

PostgreSQL CDC to Kudu

Date published: 2021-04-06

Date modified: 2024-06-03



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: PostgreSQL CDC to Kudu.....	4
Prerequisites.....	4
List of required configuration parameters for the PostgreSQL CDC to Kudu ReadyFlow.....	6

ReadyFlow: PostgreSQL CDC to Kudu

You can use the PostgreSQL CDC to Kudu ReadyFlow to retrieve CDC events from a PostgreSQL source table and stream them to a Kudu destination table.

This ReadyFlow uses Debezium to retrieve CDC events (INSERT, UPDATE, DELETE) from a PostgreSQL source table to stream the events to a Kudu destination table. Failed Kudu write operations are retried automatically to handle transient issues. Define a KPI on the failure_ModifyKuduTable connection to monitor failed write operations.



Note:

This ReadyFlow does not support schema changes or primary key field updates. Make sure to assign the correct permissions for the Kudu destination table to the specified Cloudera Workload User.

PostgreSQL CDC to Kudu ReadyFlow details	
Source	PostgreSQL Table
Source Format	PostgreSQL Table
Destination	Kudu
Destination Format	Kudu Table

Prerequisites

Learn how to collect the information you need to deploy the PostgreSQL CDC to Kudu ReadyFlow, and meet other prerequisites.

For your data ingest source



Note: Do not change primary key field values in your source table after configuring the ReadyFlow. Doing so will cause the ReadyFlow to reject updates.



Note: You need to take care of field case sensitivity when defining source and destination table structure.

- You have obtained the PostgreSQL database server hostname and port.
- You have obtained the PostgreSQL schema name and table name. Take note of the table structure, specifically field case sensitivity.
- You have obtained a username and password to access the PostgreSQL table.
- You have performed the [PostgreSQL setup tasks required to run Debezium](#).

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 data ingest target

- You have a Real-Time Data Mart cluster running Kudu, Impala, and Hue in the same environment for which Cloudera DataFlow has been enabled.
- You have the Kudu Master hostnames.
 - From Management Console, click Data Hub Clusters.
 - Select the Real-Time Data Mart cluster to which you want to ingest data into.
 - Click the Hardware tab.
 - Copy the FQDN for each Kudu Master.

- You have created the Kudu table that you want to ingest data into.
 - Navigate to your Real Time Data Mart cluster and click Hue from the Services pane.
 - Click the Tables icon on the left pane.
 - Select the default database, and click + New to create a new table.
 - In the Type field, select Manually and click Next.
 - Provide the table Name, Format, Primary keys, and any partitions.
 - Click Submit. The newly created table displays in the default database Tables pane.
 - Check the Kudu UI **Tables** tab for the name of the table you created. You will need this table name when you use the Cloudera DataFlow Deployment wizard to deploy the ReadyFlow.
- You have assigned permissions via IDBroker or in Ranger to enable the Cloudera Workload User to access the Kudu table that you want to ingest data into.
 - From the base cluster on Cloudera Public Cloud, select Ranger.
 - Select your Real Time Data Mart cluster from the **Kudu** folder.
 - Click Add New Policy policy.
 - On the **Create Policy** page, enter the Kudu table name in the topic field.
 - Add the Cloudera Workload User in the Select User field.
 - Add the Insert and Select permissions in the Permissions field.
 - Click Save.

Related Concepts

[List of required configuration parameters for the PostgreSQL CDC to Kudu ReadyFlow](#)

List of required configuration parameters for the PostgreSQL CDC to Kudu ReadyFlow

When deploying the PostgreSQL CDC to Kudu ReadyFlow, you have to provide the following parameters. Use the information you collected in *Prerequisites*.

Table 1: PostgreSQL CDC to Kudu ReadyFlow configuration parameters

Parameter Name	Description
CDP Workload User	Specify the Cloudera machine user or workload user name that you want to use to authenticate to Kudu. Ensure this user has the appropriate access rights to the Kudu table.
CDP Workload User Password	Specify the password of the Cloudera machine user or workload user you are using to authenticate to Kudu.
Destination Database Table Name	Specify the destination database table name in the form: [database_name].[table_name]
Kudu Master Hosts	Specify the Kudu Master hostnames in a comma separated list.
Source Database Name	Specify the source database name.
Source Database Password	Specify the source database password.
Source Database Server Host Name	Specify the source database server host name.
Source Database Server Port	Specify the source database server port. The default value is 5432.
Source Database Table Name	Specify the source database table name.
Source Database User	Specify the source database user in the form: [schema_name].[table_name]

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