

Cloudera Runtime 7.2.12

## Apache Hive Overview

Date published: 2019-08-21

Date modified: 2021-10-20

# CLOUDERA

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

# 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

<b>Apache Hive features.....</b>	<b>4</b>
<b>Hive low-latency analytical processing.....</b>	<b>5</b>
<b>Hive unsupported interfaces and features in public clouds.....</b>	<b>5</b>
<b>Apache Hive 3 architectural overview.....</b>	<b>6</b>
<b>Apache Hive content roadmap.....</b>	<b>7</b>

## Apache Hive features

Major changes to Apache Hive 2.x improve Apache Hive 3.x transactions and security. Knowing the major differences between these versions is critical for SQL users, including those who use Apache Spark and Apache Impala.

Hive is a data warehouse system for summarizing, querying, and analyzing huge, disparate data sets.

### ACID transaction processing

Hive 3 tables are ACID (Atomicity, Consistency, Isolation, and Durability)-compliant. Hive 3 write and read operations improve the performance on transactional tables in comparison with Hive 2. Atomic operations include simple writes and inserts, writes to multiple partitions, and multiple inserts in a single SELECT statement. A read operation is not affected by changes that occur during the operation. You can insert or delete data, and it remains consistent throughout software and hardware crashes. Creation and maintenance of Hive tables is simplified because there is no longer any need to bucket tables.

### Materialized views

Because multiple queries frequently need the same intermediate roll up or joined table, you can avoid costly, repetitious query portion sharing, by precomputing and caching intermediate tables into views.

### Query results cache

Hive filters and caches similar or identical queries. Hive does not recompute the data that has not changed. Caching repetitive queries can reduce the load substantially when hundreds or thousands of users of BI tools and web services query Hive.

### Scheduled Queries

Using SQL statements, you can schedule Hive queries to run on a recurring basis, monitor query progress, temporarily ignore a query schedule, and limit the number running in parallel. You can use scheduled queries to start compaction and periodically rebuild materialized views, for example.

### Spark integration with Hive

Spark and Hive tables interoperate using the Hive Warehouse Connector and Spark Direct Reader to access ACID managed tables. You can access external tables from Spark directly using SparkSQL.

You do not need HWC to read or write Hive external tables. Spark users just read from or write to Hive directly. You can read Hive external tables in ORC or Parquet formats. You can write Hive external tables in ORC format only.

### Security improvements

Apache Ranger secures Hive data. To meet demands for concurrency improvements, ACID support, render security, and other features, Hive tightly controls the location of the warehouse on a file system, or object store, and memory resources.

With Apache Ranger and Apache Hive ACID support, your organization will be ready to support and implement GDPR (General Data Protection Regulation).

### Connection Pooling

Hive supports HakariCP JDBC connection pooling.

### Unsupported features

CDP does not support the following features that were available in HDP and CDH platforms:

- CREATE TABLE that specifies a managed table location

Do not use the LOCATION clause to create a managed table. Hive assigns a default location in the warehouse to managed tables.

- CREATE INDEX

Hive builds and stores indexes in ORC or Parquet within the main table, instead of a different table, automatically. Set `hive.optimize.index.filter` to enable use (not recommended--use materialized views instead). Existing indexes are preserved and migrated in Parquet or ORC to CDP during upgrade.

### Related Information

[Blog: Enabling high-speed Spark direct reader for Apache Hive ACID tables](#)

## Hive low-latency analytical processing

CDP Public Cloud supports low-latency analytical processing (LLAP) of Hive queries. Using LLAP in the CDP Data Warehouse service, you can tune your data warehouse infrastructure, components, and client connection parameters to improve the performance and relevance of business intelligence and other applications.

Increasingly, enterprises want to run SQL workloads that return faster results than batch processing can provide. These enterprises often want data analytics applications to support interactive queries. Low-latency analytical processing (LLAP) can improve the performance of interactive queries. A Hive interactive query that runs on the CDP Public Cloud meets low-latency, variably gauged benchmarks to which Hive LLAP responds in 15 seconds or less. LLAP enables application development and IT infrastructure to run queries that return real-time or near-real-time results.

LLAP caches data for multiple queries and this capability does not support user impersonation.

## Hive unsupported interfaces and features in public clouds

You need to know the interfaces available in HDP or CDH platforms that are not supported in CDP.

The following interfaces are not supported:

- Hcat CLI (however, HCatalog is supported)
- Hive CLI (replaced by Beeline)
- Hive View UI feature in Ambari
- MapReduce execution engine (replaced by LLAP)
- Pig
- Spark execution engine
- Spark thrift server

Spark and Hive tables interoperate using the Hive Warehouse Connector.

- SQL Standard Authorization
- Tez View UI feature in Ambari
- WebHCat

You can use Hue in lieu of Hive View.

## Unsupported Features

CDP does not support the following features that were available in HDP and CDH platforms:

- CREATE TABLE that specifies a managed table location

Do not use the LOCATION clause to create a managed table. Hive assigns a default location in the warehouse to managed tables. That default location is configured in Hive using the `hive.metastore.warehouse.dir` configuration property, but can be overridden for the database by setting the `CREATE DATABASE MANAGEDLOCATION` parameter.

- CREATE INDEX and related index commands were removed in Hive 3, and consequently are not supported in CDP.

In CDP, you use the Hive 3 default ORC columnar file formats to achieve the performance benefits of indexing. Materialized Views with automatic query rewriting also improves performance. Indexes migrated to CDP are preserved but render any Hive tables with an undroppable index. To drop the index, google the Known Issue for CDPD-23041.

- Hive metastore (HMS) high availability (HA) load balancing in CDH

You need to set up HMS HA as described in the documentation.

- Local or Embedded Hive metastore server

CDP does not support the use of a local or embedded Hive metastore setup.

## Unsupported Connector Use

CDP does not support the Sqoop exports using the Hadoop jar command (the Java API) that Teradata documents. For more information, see [Migrating data using Sqoop](#).

# Apache Hive 3 architectural overview

Understanding Apache Hive 3 major design features, such as default ACID transaction processing, can help you use Hive to address the growing needs of enterprise data warehouse systems.

## Data storage and access control

One of the major architectural changes to support Hive 3 design gives Hive much more control over metadata memory resources and the file system, or object store. The following architectural changes from Hive 2 to Hive 3 provide improved security:

- Tightly controlled file system and computer memory resources, replacing flexible boundaries: Definitive boundaries increase predictability. Greater file system control improves security.
- Optimized workloads in shared files and YARN containers

Hive 3 is optimized for object stores in the following ways:

- Hive uses ACID to determine which files to read rather than relying on the storage system.
- In Hive 3, file movement is reduced from that in Hive 2.
- Hive caches metadata and data aggressively to reduce file system operations

The major authorization model for Hive is Ranger. Hive enforces access controls specified in Ranger. This model offers stronger security than other security schemes and more flexibility in managing policies.

This model permits only Hive to access the Hive warehouse.

## Transaction processing

You can deploy new Hive application types by taking advantage of the following transaction processing characteristics:

- Mature versions of ACID transaction processing:
  - ACID tables are the default table type.
  - ACID enabled by default causes no performance or operational overload.
- Simplified application development, operations with strong transactional guarantees, and simple semantics for SQL commands
  - You do not need to bucket ACID tables.
- Materialized view rewrites
- Automatic query cache
- Advanced optimizations

### Hive client changes

You can use the thin client Beeline for querying Hive from the command line. You can run Hive administrative commands from the command line. Beeline uses a JDBC connection to Hive to run commands. Hive parses, compiles, and runs operations. Beeline supports many of the command-line options that Hive CLI supported. Beeline does not support `hive -e set key=value` to configure the Hive Metastore.

You enter supported Hive CLI commands by invoking Beeline using the `hive` keyword, command option, and command. For example, `hive -e set`. Using Beeline instead of the thick client Hive CLI, which is no longer supported, has several advantages, including low overhead. Beeline does not use the entire Hive code base. A small number of daemons required to run queries simplifies monitoring and debugging.

Hive enforces allowlist and denylist settings that you can change using SET commands. Using the denylist, you can restrict memory configuration changes to prevent instability. Different Hive instances with different allowlists and denylists to establish different levels of stability.

### Apache Hive Metastore sharing

Hive, Impala, and other components can share a remote Hive metastore.

### Spark integration

Spark and Hive tables interoperate using the Hive Warehouse Connector.

You can use the Hive Warehouse Connector (HWC) to access Hive managed tables from Spark. HWC is specifically designed to access managed Hive tables, and supports writing to tables in ORC format only.

You do not need HWC to read from or write to Hive external tables. Spark uses native Spark to read external tables.

### Query execution of batch and interactive workloads

You can connect to Hive using a JDBC command-line tool, such as Beeline, or using an JDBC/ODBC driver with a BI tool, such as Tableau. You configure the settings file for each instance to perform either batch or interactive processing.

### Related Information

[Blog: Enabling high-speed Spark direct reader for Apache Hive ACID tables](#)

## Apache Hive content roadmap

The content roadmap provides links to the available content resources for Apache Hive.

**Table 1: Apache Hive Content roadmap**

Task	Resources	Source	Description
Understanding	<a href="#">Presentations and Papers about Hive</a>	Apache wiki	Contains meeting notes, presentations, and whitepapers from the Apache community.
Getting Started	<a href="#">Hive Tutorial</a>	Apache wiki	Provides a basic overview of Apache Hive and contains some examples on working with tables, loading data, and querying and inserting data.
Developing	<a href="#">Materialized Views</a>	Apache wiki	Covers accelerating query processing in data warehouses by pre-computing summaries using materialized views.
	<a href="#">JdbcStorageHandler</a>	Apache wiki	Describes how to read from a JDBC data source in Hive.
	<a href="#">Hive transactions</a>	Apache wiki	Describes ACID operations in Hive.
	<a href="#">Hive Streaming API</a>	Apache wiki	Explains how to use an API for pumping data continuously into Hive using clients such as NiFi and Flume.
	<a href="#">Hive Operators and Functions</a>	Apache wiki	Describes the Language Manual UDF.
	<a href="#">Beeline: HiveServer2 Client</a>	Apache wiki	Describes how to use the Beeline client.
Reference	<a href="#">SQL Language Manual</a>	Apache wiki	Language reference documentation available in the Apache wiki.
Contributing	<a href="#">Hive Developer FAQ</a>	Apache wiki	Resources available if you want to contribute to the Apache community.
	<a href="#">How to Contribute</a>	Apache wiki	
	<a href="#">Hive Developer Guide</a>	Apache wiki	
	<a href="#">Plug-in Developer Kit</a>	Apache wiki	
	<a href="#">Unit Test Parallel Execution</a>	Apache wiki	
	<a href="#">Hive Architecture Overview</a>	Apache wiki	
	<a href="#">Hive Design Docs</a>	Apache wiki	
	<a href="#">Project Bylaws</a>	Apache wiki	