

Apache Hive metastore overview

Date published: 2019-08-21

Date modified: 2021-10-20

CLOUDERA

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

Introduction to Hive metastore.....	4
Apache Hive storage in public clouds.....	4

Introduction to Hive metastore

Hive metastore (HMS) is a service that stores metadata related to Apache Hive and other services, in a backend RDBMS, such as MySQL or PostgreSQL. Impala, Spark, Hive, and other services share the metastore. The connections to and from HMS include HiveServer, Ranger, and the NameNode that represents HDFS.

Beeline, Hue, JDBC, and Impala shell clients make requests through thrift or JDBC to HiveServer. The HiveServer instance reads/writes data to HMS. By default, redundant HMS operate in active/active mode. The physical data resides in a backend RDBMS, one for HMS. All HMS instances use the same backend database. A separate RDBMS supports the security service, Ranger for example. All connections are routed to a single RDBMS service at any given time. HMS talks to the NameNode over thrift and functions as a client to HDFS.



HMS connects directly to Ranger and the NameNode (HDFS), and so does HiveServer, but this is not shown in the diagram for simplicity. One or more HMS instances on the backend can talk to other services, such as Ranger.

Related Information

[Working with Apache Hive Metastore](#)

Apache Hive storage in public clouds

Knowing the storage locations of Apache Hive data and metadata helps you troubleshoot problems and replicate data.

On public clouds, the Hive warehouse that stores Hive data is located in an object store, such as S3, by default. In the cloud, Hive uses HDFS merely for storing temporary files. The Hive Metastore (HMS) stores the schema for Hive tables. HMS uses a pre-installed MySQL database. You perform little, or no, configuration of HMS in the cloud.