

S3 to OpenSearch

Date published: 2021-04-06

Date modified: 2025-07-17



Legal Notice

© Cloudera Inc. 2025. 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: S3 to OpenSearch [Technical Preview].....	4
Prerequisites.....	4
Required parameters.....	6

ReadyFlow: S3 to OpenSearch [Technical Preview]

You can use the S3 to OpenSearch [Technical Preview] Readyflow to consume PDF documents from S3, vectorize them using an OpenAI model and write the results to OpenSearch.

This ReadyFlow consumes PDF documents from a source S3 location, partitions the PDFs, chunks the data, vectorizes the data using an OpenAI embedding model, and stores the results in OpenSearch. The default OpenAI model is 'text-embedding-ada-002'. An OpenAI API key and an OpenSearch password are required to run this flow. Define a KPI on the failure_WriteToOpenSearch connection to monitor failed write operations.

**Note:**

This ReadyFlow is considered Technical Preview and is not designed for production use. The flow uses Python processors and must be run in NiFi 2.x.

The images within PDFs will be ignored.

S3 to OpenSearch [Technical Preview] ReadyFlow details	
Source	Cloudera Managed S3
Source Format	PDF
Destination	OpenSearch
Destination Format	Vector DB

Prerequisites

Learn how to collect the information you need to deploy the S3 to OpenSearch [Technical Preview] ReadyFlow, and meet other prerequisites.

For your data ingest source

- You have the source S3 bucket and path.

- You have performed one of the following to configure access to the source S3 bucket:
 - You have configured access to the S3 bucket 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 on cloud 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 Data Flow 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 the S3 bucket using ID Broker mapping.

If your environment is not RAZ-enabled, you can configure access to the S3 bucket 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**.

For Cloudera Data Flow

- You have enabled Cloudera Data Flow for an environment.

For information on how to enable Cloudera Data Flow for an environment, see [Enabling Cloudera Data Flow 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.

1. From the Management Console, go to the environment for which Cloudera Data Flow is enabled.
2. From the Actions drop down, click **Manage Access**.
3. 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.

4. Give that user EnvironmentUser role.
- You have synchronized your user to the Cloudera on cloud environment that you enabled for Cloudera Data Flow.

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.
 1. Give a user permission to add the ReadyFlow to the Catalog.
 - a. From the Management Console, click User Management.
 - b. Enter the name of the user or group you wish to authorize in the Search field.
 - c. Select the user or group from the list that displays.
 - d. Click Roles Update Roles .
 - e. 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.
 2. Give your user or group permission to deploy flow definitions.
 - a. From the Management Console, click Environments to display the Environment List page.
 - b. Select the environment to which you want your user or group to deploy flow definitions.
 - c. Click Actions Manage Access to display the Environment Access page.
 - d. Enter the name of your user or group you wish to authorize in the Search field.
 - e. Select your user or group and click Update Roles.
 - f. Select DFFlowAdmin from the list of roles.
 - g. Click Update Roles.
 3. Give your user or group access to the Project where the ReadyFlow will be deployed.
 - a. Go to DataFlow Projects .
 - b. Select the project where you want to manage access rights and click  More Manage Access .
 4. Start typing the name of the user or group you want to add and select them from the list.
 5. Select the Resource Roles you want to grant.
 6. Click Update Roles.
 7. Click Synchronize Users.

For your data ingest target

- You have your OpenAI API key.
- You have your OpenSearch instance and index into which you want to index data.
- You have your OpenSearch user name and password.

Related Concepts

[Required parameters](#)

Required parameters

When deploying the S3 to OpenSearch [Technical Preview] ReadyFlow, you have to provide the following parameters. Use the information you collected in *Prerequisites*.

Table 1: S3 to OpenSearch [Technical Preview] ReadyFlow configuration parameters

Parameter Name	Description
CDP Workload User	Specify the Cloudera machine user or workload username that you want to use to authenticate to the object stores (via IDBroker). Ensure this user has the appropriate access rights to the object store locations.
CDP Workload User Password	Specify the password of the Cloudera machine user or workload user you are using to authenticate against the object stores (via IDBroker).

Parameter Name	Description
CDPEnvironment	The Cloudera Environment configuration resources.
OpenAI API Key	Specify the API key used to authenticate to OpenAI.
OpenAI Model Name	Specify the OpenAI model name to use for embedding the data. The default model is 'text-embedding-ada-002'.
OpenSearch Document Text Field	The name of the field in the OpenSearch document where the text is stored. The default value is `text`.
OpenSearch Document Vector Field	Specify the name of the field in the OpenSearch document where the embeddings are stored. This field needs to be a 'knn_vector' type. The default value is `vector_embedding`.
OpenSearch Host URL	Specify the URL of the destination OpenSearch instance.
OpenSearch Index	Specify the name of the destination OpenSearch index.
OpenSearch Password	Specify the password to use for authenticating to the OpenSearch server.
OpenSearch Username	Specify the username to use for authenticating to the OpenSearch server.
Source S3 Bucket	Specify the name of the source S3 bucket you want to read from.
Source S3 Path	Specify the path within the source bucket where you want to read files from.

Related Concepts

[Prerequisites](#)

Related Information

[Deploying a ReadyFlow](#)