

Installing Hive ODBC Drivers

This section describes how to install the Hive ODBC driver on a Windows or Mac OS X client machine. Make sure to install the ODBC driver version recommended by your chosen BI vendor and version.



Note: If you are connecting to AtScale from Tableau, you can alternatively use the PostgreSQL JDBC driver. For more information, see [Installing PostgreSQL JDBC Drivers](#).

About ODBC Drivers



Whether Hadoop or Google BigQuery is your data warehouse, you must use one of the Hive ODBC drivers if your client application is ODBC-compliant.

Open Database Connectivity (ODBC) is an API standard that allows client applications to connect to a relational database and issue SQL queries using a standard interface. This section explains how to connect to AtScale from ODBC-compliant client applications.

For applications that send SQL queries, AtScale uses the same protocols and drivers as a remote HiveServer2 instance.

The Hive ODBC Driver is a software library that implements the ODBC API standard for the Hive data warehouse system. An AtScale cube looks like a HiveServer2 data source to ODBC-compliant applications, and requires the appropriate HiveServer2 database driver.

AtScale supports the following ODBC drivers and versions:

- ▲ Cloudera Hive ODBC Driver 2.5.9 to 2.5.16
- ▲ Hortonworks Hive ODBC Driver 2.1.2
- ▲ MapR Hive ODBC Driver 2.1.8



Check your BI vendor's documentation to confirm the minimum Hive ODBC driver supported for the release you are on. Some BI tools, such as Tableau, require that you also upgrade your ODBC drivers when you upgrade to the newest version of the desktop client.

OLAP Data Source Descriptors For ODBC-Compliant BI Applications



The Cloudera, Hortonworks, and MapR drivers use different Hive protocols. You must change the value for `thrift.client.protocol` property in the Settings page to `cloudera`, `hortonworks`, or `mapr`; no need to restart the engine. You could also use the `atscalesimba` value for the Simba Hive ODBC driver by AtScale. For

details on how to configure this property, see [Changing Engine Settings](#).

AtScale should work with any ODBC-compliant business intelligence (BI) application that supports the appropriate Hive ODBC driver. To ODBC-compliant applications, an AtScale cube looks like one wide relational database table - with one column for each measure and dimension attribute defined in your AtScale cube.

AtScale provides additional OLAP support for popular ODBC-compliant BI applications such as Tableau by generating a data source descriptor for each published cube. These data source files not only have the cube connection details, but also the cube metadata modeled in AtScale (measures, dimensions, hierarchies, folders, drill-through column sets, descriptions, etc.).

- ▲ [Installing the Hive ODBC Driver on Mac Clients](#)
- ▲ [Installing the Hive ODBC Driver on Windows Clients](#)
- ▲ [Configure an ODBC DSN on Windows Clients](#)