

Cloudera Runtime 1.0.0

## Working with Apache Hive metastore

Date published: 2020-11-30

Date modified: 2024-12-05

# CLOUDERA

<https://docs.cloudera.com/>

# 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

<b>HMS table storage.....</b>	<b>4</b>
<b>Using Hive data connectors to support external data sources.....</b>	<b>5</b>
Creating a remote database using data connector.....	6

## HMS table storage

You need to understand how the Hive metastore (HMS) stores Hive tables when you run a `CREATE TABLE` statement or migrate a table to Cloudera. The success or failure of the statement, the resulting table type, and the table location depends on a number of factors.

### HMS table transformations

The HMS includes the following Hive metadata about tables that you create:

- A table definition
- Column names
- Data types
- Comments in a central schema repository

When you use the `EXTERNAL` keyword in the `CREATE TABLE` statement, HMS stores the table as an external table. When you omit the `EXTERNAL` keyword and create a managed table, or ingest a managed table, HMS might translate the table into an external table or the table creation can fail, depending on the table properties. An important table property that affects table transformations is the ACID or Non-ACID table type:

#### Non-ACID

Table properties do not contain any ACID related properties set to true. For example, the table does not contain such properties `transactional=true` or `insert_only=true`.

#### ACID

Table properties do contain one or more ACID properties set to true.

#### Full ACID

Table properties contain `transactional=true` but not `insert_only=true`

#### Insert-only ACID

Table properties contain `insert_only=true`.

The following matrix shows the table type and whether or not the location property is supported.

ACID	Managed	Location Property	Comments	Action
Non-ACID	Yes	Yes	Migrated to Cloudera, for example from an HDP or CDH cluster.	Table stored as external
Non-ACID	Yes	No	Table location is null	Table stored in subdirectory of external warehouse*

\* `metastore.warehouse.external.dir`

HMS detects type of client for interacting with HMS, for example Hive or Spark, and compares the capabilities of the client with the table requirement. HMS performs the following actions, depending on the results of the comparison:

Table requirement	Client meets requirements	Managed Table	ACID table type	Action
Client can write to any type of ACID table	No	Yes	Yes	CREATE TABLE fails
Client can write to full ACID table	No	Yes	<code>insert_only=true</code>	CREATE TABLE fails
Client can write to insert-only ACID table	No	Yes	<code>insert_only=true</code>	CREATE TABLE fails

If, for example, a Spark client does not have the capabilities required, the following type of error message appears:

```
Spark has no access to table `mytable`. Clients can access this table only if
they have the following capabilities: CONNECTORREAD, HIVEFULLACIDREAD, HIVE
FULLACIDWRITE,
HIVEMANAGESTATS, HIVECACHEINVALIDATE, . . .
```

## Using Hive data connectors to support external data sources

You can use Hive data connectors to map databases present in external data sources to a local Hive Metastore (HMS). The external data sources can be of different types, such as MySQL, PostgreSQL, Oracle, Redshift, Derby, or other HMS instances. You can create external tables to represent the data, and then query the tables.

### What is a data connector

Data connectors in Hive are objects that are defined using a set of properties, such as hostname, user credentials, and type that are required to connect Hive to the external data source. You define a connector and then use the same connector to map multiple databases from the remote data source to the local HMS.

### Benefits of using data connectors

Currently, you can use a JDBCStorageHandler to connect Hive to external data sources. Users define a table in the HMS for which data resides in a remote JDBC data store. The Hive table's metadata is persisted locally in the HMS. When HiveServer (HS2) runs a query against this table, data is retrieved from the remote JDBC table. However, there are certain limitations which make it difficult to map databases having a large number of tables.

- Each remote table in the remote data source has to be individually mapped to a local Hive table. It is cumbersome when you have to map an entire database having a lot of tables.
- New tables created in the remote data source are not automatically visible and should be manually mapped in Hive.
- The metadata for the mapped table is static. It does not track changes to the remote table. If the remote table is modified (added columns or dropped columns), the mapped Hive table has to be dropped and recreated. This is not feasible for organizations where tables are constantly changing.

With data connectors, you can map a database or schema in the remote data source to a Hive database. Such databases are referred to as 'Remote' databases in Hive. The benefits of creating remote databases using data connectors are as follows:

- Tables within a mapped database are automatically visible from Hive. The metadata for these tables are not persisted in the HMS. They are retrieved at runtime through an active connector.
- New tables created are automatically visible in Hive.
- The columns and their data types are mapped to compatible Hive data types during runtime. Therefore, metadata changes to tables in the remote datasource are immediately visible in hive.
- The same connector can be used to map another database to a Hive database. All the connection information is shared between these remote databases.

To create a remote database, you must first define a data connector with the required properties and then use this connector to create and map a remote database in an external data source to Hive.

## Creating a remote database using data connector

Learn how to create a Hive data connector, which you can then use to create and map a remote database to Hive. The remote database can reside in an external data source, such as MySQL, PostgreSQL, Oracle, Redshift, Derby, or other HMS instances.

### Before you begin

You must ensure that the `hive.security.temporary.authorization.for.data.connectors` property is set to "true".

In the Hive virtual warehouse details page, go to **CONFIGURATIONS Hiveserver2 hive-site** configuration file and click **+** to add this property as a custom configuration.

### Procedure

1. Create a data connector using the following syntax:

```
CREATE CONNECTOR [IF NOT EXISTS] connector_name
  [TYPE datasource_type]
  [URL datasource_url]
  [COMMENT connector_comment]
  [WITH DC PROPERTIES (property_name=property_value, ...)];
```

If you are using a cleartext password:

```
CREATE CONNECTOR postgres_local
  TYPE 'postgres'
  URL 'jdbc:postgresql://localhost:5432'
  WITH DC PROPERTIES (
    "hive.sql.dbc.username"="postgres",
    "hive.sql.dbc.password"="postgres"
  );
```

If you are using a Java keystore instead of a cleartext password:

```
CREATE CONNECTOR postgres_local_ks
  TYPE 'postgres'
  URL 'jdbc:postgresql://localhost:5432'
  WITH DC PROPERTIES (
    "hive.sql.dbc.username"="postgres",
    "hive.sql.dbc.password.keystore"="jceks://app/local/hive/secrets
.jceks",
    "hive.sql.dbc.password.key"="postgres.credential"
  );
```

2. Create a remote database in Hive using the data connector that you created in the previous step.

```
CREATE [REMOTE] (DATABASE) [IF NOT EXISTS] database_name
  [COMMENT database_comment]
  [USING connector_name]
  [WITH DBPROPERTIES (property_name=property_value, ...)];
```

```
CREATE REMOTE DATABASE postgres_hive_test
  USING postgres_local
  WITH DBPROPERTIES ("connector.remoteDbName"="remote_db_test");
```

This statement maps a remote database named "remote\_db\_test" to a Hive database named "postgres\_hive\_test" in Hive.

3. You can now use the tables that are available in the remote database.

```
USE remote_db_test;  
SHOW TABLES;  
  
DESCRIBE [formatted] <table name>;  
  
SELECT <col1> FROM <table name> WHERE <filter1> AND <filter2>;
```



**Important:**

- The metadata for these tables are not persisted in HMS.
- CREATE, ALTER, and DROP table DDLs are currently not supported in remote databases.