
Source Format	Any
Destination	Cloudera managed Amazon S3 or ADLS
Destination Format	Same as source

Prerequisites

Learn how to collect the information you need to deploy the Dropbox to S3/ADLS ReadyFlow, and meet other prerequisites.

For your data ingest source

- You have a Dropbox account. If you do not have one, you can sign up for a free account [here](#).
- You have data stored in Dropbox that you want to move to an object store.
- You have created a Dropbox App.

- You have generated the following Dropbox credentials: App Key, App Secret, Access Token, and Refresh Token

Generating App Key and App Secret

1. Log in with your Dropbox account.
2. If you already have an app, go to [Dropbox Developers](#).
3. Click App Console and select your app.

On the app's info page you can find the App key and App secret credentials.

4. If you do not have any apps, go to [Dropbox Developers](#) and click Create apps.
5. On the next page select Scoped access and Full Dropbox as access type.
6. Provide a name for your app.
7. Accept the Dropbox API Terms and Conditions and click Create app.

Your app is created and you are taken to the app's info page. You can find the App key and App secret credentials on the Settings tab.

Setting required permissions for your app

1. Go to the Permissions tab of your app's info page.
2. Set the required permission for your app by enabling files.content.read.

With this permission, your app will be able to read the files in Dropbox.

3. Click Submit.



Note: In case you already have an Access Token and Refresh Token, those tokens have to be regenerated after the permission change.

For more information, see *Generating Access Token and Refresh Token*.

Generating Access Token and Refresh Token

1. Go to the following web page: https://www.dropbox.com/oauth2/authorize?token_access_type=offline&response_type=code&client_id=your_app_key
2. Click Next and then click Allow on the next page.

An access code is generated for you, and it is displayed on the next page.

3. Execute the following command from the terminal to fetch the access and refresh tokens. `curl https://api.dropbox.com/oauth2/token -d code=your_generated_access_code -d grant_type=authorization_code -u your_app_key:your_app_secret`



Important: Make sure you execute the curl command right after the access code generation, since the code expires very quickly. In case the curl command returns invalid grant error, generate a new access code.

The curl command results in a json file that contains the access_token and the refresh_token.

```
{
  "access_token": "sl.xxxxxxxxxxxx",
  "expires_in": 14400,
  "refresh_token": "xxxxxxx",
  "scope": "files.content.read files.metadata.read",
  "uid": "xxxxxx",
  "account_id": "dbid:xxxx"
}
```

For more information, see the [Dropbox Getting Started](#).

For Cloudera DataFlow

- You have enabled Cloudera DataFlow for an environment.

For information on how to enable Cloudera DataFlow for an environment, see [Enabling Cloudera DataFlow for an Environment](#).

- You have created a Machine User to use as the Cloudera Workload User.
- You have given the Cloudera Workload User the EnvironmentUser role.
 - From the Management Console, go to the environment for which Cloudera DataFlow is enabled.
 - From the Actions drop down, click Manage Access.
 - Identify the user you want to use as a Workload User.



Note:

The Cloudera Workload User can be a machine user or your own user name. It is best practice to create a dedicated Machine user for this.


- Give that user EnvironmentUser role.
- You have synchronized your user to the Cloudera Public Cloud environment that you enabled for Cloudera DataFlow.

For information on how to synchronize your user to FreeIPA, see [Performing User Sync](#).

- You have granted your Cloudera user the DFCatalogAdmin and DFFlowAdmin roles to enable your user to add the ReadyFlow to the Catalog and deploy the flow definition.
 - Give a user permission to add the ReadyFlow to the Catalog.
 - From the Management Console, click User Management.
 - Enter the name of the user or group you wish to authorize in the Search field.
 - Select the user or group from the list that displays.
 - Click Roles Update Roles .
 - From Update Roles, select DFCatalogAdmin and click Update.



Note: If the ReadyFlow is already in the Catalog, then you can give your user just the DFCatalogViewer role.

- Give your user or group permission to deploy flow definitions.
 - From the Management Console, click Environments to display the Environment List page.
 - Select the environment to which you want your user or group to deploy flow definitions.
 - Click Actions Manage Access to display the Environment Access page.
 - Enter the name of your user or group you wish to authorize in the Search field.
 - Select your user or group and click Update Roles.
 - Select DFFlowAdmin from the list of roles.
 - Click Update Roles.
- Give your user or group access to the Project where the ReadyFlow will be deployed.
 - Go to DataFlow Projects .
 - Select the project where you want to manage access rights and click  More Manage Access .
- Start typing the name of the user or group you want to add and select them from the list.
- Select the Resource Roles you want to grant.
- Click Update Roles.
- Click Synchronize Users.

For your ADLS data ingest target

- You have your ADLS container and path into which you want to ingest data.

- You have performed one of the following to configure access to your ADLS folder:
 - You have configured access to the ADLS folders with a RAZ enabled environment.

It is a best practice to enable RAZ to control access to your object store folders. This allows you to use your Cloudera Public Cloud credentials to access ADLS folders, increases auditability, and makes object store data ingest workflows portable across cloud providers.

1. Ensure that Fine-grained access control is enabled for your Cloudera DataFlow environment.
2. From the Ranger UI, navigate to the ADLS repository.
3. Create a policy to govern access to the ADLS container and path used in your ingest workflow. For example: adls-to-adls-avro-ingest



Tip: The Path field must begin with a forward slash (/).

4. Add the machine user that you have created for your ingest workflow to ingest the policy you just created.

For more information, see *Ranger policies for RAZ-enabled Azure environment*.

- You have configured access to ADLS folders using ID Broker mapping.

If your environment is not RAZ-enabled, you can configure access to ADLS folders using ID Broker mapping.

1. Access IDBroker mappings.
 - a. To access IDBroker mappings in your environment, click **Actions Manage Access**.
 - b. Choose the IDBroker Mappings tab where you can provide mappings for users or groups and click **Edit**.
2. Add your Cloudera Workload User and the corresponding Azure role that provides write access to your folder in ADLS to the Current Mappings section by clicking the blue + sign.



Note: You can get the Azure Managed Identity Resource ID from the Azure Portal by navigating to **Managed Identities Your Managed Identity Properties Resource ID**. The selected Azure MSI role must have a trust policy allowing IDBroker to assume this role.

3. Click **Save and Sync**.

For your S3 data ingest target

- You have your source S3 path and bucket.

- Perform one of the following to configure access to S3 buckets:

- You have configured access to S3 buckets with a RAZ enabled environment.

It is a best practice to enable RAZ to control access to your object store buckets. This allows you to use your Cloudera credentials to access S3 buckets, increases auditability, and makes object store data ingest workflows portable across cloud providers.

1. Ensure that Fine-grained access control is enabled for your Cloudera DataFlow environment.
2. From the Ranger UI, navigate to the S3 repository.
3. Create a policy to govern access to the S3 bucket and path used in your ingest workflow.



Tip:

The Path field must begin with a forward slash (/).

4. Add the machine user that you have created for your ingest workflow to the policy you just created.

For more information, see *Creating Ranger policy to use in RAZ-enabled AWS environment*.

- You have configured access to S3 buckets using ID Broker mapping.

If your environment is not RAZ-enabled, you can configure access to S3 buckets using ID Broker mapping.

1. Access IDBroker mappings.
 - a. To access IDBroker mappings in your environment, click **Actions Manage Access**.
 - b. Choose the IDBroker Mappings tab where you can provide mappings for users or groups and click **Edit**.
2. Add your Cloudera Workload User and the corresponding AWS role that provides write access to your folder in your S3 bucket to the **Current Mappings** section by clicking the blue + sign.



Note: You can get the AWS IAM role ARN from the Roles Summary page in AWS and can copy it into the IDBroker role field. The selected AWS IAM role must have a trust policy allowing IDBroker to assume this role.

3. Click **Save and Sync**.

Related Concepts

[List of required configuration parameters for the Dropbox to S3/ADLS ReadyFlow](#)

List of required configuration parameters for the Dropbox to S3/ADLS ReadyFlow

When deploying the Dropbox to S3/ADLS ReadyFlow, you have to provide the following parameters. Use the information you collected in *Prerequisites*.

Table 1: Dropbox to S3/ADLS ReadyFlow configuration parameters

Parameter name	Description
CDP Workload User	Specify the CDP machine user or workload username that you want to use to authenticate to the object stores. Ensure this user has the appropriate access rights to the object store locations in Ranger or IDBroker.
CDP Workload User Password	Specify the password of the CDP machine user or workload user you are using to authenticate against the object stores (via IDBroker).
Destination S3 or ADLS Path	Specify the name of the destination S3 or ADLS path you want to write to. Make sure that the path starts with "/".

Parameter name	Description
Destination S3 or ADLS Storage Location	Specify the name of the destination S3 bucket or ADLS container you want to write to. <ul style="list-style-type: none"> For S3, enter a value in the form: s3a:///***<i>Destination S3 Bucket</i>***] For ADLS, enter a value in the form: abfs:///***<i>Destination ADLS File System</i>***]@***<i>Destination ADLS Storage Account</i>***].dfs.core.windows.net
Dropbox App Access Token	Specify your Dropbox App Access Token.
Dropbox App Key	Specify your Dropbox App Key.
Dropbox App Refresh Token	Specify your Dropbox App Refresh Token.
Dropbox App Secret	Specify your Dropbox App Secret.
Dropbox Folder ID or Path	Specify the identifier or path of the Dropbox folder you want to read from. 'Folder' should match the following regular expression pattern: /. id:. Example folder identifier: id: odTIUvbpIEAAAAAAAAAAGGQ Example folder path: /Folder1/Folder2
Include Subfolder Files	Specify whether to include a list of files from subfolders. Set to "true" to include files from subfolders. Set to "false" to only list files from the specified Dropbox folder. The default value is "true".

Related Concepts[Prerequisites](#)[Related Information](#)[Deploying a ReadyFlow](#)