

# Cloudera ODBC Driver for Apache Hive Version 2.5.15



## Important Notice

© 2010-2015 Cloudera, Inc. All rights reserved.

Cloudera, the Cloudera logo, Cloudera Impala, Impala, and any other product or service names or slogans contained in this document, except as otherwise disclaimed, are trademarks of Cloudera and its suppliers or licensors, and may not be copied, imitated or used, in whole or in part, without the prior written permission of Cloudera or the applicable trademark holder.

Hadoop and the Hadoop elephant logo are trademarks of the Apache Software Foundation. All other trademarks, registered trademarks, product names and company names or logos mentioned in this document are the property of their respective owners. Reference to any products, services, processes or other information, by trade name, trademark, manufacturer, supplier or otherwise does not constitute or imply endorsement, sponsorship or recommendation thereof by us.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Cloudera.

Cloudera may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Cloudera, the furnishing of this document does not give you any license to these patents, trademarks copyrights, or other intellectual property.

The information in this document is subject to change without notice. Cloudera shall not be liable for any damages resulting from technical errors or omissions which may be present in this document, or from use of this document.

**Cloudera, Inc.**  
**1001 Page Mill Road, Building 2**  
**Palo Alto, CA 94304-1008**  
[info@cloudera.com](mailto:info@cloudera.com)  
**US: 1-888-789-1488**  
**Intl: 1-650-843-0595**  
[www.cloudera.com](http://www.cloudera.com)

## Release Information

Version: 2.5.15

Date: May 12, 2015

# Table of Contents

- INTRODUCTION** ..... 5
- WINDOWS DRIVER** ..... 5
  - SYSTEM REQUIREMENTS ..... 5
  - INSTALLING THE DRIVER ..... 6
  - VERIFYING THE VERSION NUMBER ..... 6
  - CREATING A DATA SOURCE NAME ..... 6
  - CONFIGURING A DSN-LESS CONNECTION ..... 9
  - CONFIGURING AUTHENTICATION ..... 11
  - CONFIGURING ADVANCED OPTIONS ..... 13
  - CONFIGURING SERVER-SIDE PROPERTIES ..... 14
  - CONFIGURING HTTP OPTIONS ..... 15
  - CONFIGURING SSL VERIFICATION ..... 16
  - CONFIGURING THE TEMPORARY TABLE FEATURE ..... 17
- LINUX DRIVER** ..... 18
  - SYSTEM REQUIREMENTS ..... 18
  - INSTALLING THE DRIVER ..... 18
  - VERIFYING THE VERSION NUMBER ..... 19
  - SETTING THE LD\_LIBRARY\_PATH ENVIRONMENT VARIABLE ..... 20
- MAC OS X DRIVER** ..... 20
  - SYSTEM REQUIREMENTS ..... 20
  - INSTALLING THE DRIVER ..... 20
  - VERIFYING THE VERSION NUMBER ..... 21
  - SETTING THE DYLD\_LIBRARY\_PATH ENVIRONMENT VARIABLE ..... 21
- AIX DRIVER** ..... 22
  - SYSTEM REQUIREMENTS ..... 22
  - INSTALLING THE DRIVER ..... 22
  - VERIFYING THE VERSION NUMBER ..... 23
  - SETTING THE LD\_LIBRARY\_PATH ENVIRONMENT VARIABLE ..... 23
- DEBIAN DRIVER** ..... 23
  - SYSTEM REQUIREMENTS ..... 23
  - INSTALLING THE DRIVER ..... 24
  - VERIFYING THE VERSION NUMBER ..... 25
  - SETTING THE LD\_LIBRARY\_PATH ENVIRONMENT VARIABLE ..... 25
- CONFIGURING ODBC CONNECTIONS FOR NON-WINDOWS PLATFORMS** ..... 25
  - FILES ..... 26

SAMPLE FILES .....	26
CONFIGURING THE ENVIRONMENT .....	27
CONFIGURING THE ODBC.INI FILE .....	27
CONFIGURING THE ODBCINST.INI FILE .....	29
CONFIGURING THE CLUDERA.HIVEODBC.INI FILE .....	30
CONFIGURING SERVICE DISCOVERY MODE .....	31
CONFIGURING AUTHENTICATION .....	31
CONFIGURING SSL VERIFICATION .....	33
<b>FEATURES .....</b>	<b>34</b>
SQL QUERY VERSUS HIVEQL QUERY .....	34
SQL CONNECTOR .....	34
DATA TYPES .....	35
CATALOG AND SCHEMA SUPPORT .....	36
HIVE_SYSTEM TABLE .....	36
SERVER-SIDE PROPERTIES .....	36
TEMPORARY TABLE .....	37
GET TABLES WITH QUERY .....	38
ACTIVE DIRECTORY .....	38
WRITE-BACK .....	38
DYNAMIC SERVICE DISCOVERY USING ZOOKEEPER .....	39
<b>CONTACT US .....</b>	<b>39</b>
<b>APPENDIX A USING A CONNECTION STRING .....</b>	<b>40</b>
DSN CONNECTIONS .....	40
DSN-LESS CONNECTIONS .....	40
<b>APPENDIX B AUTHENTICATION OPTIONS .....</b>	<b>41</b>
<b>APPENDIX C CONFIGURING KERBEROS AUTHENTICATION FOR WINDOWS .....</b>	<b>43</b>
ACTIVE DIRECTORY .....	43
MIT KERBEROS .....	43
<b>APPENDIX D DRIVER CONFIGURATION OPTIONS .....</b>	<b>47</b>
CONFIGURATION OPTIONS APPEARING IN THE USER INTERFACE .....	47
CONFIGURATION OPTIONS HAVING ONLY KEY NAMES .....	63
<b>APPENDIX E ODBC API CONFORMANCE LEVEL .....</b>	<b>65</b>

## Introduction

The Cloudera ODBC Driver for Apache Hive is used for direct SQL and HiveQL access to Apache Hadoop / Hive distributions, enabling Business Intelligence (BI), analytics, and reporting on Hadoop / Hive-based data. The driver efficiently transforms an application's SQL query into the equivalent form in HiveQL, which is a subset of SQL-92. If an application is Hive-aware, then the driver is configurable to pass the query through to the database for processing. The driver interrogates Hive to obtain schema information to present to a SQL-based application. Queries, including joins, are translated from SQL to HiveQL. For more information about the differences between HiveQL and SQL, see "Features" on page 34.

The Cloudera ODBC Driver for Apache Hive complies with the ODBC 3.52 data standard and adds important functionality such as Unicode and 32- and 64-bit support for high-performance computing environments.

ODBC is one of the most established and widely supported APIs for connecting to and working with databases. At the heart of the technology is the ODBC driver, which connects an application to the database. For more information about ODBC, see <http://www.simba.com/resources/data-access-standards-library>. For complete information about the ODBC specification, see the *ODBC API Reference* at [http://msdn.microsoft.com/en-us/library/windows/desktop/ms714562\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/ms714562(v=vs.85).aspx).

The *Installation and Configuration Guide* is suitable for users who are looking to access data residing within Hive from their desktop environment. Application developers may also find the information helpful. Refer to your application for details on connecting via ODBC.

## Windows Driver

### System Requirements

You install the Cloudera ODBC Driver for Apache Hive on client computers accessing data in a Hadoop cluster with the Hive service installed and running. Each computer where you install the driver must meet the following minimum system requirements:

- One of the following operating systems (32- and 64-bit editions are supported):
  - Windows® XP with SP3
  - Windows® Vista
  - Windows® 7 Professional and Enterprise
  - Windows® 8 Pro and Enterprise
  - Windows® Server 2008 R2
- 25 MB of available disk space

**Important:**

To install the driver, you must have Administrator privileges on the computer.

The driver is suitable for use with all versions of Apache Hive.

## Installing the Driver

On 64-bit Windows operating systems, you can execute 32- and 64-bit applications transparently. You must use the version of the driver matching the bitness of the client application accessing data in Hadoop / Hive:

- **ClouderaHiveODBC32.msi** for 32-bit applications
- **ClouderaHiveODBC64.msi** for 64-bit applications

You can install both versions of the driver on the same computer.

### Note:

For an explanation of how to use ODBC on 64-bit editions of Windows, see <http://www.simba.com/wp-content/uploads/2010/10/HOW-TO-32-bit-vs-64-bit-ODBC-Data-Source-Administrator.pdf>

### To install the Cloudera ODBC Driver for Apache Hive:

1. Depending on the bitness of your client application, double-click to run **ClouderaHiveODBC32.msi** or **ClouderaHiveODBC64.msi**
2. Click **Next**
3. Select the check box to accept the terms of the License Agreement if you agree, and then click **Next**
4. To change the installation location, click **Change**, then browse to the desired folder, and then click **OK**. To accept the installation location, click **Next**
5. Click **Install**
6. When the installation completes, click **Finish**

## Verifying the Version Number

If you need to verify the version of the Cloudera ODBC Driver for Apache Hive that is installed on your Windows machine, you can find the version number in the ODBC Data Source Administrator.

### To verify the version number:

1. Click the **Start** button, then click **All Programs**, then click the **Cloudera ODBC Driver for Apache Hive 2.5** program group corresponding to the bitness of the client application accessing data in Hadoop / Hive, and then click **ODBC Administrator**
2. In the ODBC Data Source Administrator, click the **Drivers** tab and then find the Cloudera ODBC Driver for Apache Hive in the list of ODBC drivers that are installed on your system. The version number is displayed in the Version column.

## Creating a Data Source Name

After installing the Cloudera ODBC Driver for Apache Hive, you need to create a Data Source Name (DSN).

Alternatively, for information about DSN-less connections, see "Configuring a DSN-less Connection" on page 9.

**To create a Data Source Name:**

1. Click the **Start** button, then click **All Programs**, then click the **Cloudera ODBC Driver for Apache Hive 2.5** program group corresponding to the bitness of the client application accessing data in Hadoop / Hive, and then click **ODBC Administrator**
2. In the ODBC Data Source Administrator, click the **Drivers** tab, and then scroll down as needed to confirm that the Cloudera ODBC Driver for Apache Hive appears in the alphabetical list of ODBC drivers that are installed on your system.
3. To create a DSN that only the user currently logged into Windows can use, click the **User DSN** tab.

OR

To create a DSN that all users who log into Windows can use, click the **System DSN** tab.

4. Click **Add**
5. In the Create New Data Source dialog box, select **Cloudera ODBC Driver for Apache Hive** and then click **Finish**
6. Use the options in the Cloudera ODBC Driver for Apache Hive DSN Setup dialog box to configure your DSN:
  - a) In the Data Source Name field, type a name for your DSN.
  - b) Optionally, in the Description field, type relevant details about the DSN.
  - c) In the Hive Server Type list, select **Hive Server 1** or **Hive Server 2**

**Note:**

If you are connecting through Apache ZooKeeper, then **Hive Server 1** is not supported.

- d) To connect to Hive without using the Apache ZooKeeper service, in the Service Discovery Mode list, select **No Service Discovery**

OR

To enable the driver to discover Hive Server 2 services via the ZooKeeper service, in the Service Discovery Mode list, select **ZooKeeper**

- e) In the Host(s) field, if you selected **No Service Discovery** in step d, then type the IP address or host name of the Hive server.

OR

If you selected **ZooKeeper** in step d, then type a comma-separated list of ZooKeeper servers. Use the following format, where *zk\_host* is the IP address or host name of the ZooKeeper server and *zk\_port* is the number of the port that the ZooKeeper server uses:

*zk\_host1:zk\_port1, zk\_host2:zk\_port2*

- f) In the Port field, if you selected **No Service Discovery** in step d, then type the number of the TCP port on which the Hive server is listening. Otherwise, do not type a value in the field.
- g) In the Database field, type the name of the database schema to use when a schema is not explicitly specified in a query.

**Note:**

You can still issue queries on other schemas by explicitly specifying the schema in the query. To inspect your databases and determine the appropriate schema to use, type the **show databases** command at the Hive command prompt.

- h) In the ZooKeeper Namespace field, if you selected **ZooKeeper** in step d, then type the namespace on ZooKeeper under which Hive Server 2 znodes are added. Otherwise, do not type a value in the field.
- i) In the Authentication area, configure authentication as needed. For more information, see "Configuring Authentication" on page 11.

**Note:**

Hive Server 1 does not support authentication. Most default configurations of Hive Server 2 require User Name authentication. To verify the authentication mechanism that you need to use for your connection, check the configuration of your Hadoop / Hive distribution. For more information, see "Authentication Options" on page 41.

- j) Optionally, if the operations against Hive are to be done on behalf of a user that is different than the authenticated user for the connection, type the name of the user to be delegated in the Delegation UID field.

**Note:**

This option is applicable only when connecting to a Hive Server 2 instance that supports this feature.

- k) In the Thrift Transport list, select the transport protocol to use in the Thrift layer.
- l) To configure HTTP options such as custom headers, click **HTTP Options**. For more information, see "Configuring HTTP Options" on page 15.

**Note:**

The HTTP options are available only when the Thrift Transport option is set to HTTP.

- m) To configure client-server verification over SSL, click **SSL Options**. For more information, see "Configuring SSL Verification" on page 16.
- n) To configure advanced driver options, click **Advanced Options**. For more information, see "Configuring Advanced Options" on page 13.



- o) To configure server-side properties, click **Advanced Options** and then click **Server Side Properties**. For more information, see "Configuring Server-Side Properties" on page 14.
- p) To configure the Temporary Table feature, click **Advanced Options** and then click **Temporary Table Configuration**. For more information, see "Configuring the Temporary Table Feature" on page 17 and "Temporary Table" on page 37.

**Important:**

When connecting to Hive 0.14 or later, the Temporary Tables feature is always enabled and you do not need to configure it in the driver.

7. To test the connection, click **Test**. Review the results as needed, and then click **OK**.

**Note:**

If the connection fails, then confirm that the settings in the Cloudera ODBC Driver for Apache Hive DSN Setup dialog box are correct. Contact your Hive server administrator as needed.

8. To save your settings and close the Cloudera ODBC Driver for Apache Hive DSN Setup dialog box, click **OK**
9. To close the ODBC Data Source Administrator, click **OK**

## Configuring a DSN-less Connection

Some client applications provide support for connecting to a data source using a driver without a Data Source Name (DSN). To configure a DSN-less connection, you can use a connection string or the Cloudera Hive ODBC Driver Configuration tool that is installed with the Cloudera ODBC Driver for Apache Hive. The following section explains how to use the driver configuration tool. For information about using connection strings, see "DSN-less Connections" on page 40.

### To configure a DSN-less connection using the driver configuration tool:

1. Click the **Start** button, then click **All Programs**, and then click the **Cloudera ODBC Driver for Apache Hive 2.5** program group corresponding to the bitness of the client application accessing data in Hadoop / Hive.
2. Click **Driver Configuration**, and then click **OK** if prompted for administrator permission to make modifications to the computer.

**Note:**

You must have administrator access to the computer in order to run this application because it makes changes to the registry.

3. In the Hive Server Type list, select **Hive Server 1** or **Hive Server 2**

**Note:**

If you are connecting through Apache ZooKeeper, then **Hive Server 1** is not supported.

4. To connect to Hive without using the Apache ZooKeeper service, in the Service Discovery Mode list, select **No Service Discovery**

OR

To enable the driver to discover Hive Server 2 services via the ZooKeeper service, in the Service Discovery Mode list, select **ZooKeeper**

5. In the ZooKeeper Namespace field, if you selected **ZooKeeper** in step 4, then type the namespace on ZooKeeper under which Hive Server 2 znodes are added. Otherwise, do not type a value in the field.
6. In the Authentication area, configure authentication as needed. For more information, see "Configuring Authentication" on page 11.

**Note:**

Hive Server 1 does not support authentication. Most default configurations of Hive Server 2 require User Name authentication. To verify the authentication mechanism that you need to use for your connection, check the configuration of your Hadoop / Hive distribution. For more information, see "Authentication Options" on page 41.

7. Optionally, if the operations against Hive are to be done on behalf of a user that is different than the authenticated user for the connection, type the name of the user to be delegated in the Delegation UID field.

**Note:**

This option is applicable only when connecting to a Hive Server 2 instance that supports this feature.

8. In the Thrift Transport list, select the transport protocol to use in the Thrift layer.
9. To configure HTTP options such as custom headers, click **HTTP Options**. For more information, see "Configuring HTTP Options" on page 15.

**Note:**

The HTTP options are available only when the Thrift Transport option is set to HTTP.

10. To configure client-server verification over SSL, click **SSL Options**. For more information, see "Configuring SSL Verification" on page 16.
11. To configure advanced options, click **Advanced Options**. For more information, see "Configuring Advanced Options" on page 13.
12. To configure server-side properties, click **Advanced Options** and then click **Server Side Properties**. For more information, see "Configuring Server-Side Properties" on page 14.
13. To configure the Temporary Table feature, click **Advanced Options** and then click **Temporary Table Configuration**. For more information, see "Temporary Table" on page 37 and "Configuring the Temporary Table Feature" on page 17.

**Important:**

When connecting to Hive 0.14 or later, the Temporary Tables feature is always enabled and you do not need to configure it in the driver.

14. To save your settings and close the Cloudera Hive ODBC Driver Configuration tool, click **OK**

## Configuring Authentication

Some Hive servers are configured to require authentication for access. To connect to a Hive server, you must configure the Cloudera ODBC Driver for Apache Hive to use the authentication mechanism that matches the access requirements of the server and provides the necessary credentials.

For information about how to determine the type of authentication your Hive server requires, see Appendix B "Authentication Options" on page 41.

ODBC applications that connect to Hive Server 2 using a DSN can pass in authentication credentials by defining them in the DSN. To configure authentication for a connection that uses a DSN, use the ODBC Data Source Administrator.

Normally, applications that are not Hive Server 2 aware and that connect using a DSN-less connection do not have a facility for passing authentication credentials to the Cloudera ODBC Driver for Apache Hive for a connection. However, the Cloudera Hive ODBC Driver Configuration tool enables you to configure authentication without using a DSN.

**Important:**

Credentials defined in a DSN take precedence over credentials configured using the driver configuration tool. Credentials configured using the driver configuration tool apply for all connections that are made using a DSN-less connection unless the client application is Hive Server 2 aware and requests credentials from the user.

## Using No Authentication

When connecting to a Hive server of type Hive Server 1, you must use No Authentication. When you use No Authentication, SASL is not supported.

### To configure a connection without authentication:

1. To access authentication options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, and then click **Configure**

OR

To access authentication options for a DSN-less connection, open the Cloudera Hive ODBC Driver Configuration tool.

2. In the Mechanism list, select **No Authentication**
3. In the Thrift Transport list, select the transport protocol to use in the Thrift layer.

**Important:**

When using this authentication mechanism, SASL is not supported.

4. If the Hive server is configured to use SSL, then click **SSL Options** to configure SSL for the connection. For more information, see "Configuring SSL Verification" on page 16.
5. To save your settings and close the dialog box, click **OK**

## Using Kerberos

Kerberos must be installed and configured before you can use this authentication mechanism. For more information, see Appendix C "Configuring Kerberos Authentication for Windows" on page 43.

This authentication mechanism is available only for Hive Server 2. When you use Kerberos authentication, the Binary transport protocol is not supported.

### To configure Kerberos authentication:

1. To access authentication options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, and then click **Configure**

OR

To access authentication options for a DSN-less connection, open the Cloudera Hive ODBC Driver Configuration tool.

2. In the Mechanism list, select **Kerberos**
3. If your Kerberos setup does not define a default realm or if the realm of your Hive Server 2 host is not the default, then type the Kerberos realm of the Hive Server 2 host in the **Realm** field.

OR

To use the default realm defined in your Kerberos setup, leave the Realm field empty.

4. In the Host FQDN field, type the fully qualified domain name of the Hive Server 2 host.
5. In the Service Name field, type the service name of the Hive server.
6. In the Thrift Transport list, select the transport protocol to use in the Thrift layer.

**Important:**

When using this authentication mechanism, the Binary transport protocol is not supported.

7. If the Hive server is configured to use SSL, then click **SSL Options** to configure SSL for the connection. For more information, see "Configuring SSL Verification" on page 16.
8. To save your settings and close the dialog box, click **OK**

## Using User Name

This authentication mechanism requires a user name but not a password. The user name labels the session, facilitating database tracking.

This authentication mechanism is available only for Hive Server 2. Most default configurations of Hive Server 2 require User Name authentication. When you use User Name authentication, SSL is not supported and you must use SASL as the Thrift transport protocol.

**To configure User Name authentication:**

1. To access authentication options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, and then click **Configure**

OR

To access authentication options for a DSN-less connection, open the Cloudera Hive ODBC Driver Configuration tool.

2. In the Mechanism list, select **User Name**
3. In the User Name field, type an appropriate user name for accessing the Hive server.
4. In the Thrift Transport list, select **SASL**
5. To save your settings and close the dialog box, click **OK**

### Using User Name and Password

This authentication mechanism requires a user name and a password.

This authentication mechanism is available only for Hive Server 2.

**To configure User Name and Password authentication:**

1. To access authentication options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, and then click **Configure**

OR

To access authentication options for a DSN-less connection, open the Cloudera Hive ODBC Driver Configuration tool.

2. In the Mechanism list, select **User Name and Password**
3. In the User Name field, type an appropriate user name for accessing the Hive server.
4. In the Password field, type the password corresponding to the user name you typed in step 3.
5. In the Thrift Transport list, select the transport protocol to use in the Thrift layer.
6. If the Hive server is configured to use SSL, then click **SSL Options** to configure SSL for the connection. For more information, see "Configuring SSL Verification" on page 16.
7. To save your settings and close the dialog box, click **OK**

### Configuring Advanced Options

You can configure advanced options to modify the behavior of the driver.

**To configure advanced options:**

1. To access advanced options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **Advanced Options**

OR

To access advanced options for a DSN-less connection, open the Cloudera Hive ODBC Driver Configuration tool, and then click **Advanced Options**

2. To disable the SQL Connector feature, select the **Use Native Query** check box.
3. To defer query execution to SQLExecute, select the **Fast SQLPrepare** check box.
4. To allow driver-wide configurations to take precedence over connection and DSN settings, select the **Driver Config Take Precedence** check box.
5. To use the asynchronous version of the API call against Hive for executing a query, select the **Use Async Exec** check box.

**Note:**

This option is applicable only when connecting to a Hive cluster running Hive 0.12.0 or later.

6. To retrieve the names of tables in a database by using the SHOW TABLES query, select the **Get Tables With Query** check box.

**Note:**

This option is applicable only when connecting to Hive Server 2.

7. To enable the driver to return SQL\_WVARCHAR instead of SQL\_VARCHAR for STRING and VARCHAR columns, and SQL\_WCHAR instead of SQL\_CHAR for CHAR columns, select the **Unicode SQL character types** check box.
8. To enable the driver to return the hive\_system table for catalog function calls such as SQLTables and SQLColumns, select the **Show System Table** check box.
9. In the Rows fetched per block field, type the number of rows to be fetched per block.
10. In the Default string column length field, type the maximum data length for STRING columns.
11. In the Binary column length field, type the maximum data length for BINARY columns.
12. In the Decimal column scale field, type the maximum number of digits to the right of the decimal point for numeric data types.
13. To save your settings and close the Advanced Options dialog box, click **OK**

## Configuring Server-Side Properties

You can use the driver to apply configuration properties to the Hive server.

**To configure server-side properties:**

1. To configure server-side properties for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN and click **Configure**, then click **Advanced Options**, and then click **Server Side Properties**

OR

To configure server-side properties for a DSN-less connection, open the Cloudera Hive ODBC Driver Configuration tool, then click **Advanced Options**, and then click **Server Side Properties**

2. To create a server-side property, click **Add**, then type appropriate values in the Key and Value fields, and then click **OK**

**Note:**

For a list of all Hadoop and Hive server-side properties that your implementation supports, type **set -v** at the Hive CLI command line or Beeline. You can also execute the set -v query after connecting using the driver.

3. To edit a server-side property, select the property from the list, then click **Edit**, then update the Key and Value fields as needed, and then click **OK**
4. To delete a server-side property, select the property from the list, and then click **Remove**. In the confirmation dialog box, click **Yes**
5. To configure the driver to apply each server-side property by executing a query when opening a session to the Hive server, select the **Apply Server Side Properties with Queries** check box.

OR

To configure the driver to use a more efficient method for applying server-side properties that does not involve additional network round-tripping, clear the **Apply Server Side Properties with Queries** check box.

**Note:**

The more efficient method is not available for Hive Server 1, and it might not be compatible with some Hive Server 2 builds. If the server-side properties do not take effect when the check box is clear, then select the check box.

6. To force the driver to convert server-side property key names to all lower case characters, select the **Convert Key Name to Lower Case** check box.
7. To save your settings and close the Server Side Properties dialog box, click **OK**

**Configuring HTTP Options**

You can configure options such as custom headers when using the HTTP transport protocol in the Thrift layer.

**To configure HTTP options:**

1. If you are configuring HTTP for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then ensure that the Thrift Transport option is set to **HTTP**

OR

If you are configuring HTTP for a DSN-less connection, open the Cloudera Hive ODBC Driver Configuration tool and then ensure that the Thrift Transport option is set to **HTTP**

2. To access HTTP options, click **HTTP Options**

**Note:**

The HTTP options are available only when the Thrift Transport option is set to HTTP.

3. In the HTTP Path field, type the partial URL corresponding to the Hive server.
4. To create a custom HTTP header, click **Add**, then type appropriate values in the Key and Value fields, and then click **OK**
5. To edit a custom HTTP header, select the header from the list, then click **Edit**, then update the Key and Value fields as needed, and then click **OK**
6. To delete a custom HTTP header, select the header from the list, and then click **Remove**. In the confirmation dialog box, click **Yes**
7. To save your settings and close the HTTP Options dialog box, click **OK**

## Configuring SSL Verification

You can configure verification between the client and the Hive server over SSL.

**To configure SSL verification:**

1. To access SSL options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **SSL Options**

OR

To access advanced options for a DSN-less connection, open the Cloudera Hive ODBC Driver Configuration tool, and then click **SSL Options**

2. Select the **Enable SSL** check box.
3. To allow self-signed certificates from the server, select the **Allow Self-signed Server Certificate** check box.
4. To allow the common name of a CA-issued SSL certificate to not match the host name of the Hive server, select the **Allow Common Name Host Name Mismatch** check box.
5. To configure the driver to load SSL certificates from a specific PEM file when verifying the server, specify the full path to the file in the Trusted Certificates field.

OR

To use the trusted CA certificates PEM file that is installed with the driver, leave the Trusted Certificates field empty.



6. If you want to configure two-way SSL verification, select the **Two Way SSL** check box and then do the following:
  - a) In the Client Certificate File field, specify the full path of the PEM file containing the client's certificate.
  - b) In the Client Private Key File field, specify the full path of the file containing the client's private key.
  - c) If the private key file is protected with a password, type the password in the Client Private Key Password field. To save the password, select the **Save Password (Encrypted)** check box.

**Important:**

The password will be obscured (not saved in plain text). However, it is still possible for the encrypted password to be copied and used.

7. To save your settings and close the SSL Options dialog box, click **OK**

## Configuring the Temporary Table Feature

You can configure the driver to create temporary tables. For more information about this feature, including details about the statement syntax used for temporary tables, see "Temporary Table" on page 37.

**Important:**

When connecting to Hive 0.14 or later, the Temporary Tables feature is always enabled and you do not need to configure it in the driver.

### To configure the Temporary Table feature:

1. To configure the temporary table feature for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN and click **Configure**, then click **Advanced Options**, and then click **Temporary Table Configuration**

OR

To configure the temporary table feature for a DSN-less connection, open the Cloudera Hive ODBC Driver Configuration tool, then click **Advanced Options**, and then click **Temporary Table Configuration**

2. To enable the Temporary Table feature, select the **Enable Temporary Table** check box.
3. In the Web HDFS Host field, type the host name or IP address of the machine hosting both the namenode of your Hadoop cluster and the WebHDFS service. If this field is left blank, then the host name of the Hive server will be used.
4. In the Web HDFS Port field, type the WebHDFS port for the namenode.
5. In the HDFS User field, type the name of the HDFS user that the driver will use to create the necessary files for supporting the Temporary Table feature.
6. In the Data file HDFS dir field, type the HDFS directory that the driver will use to store the necessary files for supporting the Temporary Table feature.

**Note:**

Due to a problem in Hive (see <https://issues.apache.org/jira/browse/HIVE-4554>), HDFS paths with space characters do not work with versions of Hive prior to 0.12.0.

7. In the Temp Table TTL field, type the number of minutes that a temporary table is guaranteed to exist in Hive after it is created.
8. To save your settings and close the Temporary Table Configuration dialog box, click **OK**

## Linux Driver

### System Requirements

You install the Cloudera ODBC Driver for Apache Hive on client computers accessing data in a Hadoop cluster with the Hive service installed and running. Each computer where you install the driver must meet the following minimum system requirements:

- One of the following distributions (32- and 64-bit editions are supported):
  - Red Hat® Enterprise Linux® (RHEL) 5.0 or 6.0
  - CentOS 5.0 or 6.0
  - SUSE Linux Enterprise Server (SLES) 11
- 45 MB of available disk space
- One of the following ODBC driver managers installed:
  - iODBC 3.52.7 or later
  - unixODBC 2.2.12 or later

The driver is suitable for use with all versions of Hive.

### Installing the Driver

There are two versions of the driver for Linux:

- **ClouderaHiveODBC-32bit-Version-Release.LinuxDistro.i686.rpm** for the 32-bit driver
- **ClouderaHiveODBC-Version-Release.LinuxDistro.x86\_64.rpm** for the 64-bit driver

*Version* is the version number of the driver, and *Release* is the release number for this version of the driver.

The bitness of the driver that you select should match the bitness of the client application accessing your Hadoop / Hive-based data. For example, if the client application is 64-bit, then you should install the 64-bit driver. Note that 64-bit editions of Linux support both 32- and 64-bit applications. Verify the bitness of your intended application and install the appropriate version of the driver.

**Important:**

Ensure that you install the driver using the RPM corresponding to your Linux distribution.

The Cloudera ODBC Driver for Apache Hive driver files are installed in the following directories:

- **/opt/cloudera/hiveodbc** contains release notes, the *Cloudera ODBC Driver for Apache Hive Installation and Configuration Guide* in PDF format, and a Readme.txt file that provides plain text installation and configuration instructions.
- **/opt/cloudera/hiveodbc/ErrorMessage**s contains error message files required by the driver.
- **/opt/cloudera/hiveodbc/Setup** contains sample configuration files named `odbc.ini` and `odbcinst.ini`
- **/opt/cloudera/hiveodbc/lib/32** contains the 32-bit shared libraries and the `cloudera.hiveodbc.ini` configuration file.
- **/opt/cloudera/hiveodbc/lib/64** contains the 64-bit shared libraries and the `cloudera.hiveodbc.ini` configuration file.

**To install the Cloudera ODBC Driver for Apache Hive:**

1. In Red Hat Enterprise Linux or CentOS, log in as the root user, then navigate to the folder containing the driver RPM packages to install, and then type the following at the command line, where *RPMFileName* is the file name of the RPM package containing the version of the driver that you want to install:

```
yum --nogpgcheck localinstall RPMFileName
```

OR

In SUSE Linux Enterprise Server, log in as the root user, then navigate to the folder containing the driver RPM packages to install, and then type the following at the command line, where *RPMFileName* is the file name of the RPM package containing the version of the driver that you want to install:

```
zypper install RPMFileName
```

The Cloudera ODBC Driver for Apache Hive depends on the following resources:

- `cyrus-sasl-2.1.22-7` or above
- `cyrus-sasl-gssapi-2.1.22-7` or above
- `cyrus-sasl-plain-2.1.22-7` or above

If the package manager in your Linux distribution cannot resolve the dependencies automatically when installing the driver, then download and manually install the packages required by the version of the driver that you want to install.

**Verifying the Version Number**

If you need to verify the version of the Cloudera ODBC Driver for Apache Hive that is installed on your Linux machine, you can query the version number through the command-line interface.

### To verify the version number:

- At the command prompt, run the following command:

```
yum list | grep ClouderaHiveODBC
```

OR

Run the following command:

```
rpm -qa | grep ClouderaHiveODBC
```

The command returns information about the Cloudera ODBC Driver for Apache Hive that is installed on your machine, including the version number.

## Setting the LD\_LIBRARY\_PATH Environment Variable

The LD\_LIBRARY\_PATH environment variable must include the paths to the installed ODBC driver manager libraries.

For example, if ODBC driver manager libraries are installed in /usr/local/lib, then set LD\_LIBRARY\_PATH as follows:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/lib
```

For information about how to set environment variables permanently, refer to your Linux shell documentation.

For information about creating ODBC connections using the Cloudera ODBC Driver for Apache Hive, see "Configuring ODBC Connections for Non-Windows Platforms" on page 25.

## Mac OS X Driver

### System Requirements

You install the Cloudera ODBC Driver for Apache Hive on client computers accessing data in a Hadoop cluster with the Hive service installed and running. Each computer where you install the driver must meet the following minimum system requirements:

- Mac OS X version 10.6.8 or later
- 100 MB of available disk space
- iODBC 3.52.7 or later

The driver is suitable for use with all versions of Hive. The driver supports both 32- and 64-bit client applications.

### Installing the Driver

The Cloudera ODBC Driver for Apache Hive driver files are installed in the following directories:

- **/opt/cloudera/hiveodbc** contains release notes and the *Cloudera ODBC Driver for Apache Hive Installation and Configuration Guide* in PDF format.
- **/opt/cloudera/hiveodbc/ErrorMessage**s contains error messages required by the driver.

- **/opt/cloudera/hiveodbc/Setup** contains sample configuration files named `odbc.ini` and `odbcinst.ini`
- **/opt/cloudera/hiveodbc/lib/universal** contains the driver binaries and the `cloudera.hiveodbc.ini` configuration file.

#### To install the Cloudera ODBC Driver for Apache Hive:

1. Double-click **ClouderaHiveODBC.dmg** to mount the disk image.
2. Double-click **ClouderaHiveODBC.pkg** to run the installer.
3. In the installer, click **Continue**
4. On the Software License Agreement screen, click **Continue**, and when the prompt appears, click **Agree** if you agree to the terms of the License Agreement.
5. Optionally, to change the installation location, click **Change Install Location**, then select the desired location, and then click **Continue**
6. To accept the installation location and begin the installation, click **Install**
7. When the installation completes, click **Close**

### Verifying the Version Number

If you need to verify the version of the Cloudera ODBC Driver for Apache Hive that is installed on your Mac OS X machine, you can query the version number through the Terminal.

#### To verify the version number:

- At the Terminal, run the following command:

```
pkgutil --info cloudera.hiveodbc
```

The command returns information about the Cloudera ODBC Driver for Apache Hive that is installed on your machine, including the version number.

### Setting the DYLD\_LIBRARY\_PATH Environment Variable

The `DYLD_LIBRARY_PATH` environment variable must include the paths to the installed ODBC driver manager libraries.

For example, if ODBC driver manager libraries are installed in `/usr/local/lib`, then set `DYLD_LIBRARY_PATH` as follows:

```
export DYLD_LIBRARY_PATH=$DYLD_LIBRARY_PATH:/usr/local/lib
```

For information about how to set environment variables permanently, refer to your Mac OS X shell documentation.

For information about creating ODBC connections using the Cloudera ODBC Driver for Apache Hive, see "Configuring ODBC Connections for Non-Windows Platforms" on page 25.

## AIX Driver

### System Requirements

You install the Cloudera ODBC Driver for Apache Hive on client computers accessing data in a Hadoop cluster with the Hive service installed and running. Each computer where you install the driver must meet the following minimum system requirements:

- IBM AIX 5.3, 6.1, or 7.1 (32- and 64-bit editions are supported)
- 150 MB of available disk space
- One of the following ODBC driver managers installed:
  - iODBC 3.52.7 or later
  - unixODBC 2.3.0 or later

The driver is suitable for use with all versions of Hive.

### Installing the Driver

There are two versions of the driver for AIX:

- **ClouderaHiveODBC-32bit-Version-Release.ppc.rpm** for the 32-bit driver
- **ClouderaHiveODBC-Version-Release.ppc.rpm** for the 64-bit driver

*Version* is the version number of the driver, and *Release* is the release number for this version of the driver.

The bitness of the driver that you select should match the bitness of the client application accessing your Hadoop / Hive-based data. For example, if the client application is 64-bit, then you should install the 64-bit driver. Note that 64-bit editions of AIX support both 32- and 64-bit applications. Verify the bitness of your intended application and install the appropriate version of the driver.

The Cloudera ODBC Driver for Apache Hive driver files are installed in the following directories:

- **/opt/cloudera/hiveodbc** contains release notes, the *Cloudera ODBC Driver for Apache Hive Installation and Configuration Guide* in PDF format, and a Readme.txt file that provides plain text installation and configuration instructions.
- **/opt/cloudera/hiveodbc/ErrorMessage**s contains error message files required by the driver.
- **/opt/cloudera/hiveodbc/Setup** contains sample configuration files named `odbc.ini` and `odbcinst.ini`
- **/opt/cloudera/hiveodbc/lib/32** contains the 32-bit driver and the `cloudera.hiveodbc.ini` configuration file.
- **/opt/cloudera/hiveodbc/lib/64** contains the 64-bit driver and the `cloudera.hiveodbc.ini` configuration file.

**To install the Cloudera ODBC Driver for Apache Hive:**

1. Log in as the root user, then navigate to the folder containing the driver RPM packages to install, and then type the following at the command line, where *RPMFileName* is the file name of the RPM package containing the version of the driver that you want to install:

```
rpm --install RPMFileName
```

**Verifying the Version Number**

If you need to verify the version of the Cloudera ODBC Driver for Apache Hive that is installed on your AIX machine, you can query the version number through the command-line interface.

**To verify the version number:**

- At the command prompt, run the following command:

```
rpm -qa | grep ClouderaHiveODBC
```

The command returns information about the Cloudera ODBC Driver for Apache Hive that is installed on your machine, including the version number.

**Setting the LD\_LIBRARY\_PATH Environment Variable**

The LD\_LIBRARY\_PATH environment variable must include the path to the installed ODBC driver manager libraries.

For example, if ODBC driver manager libraries are installed in /usr/local/lib, then set LD\_LIBRARY\_PATH as follows:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/lib
```

For information about how to set environment variables permanently, refer to your AIX shell documentation.

For information about creating ODBC connections using the Cloudera ODBC Driver for Apache Hive, see "Configuring ODBC Connections for Non-Windows Platforms" on page 25.

**Debian Driver****System Requirements**

You install the Cloudera ODBC Driver for Apache Hive on client computers accessing data in a Hadoop cluster with the Hive service installed and running. Each computer where you install the driver must meet the following minimum system requirements:

- Debian 7 (Ubuntu 12.04 LTS and Ubuntu 14.04 LTS)
- 45 MB of available disk space
- One of the following ODBC driver managers installed:
  - iODBC 3.52.7 or later
  - unixODBC 2.2.12 or later

The driver is suitable for use with all versions of Hive. It supports both 32- and 64-bit client applications.

## Installing the Driver

There are two versions of the driver for Debian:

- **ClouderaHiveODBC-32bit-Version-Release\_i386.deb** for the 32-bit driver
- **ClouderaHiveODBC-Version-Release\_amd64.deb** for the 64-bit driver

*Version* is the version number of the driver, and *Release* is the release number for this version of the driver.

The bitness of the driver that you select should match the bitness of the client application accessing your Hadoop / Hive-based data. For example, if the client application is 64-bit, then you should install the 64-bit driver. Note that 64-bit editions of Debian support both 32- and 64-bit applications. Verify the bitness of your intended application and install the appropriate version of the driver.

The Cloudera ODBC Driver for Apache Hive driver files are installed in the following directories:

- **/opt/cloudera/hiveodbc** contains release notes, the *Cloudera ODBC Driver for Apache Hive Installation and Configuration Guide* in PDF format, and a Readme.txt file that provides plain text installation and configuration instructions.
- **/opt/cloudera/hiveodbc/ErrorMessage**s contains error message files required by the driver.
- **/opt/cloudera/hiveodbc/Setup** contains sample configuration files named odbc.ini and odbcinst.ini
- **/opt/cloudera/hiveodbc/lib/32** contains the 32-bit shared libraries and the cloudera.hiveodbc.ini configuration file.
- **/opt/cloudera/hiveodbc/lib/64** contains the 64-bit shared libraries and the cloudera.hiveodbc.ini configuration file.

### To install the Cloudera ODBC Driver for Apache Hive:

1. In Ubuntu, log in as the root user, then navigate to the folder containing the driver Debian packages to install, and double-click **ClouderaHiveODBC-32bit-Version-Release\_i386.deb** or **ClouderaHiveODBC-Version-Release\_amd64.deb**
2. Follow the instructions in the installer to complete the installation process.
3. If you received a license file via e-mail, then copy the license file into the `/opt/cloudera/hiveodbc/lib/32` or `/opt/cloudera/hiveodbc/lib/64` folder, depending on the version of the driver that you installed.

**Note:**

To avoid security issues, you may need to save the license file on your local computer prior to copying the file into the folder.



The Cloudera ODBC Driver for Apache Hive depends on the following resources:

- cyrus-sasl-2.1.22-7 or above
- cyrus-sasl-gssapi-2.1.22-7 or above
- cyrus-sasl-plain-2.1.22-7 or above

If the package manager in your Ubuntu distribution cannot resolve the dependencies automatically when installing the driver, then download and manually install the packages required by the version of the driver that you want to install.

## Verifying the Version Number

If you need to verify the version of the Cloudera ODBC Driver for Apache Hive that is installed on your Debian machine, you can query the version number through the command-line interface.

**To verify the version number:**

- At the command prompt, run the following command:

```
dpkg -l | grep ClouderaHiveODBC
```

The command returns information about the Cloudera ODBC Driver for Apache Hive that is installed on your machine, including the version number.

## Setting the LD\_LIBRARY\_PATH Environment Variable

The LD\_LIBRARY\_PATH environment variable must include the path to the installed ODBC driver manager libraries.

For example, if ODBC driver manager libraries are installed in /usr/local/lib, then set LD\_LIBRARY\_PATH as follows:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/lib
```

For information about how to set environment variables permanently, refer to your Ubuntu shell documentation.

For information about creating ODBC connections using the Cloudera ODBC Driver for Apache Hive, see "Configuring ODBC Connections for Non-Windows Platforms" on page 25.

## Configuring ODBC Connections for Non-Windows Platforms

The following sections describe how to configure ODBC connection when using the Cloudera ODBC Driver for Apache Hive with non-Windows platforms:

- "Files" on page 26
- "Sample Files" on page 26
- "Configuring the Environment" on page 27
- "Configuring the odbc.ini File" on page 27
- "Configuring the odbcinst.ini File" on page 29

- "Configuring the cloudera.hiveodbc.ini File" on page 30
- "Configuring Service Discovery Mode" on page 31
- "Configuring Authentication" on page 31

### Files

ODBC driver managers use configuration files to define and configure ODBC data sources and drivers. By default, the following configuration files residing in the user's home directory are used:

- **.odbc.ini** is used to define ODBC data sources, and it is required for DSNs.
- **.odbcinst.ini** is used to define ODBC drivers, and it is optional.

Also, by default the Cloudera ODBC Driver for Apache Hive is configured using the `cloudera.hiveodbc.ini` file, which is located in one of the following directories depending on the version of the driver that you are using:

- **/opt/cloudera/hiveodbc/lib/32** for the 32-bit driver on Linux/AIX/Debian
- **/opt/cloudera/hiveodbc/lib/64** for the 64-bit driver on Linux/AIX/Debian
- **/opt/cloudera/hiveodbc/lib/universal** for the driver on Mac OS X

The `cloudera.hiveodbc.ini` file is required.

#### Note:

The `cloudera.hiveodbc.ini` file in the `/lib` subfolder provides default settings for most configuration options available in the Cloudera ODBC Driver for Apache Hive.

You can set driver configuration options in your `odbc.ini` and `cloudera.hiveodbc.ini` files. Configuration options set in a `cloudera.hiveodbc.ini` file apply to all connections, whereas configuration options set in an `odbc.ini` file are specific to a connection. Configuration options set in `odbc.ini` take precedence over configuration options set in `cloudera.hiveodbc.ini`. For information about the configuration options available for controlling the behavior of DSNs that are using the Cloudera ODBC Driver for Apache Hive, see Appendix D "Driver Configuration Options" on page 47.

### Sample Files

The driver installation contains the following sample configuration files in the Setup directory:

- `odbc.ini`
- `odbcinst.ini`

These sample configuration files provide preset values for settings related to the Cloudera ODBC Driver for Apache Hive.

The names of the sample configuration files do not begin with a period (.) so that they will appear in directory listings by default. A filename beginning with a period (.) is hidden. For `odbc.ini` and `odbcinst.ini`, if the default location is used, then the filenames must begin with a period (.).

If the configuration files do not exist in the home directory, then you can copy the sample configuration files to the home directory, and then rename the files. If the configuration files already exist in the home directory, then use the sample configuration files as a guide to modify the existing configuration files.

## Configuring the Environment

Optionally, you can use three environment variables—ODBCINI, ODBCYSINI, and CLOUDERAHIVEINI—to specify different locations for the `odbc.ini`, `odbcinst.ini`, and `cloudera.hiveodbc.ini` configuration files by doing the following:

- Set ODBCINI to point to your `odbc.ini` file.
- Set ODBCYSINI to point to the directory containing the `odbcinst.ini` file.
- Set CLOUDERAHIVEINI to point to your `cloudera.hiveodbc.ini` file.

For example, if your `odbc.ini` and `cloudera.hiveodbc.ini` files are located in `/etc` and your `odbcinst.ini` file is located in `/usr/local/odbc`, then set the environment variables as follows:

```
export ODBCINI=/etc/odbc.ini
export ODBCYSINI=/usr/local/odbc
export CLOUDERAHIVEINI=/etc/cloudera.hiveodbc.ini
```

The following search order is used to locate the `cloudera.hiveodbc.ini` file:

1. If the CLOUDERAHIVEINI environment variable is defined, then the driver searches for the file specified by the environment variable.

### Important:

CLOUDERAHIVEINI must specify the full path, including the file name.

2. The directory containing the driver's binary is searched for a file named `cloudera.hiveodbc.ini` (not beginning with a period).
3. The current working directory of the application is searched for a file named `cloudera.hiveodbc.ini` (not beginning with a period).
4. The directory `~/` (that is, `$HOME`) is searched for a hidden file named `.cloudera.hiveodbc.ini`
5. The directory `/etc` is searched for a file named `cloudera.hiveodbc.ini` (not beginning with a period).

## Configuring the `odbc.ini` File

### Note:

If you are using a DSN-less connection, then you do not need to configure the `odbc.ini` file. For information about configuring a DSN-less connection, see "DSN-less Connections" on page 40.

ODBC Data Source Names (DSNs) are defined in the `odbc.ini` configuration file. The file is divided into several sections:

- **[ODBC]** is optional and used to control global ODBC configuration, such as ODBC tracing.
- **[ODBC Data Sources]** is required, listing DSNs and associating DSNs with a driver.
- A section having the same name as the data source specified in the **[ODBC Data Sources]** section is required to configure the data source.

The following is an example of an `odbc.ini` configuration file for Linux/AIX/Debian:

```
[ODBC Data Sources]
Sample Cloudera Hive DSN 32=Cloudera Hive ODBC Driver 32-bit
[Sample Cloudera Hive DSN 32]
Driver=/opt/cloudera/hiveodbc/lib/32/libclouderahiveodbc32.so
HOST=MyHiveServer
PORT=10000
```

*MyHiveServer* is the IP address or host name of the Hive server.

The following is an example of an `odbc.ini` configuration file for Mac OS X:

```
[ODBC Data Sources]
Sample Cloudera Hive DSN=Cloudera Hive ODBC Driver
[Sample Cloudera Hive DSN]
Driver=/opt/cloudera/hiveodbc/lib/universal/libclouderahiveodbc.dyl
ib
HOST=MyHiveServer
PORT=10000
```

*MyHiveServer* is the IP address or host name of the Hive server.

### To create a Data Source Name:

1. Open the `.odbc.ini` configuration file in a text editor.
2. In the **[ODBC Data Sources]** section, add a new entry by typing the Data Source Name (DSN), then an equal sign (=), and then the driver name.
3. In the `.odbc.ini` file, add a new section with a name that matches the DSN you specified in step 2, and then add configuration options to the section. Specify configuration options as key-value pairs.

#### Note:

Hive Server 1 does not support authentication. Most default configurations of Hive Server 2 require User Name authentication, which you configure by setting the `AuthMech` key to 2. To verify the authentication mechanism that you need to use for your connection, check the configuration of your Hadoop / Hive distribution. For more information, see "Authentication Options" on page 41.

4. Save the .odbc.ini configuration file.

For information about the configuration options available for controlling the behavior of DSNs that are using the Cloudera ODBC Driver for Apache Hive, see Appendix D "Driver Configuration Options" on page 47.

## Configuring the odbcinst.ini File

ODBC drivers are defined in the odbcinst.ini configuration file. The configuration file is optional because drivers can be specified directly in the odbc.ini configuration file, as described in "Configuring the odbc.ini File" on page 27.

The odbcinst.ini file is divided into the following sections:

- **[ODBC Drivers]** lists the names of all the installed ODBC drivers.
- A section having the same name as the driver name specified in the [ODBC Drivers] section lists driver attributes and values.

The following is an example of an odbcinst.ini configuration file for Linux/AIX/Debian:

```
[ODBC Drivers]
Cloudera Hive ODBC Driver 32-bit=Installed
Cloudera Hive ODBC Driver 64-bit=Installed
[Cloudera Hive ODBC Driver 32-bit]
Description=Cloudera Hive ODBC Driver (32-bit)
Driver=/opt/cloudera/hiveodbc/lib/32/libclouderahiveodbc32.so
[Cloudera Hive ODBC Driver 64-bit]
Description=Cloudera Hive ODBC Driver (64-bit)
Driver=/opt/cloudera/hiveodbc/lib/64/libclouderahiveodbc64.so
```

The following is an example of an odbcinst.ini configuration file for Mac OS X:

```
[ODBC Drivers]
Cloudera Hive ODBC Driver=Installed
[Cloudera Hive ODBC Driver]
Description=Cloudera Hive ODBC Driver
Driver=/opt/cloudera/hiveodbc/lib/universal/libclouderahiveodbc.dylib
```

### To define a driver:

1. Open the .odbcinst.ini configuration file in a text editor.
2. In the [ODBC Drivers] section, add a new entry by typing the driver name and then typing **=Installed**

**Note:**

Type a symbolic name that you want to use to refer to the driver in connection strings or DSNs.

3. In the `.odbcinst.ini` file, add a new section with a name that matches the driver name you typed in step 2, and then add configuration options to the section based on the sample `odbcinst.ini` file provided in the Setup directory. Specify configuration options as key-value pairs.
4. Save the `.odbcinst.ini` configuration file.

## Configuring the `cloudera.hiveodbc.ini` File

The `cloudera.hiveodbc.ini` file contains configuration settings for the Cloudera ODBC Driver for Apache Hive. Settings that you define in the `cloudera.hiveodbc.ini` file apply to all connections that use the driver.

### To configure the Cloudera ODBC Driver for Apache Hive to work with your ODBC driver manager:

1. Open the `cloudera.hiveodbc.ini` configuration file in a text editor.
2. Edit the `DriverManagerEncoding` setting. The value is usually **UTF-16** or **UTF-32** if you are using Linux/Mac OS X, depending on the ODBC driver manager you use. iODBC uses **UTF-32**, and unixODBC uses **UTF-16**. To determine the correct setting to use, refer to your ODBC Driver Manager documentation.

OR

If you are using AIX and the unixODBC driver manager, then set the value to **UTF-16**. If you are using AIX and the iODBC driver manager, then set the value to **UTF-16** for the 32-bit driver or **UTF-32** for the 64-bit driver.

3. Edit the `ODBCInstLib` setting. The value is the name of the `ODBCInst` shared library for the ODBC driver manager you use. To determine the correct library to specify, refer to your ODBC driver manager documentation.

The configuration file defaults to the shared library for iODBC. In Linux/AIX/Debian, the shared library name for iODBC is **libiodbcinst.so**. In Mac OS X, the shared library name for iODBC is **libiodbcinst.dylib**.

**Note:**

You can specify an absolute or relative filename for the library. If you intend to use the relative filename, then the path to the library must be included in the library path environment variable. In Linux/AIX/Debian, the library path environment variable is named `LD_LIBRARY_PATH`. In Mac OS X, the library path environment variable is named `DYLD_LIBRARY_PATH`.

4. Save the `cloudera.hiveodbc.ini` configuration file.

## Configuring Service Discovery Mode

You can configure the Cloudera ODBC Driver for Apache Hive to discover Hive Server 2 services via ZooKeeper.

### To enable service discovery via ZooKeeper:

1. Open the `odbc.ini` configuration file in a text editor.
2. Set the `ServiceDiscoveryMode` connection attribute to 1
3. Set the `ZKNamespace` connection attribute to specify the namespace on ZooKeeper under which Hive Server 2 `znodes` are added.
4. Set the `Host` connection attribute to specify the ZooKeeper ensemble as a comma-separated list of ZooKeeper servers. For example, type the following, where `zk_host` is the IP address or host name of the ZooKeeper server and `zk_port` is the number of the port that the ZooKeeper server uses:

```
zk_host1:zk_port1,zk_host2:zk_port2
```

#### Important:

When `ServiceDiscoveryMode` is set to 1, connections to Hive Server 1 are not supported and the `Port` connection attribute is not applicable.

Depending on whether service discovery mode is enabled or disabled, you may need to provide different connection attributes or values in your connection string or DSN. For more information about connection attributes, see "Driver Configuration Options" on page 47.

## Configuring Authentication

Some Hive servers are configured to require authentication for access. To connect to a Hive server, you must configure the Cloudera ODBC Driver for Apache Hive to use the authentication mechanism that matches the access requirements of the server and provides the necessary credentials.

For information about how to determine the type of authentication your Hive server requires, see Appendix B "Authentication Options" on page 41.

You can select the type of authentication to use for a connection by defining the `AuthMech` connection attribute in a connection string or in a DSN (in the `odbc.ini` file). Depending on the authentication mechanism you use, there may be additional connection attributes that you must define. For more information about the attributes involved in configuring authentication, see Appendix D "Driver Configuration Options" on page 47.

### Using No Authentication

When connecting to a Hive server of type Hive Server 1, you must use No Authentication. When you use No Authentication, SASL is not supported.

**To configure a connection without authentication:**

1. Set the AuthMech connection attribute to 0
2. Set the ThriftTransport connection attribute to the transport protocol to use in the Thrift layer.

**Important:**

When using this authentication mechanism, SASL (**ThriftTransport=1**) is not supported.

3. If the Hive server is configured to use SSL, then configure SSL for the connection. For more information, see "Configuring SSL Verification" on page 33.

**Using Kerberos**

Kerberos must be installed and configured before you can use this authentication mechanism. For more information, refer to the MIT Kerberos documentation.

This authentication mechanism is available only for Hive Server 2. When you use Kerberos authentication, the Binary transport protocol is not supported.

**To configure Kerberos authentication:**

1. Set the AuthMech connection attribute to 1
2. If your Kerberos setup does not define a default realm or if the realm of your Hive server is not the default, then set the appropriate realm using the KrbRealm attribute.

OR

To use the default realm defined in your Kerberos setup, do not set the KrbRealm attribute.

3. Set the KrbHostFQDN attribute to the fully qualified domain name of the Hive Server 2 host.
4. Set the KrbServiceName attribute to the service name of the Hive server.
5. Set the ThriftTransport connection attribute to the transport protocol to use in the Thrift layer.

**Important:**

When using this authentication mechanism, Binary (**ThriftTransport=0**) is not supported.

6. If the Hive server is configured to use SSL, then configure SSL for the connection. For more information, see "Configuring SSL Verification" on page 33.

**Using User Name**

This authentication mechanism requires a user name but does not require a password. The user name labels the session, facilitating database tracking.

This authentication mechanism is available only for Hive Server 2. Most default configurations of Hive Server 2 require User Name authentication. When you use User Name authentication, SSL is not supported and you must use SASL as the Thrift transport protocol.



**To configure User Name authentication:**

1. Set the AuthMech connection attribute to 2
2. Set the UID attribute to an appropriate user name for accessing the Hive server.
3. Set the ThriftTransport connection attribute to 1

**Using User Name and Password**

This authentication mechanism requires a user name and a password.

This authentication mechanism is available only for Hive Server 2.

**To configure User Name and Password authentication:**

1. Set the AuthMech connection attribute to 3
2. Set the UID attribute to an appropriate user name for accessing the Hive server.
3. Set the PWD attribute to the password corresponding to the user name you provided in step 2.
4. Set the ThriftTransport connection attribute to the transport protocol to use in the Thrift layer.
5. If the Hive server is configured to use SSL, then configure SSL for the connection. For more information, see "Configuring SSL Verification" on page 33.

**Configuring SSL Verification**

You can configure verification between the client and the Hive server over SSL.

**To configure SSL verification:**

1. Open the odbc.ini configuration file in a text editor.
2. To enable SSL, set the SSL attribute to 1
3. To allow self-signed certificates from the server, set the AllowSelfSignedServerCert attribute to 1.
4. To allow the common name of a CA-issued SSL certificate to not match the host name of the Hive server, set the CAIssuedCertNamesMismatch attribute to 1.
5. To configure the driver to load SSL certificates from a specific PEM file when verifying the server, set the TrustedCerts attribute to the full path of the PEM file.

OR

To use the trusted CA certificates PEM file that is installed with the driver, do not specify a value for the TrustedCerts attribute.

6. If you want to configure two-way SSL verification, set the TwoWaySSL attribute to 1 and then do the following:
  - a) Set the ClientCert attribute to the full path of the PEM file containing the client's certificate.
  - b) Set the ClientPrivateKey attribute to the full path of the file containing the client's private key.

- c) If the private key file is protected with a password, set the ClientPrivateKeyPassword attribute to the password.
7. Save the odbc.ini configuration file.

## Features

More information is provided on the following features of the Cloudera ODBC Driver for Apache Hive:

- "SQL Query versus HiveQL Query" on page 34
- "SQL Connector" on page 34
- "Data Types" on page 35
- "Catalog and Schema Support" on page 36
- "hive\_system Table" on page 36
- "Server-Side Properties" on page 36
- "Get Tables With Query" on page 38
- "Active Directory" on page 38
- "Write-back" on page 38
- "Dynamic Service Discovery using ZooKeeper" on page 39

## SQL Query versus HiveQL Query

The native query language supported by Hive is HiveQL. For simple queries, HiveQL is a subset of SQL-92. However, the syntax is different enough that most applications do not work with native HiveQL.

## SQL Connector

To bridge the difference between SQL and HiveQL, the SQL Connector feature translates standard SQL-92 queries into equivalent HiveQL queries. The SQL Connector performs syntactical translations and structural transformations. For example:

- **Quoted Identifiers** — The double quotes (") that SQL uses to quote identifiers are translated into back quotes (`) to match HiveQL syntax. The SQL Connector needs to handle this translation because even when a driver reports the back quote as the quote character, some applications still generate double-quoted identifiers.
- **Table Aliases** — Support is provided for the AS keyword between a table reference and its alias, which HiveQL normally does not support.
- **JOIN, INNER JOIN, and CROSS JOIN** — SQL JOIN, INNER JOIN, and CROSS JOIN syntax is translated to HiveQL JOIN syntax.
- **TOP N/LIMIT** — SQL TOP N queries are transformed to HiveQL LIMIT queries.

## Data Types

The Cloudera ODBC Driver for Apache Hive supports many common data formats, converting between Hive data types and SQL data types.

Table 1 lists the supported data type mappings.

**Table 1. Supported Data Types**

Hive Type	SQL Type
TINYINT	SQL_TINYINT
SMALLINT	SQL_SMALLINT
INT	SQL_INTEGER
BIGINT	SQL_BIGINT
FLOAT	SQL_REAL
DOUBLE	SQL_DOUBLE
DECIMAL	SQL_DECIMAL
BOOLEAN	SQL_BIT
STRING	SQL_VARCHAR <div data-bbox="532 1157 1313 1350" style="border: 1px solid orange; padding: 5px; margin-top: 10px;"> <p><b>Note:</b></p> <p>SQL_WVARCHAR is returned instead if the Unicode SQL Character Types configuration option (the UseUnicodeSqlCharacterTypes key) is enabled.</p> </div>
TIMESTAMP	SQL_TYPE_TIMESTAMP
VARCHAR(n)	SQL_VARCHAR
DATE	SQL_TYPE_DATE
DECIMAL(p,s)	SQL_DECIMAL
CHAR(n)	SQL_CHAR

Hive Type	SQL Type
	<div style="border: 1px solid orange; padding: 5px;"> <p><b>Note:</b></p> <p>SQL_WCHAR is returned instead if the Unicode SQL Character Types configuration option (the UseUnicodeSqlCharacterTypes key) is enabled.</p> </div>
BINARY	SQL_VARBINARY

**Note:**

The aggregate types (ARRAY, MAP, and STRUCT) are not yet supported. Columns of aggregate types are treated as STRING columns.

## Catalog and Schema Support

The Cloudera ODBC Driver for Apache Hive supports both catalogs and schemas in order to make it easy for the driver to work with various ODBC applications. Since Hive only organizes tables into schemas/databases, the driver provides a synthetic catalog called “HIVE” under which all of the schemas/databases are organized. The driver also maps the ODBC schema to the Hive schema/database.

## hive\_system Table

A pseudo-table called hive\_system can be used to query for Hive cluster system environment information. The pseudo-table is under the pseudo-schema called hive\_system. The table has two STRING type columns, envkey and envvalue. Standard SQL can be executed against the hive\_system table. For example:

```
SELECT * FROM HIVE.hive_system.hive_system WHERE envkey LIKE '%hive%'
```

The above query returns all of the Hive system environment entries whose key contains the word “hive.” A special query, set -v, is executed to fetch system environment information. Some versions of Hive do not support this query. For versions of Hive that do not support querying system environment information, the driver returns an empty result set.

## Server-Side Properties

The Cloudera ODBC Driver for Apache Hive allows you to set server-side properties via a DSN. Server-side properties specified in a DSN affect only the connection that is established using the DSN.

You can also specify server-side properties for connections that do not use a DSN. To do this, use the Cloudera Hive ODBC Driver Configuration tool that is installed with the Windows version of the driver, or set the appropriate configuration options in your connection string or the cloudera.hiveodbc.ini file. Properties specified in the driver configuration tool or the

cloudera.hiveodbc.ini file apply to all connections that use the Cloudera ODBC Driver for Apache Hive.

For more information about setting server-side properties when using the Windows driver, see "Configuring Server-Side Properties" on page 14. For information about setting server-side properties when using the driver on a non-Windows platform, see "Driver Configuration Options" on page 47.

## Temporary Table

The Temporary Table feature adds support for creating temporary tables and inserting literal values into temporary tables. Temporary tables are only accessible by the ODBC connection that created them and they will be dropped upon disconnect.

### CREATE TABLE Statement for Temporary Tables

The driver supports the following DDL syntax for creating temporary tables:

```
<create table statement> := CREATE TABLE <temporary table name>
<left paren><column definition list><right paren>
<column definition list> := <column definition>[, <column
definition>]*
<column definition> := <column name> <data type>
<temporary table name> := <double quote><number sign><table
name><double quote>
<left paren> := (
<right paren> := )
<double quote> := "
<number sign> := #
```

The following is an example of a SQL statement for creating a temporary table:

```
CREATE TABLE "#TEMPTABLE1" (C1 DATATYPE_1, C2 DATATYPE_2, ..., Cn
DATATYPE_n)
```

The temporary table name in a SQL query must be surrounded by double quotes ("), and the name must begin with a number sign (#).

**Note:**

You can only use data types that are supported by Hive.

### INSERT Statement for Temporary Tables

The driver supports the following DDL syntax for inserting data into temporary tables:

```
<insert statement> := INSERT INTO <temporary table name> <left
paren><column name list><right paren> VALUES <left paren><literal
value list><right paren>
<column name list> := <column name>[, <column name>]*
```

```
<literal value list> := <literal value>[, <literal value>]*  
<temporary table name> := <double quote><number sign><table  
name><double quote>  
<left paren> := (  
<right paren> := )  
<double quote> := "  
<number sign> := #
```

The following is an example of a SQL statement for inserting data into temporary tables:

```
INSERT INTO "#TEMPTABLE1" values (VAL(C1), VAL(C2) ... VAL(Cn) )
```

VAL(C1) is the literal value for the first column in the table, and VAL(Cn) is the literal value for the nth column in the table.

**Note:**

The INSERT statement is only supported for temporary tables.

## Get Tables With Query

The Get Tables With Query configuration option allows you to choose whether to use the SHOW TABLES query or the GetTables API call to retrieve table names from a database.

Hive Server 2 has a limit on the number of tables that can be in a database when handling the GetTables API call. When the number of tables in a database is above the limit, the API call will return a stack overflow error or a timeout error. The exact limit and the error that appears depends on the JVM settings.

As a workaround for this issue, enable the Get Tables with Query configuration option (or GetTablesWithQuery key) to use the query instead of the API call.

## Active Directory

The Cloudera ODBC Driver for Apache Hive supports Active Directory Kerberos on Windows. There are two prerequisites for using Active Directory Kerberos on Windows:

- MIT Kerberos is **not** installed on the client Windows machine.
- The MIT Kerberos Hadoop realm has been configured to trust the Active Directory realm, according to Cloudera's documentation, so that users in the Active Directory realm can access services in the MIT Kerberos Hadoop realm.

## Write-back

The Cloudera ODBC Driver for Apache Hive supports translation for INSERT, UPDATE, and DELETE syntax when connecting to a Hive Server 2 instance that is running Hive 0.14 or later.

## Dynamic Service Discovery using ZooKeeper

The Cloudera ODBC Driver for Apache Hive can be configured to discover Hive Server 2 services via the ZooKeeper service.

For information about configuring this feature in the Windows driver, see "Creating a Data Source Name" on page 6 or "Configuring a DSN-less Connection" on page 9. For information about configuring this feature when using the driver on a non-Windows platform, see "Configuring Service Discovery Mode" on page 31.

## Contact Us

If you have difficulty using the driver, please contact our Technical Support staff. We welcome your questions, comments, and feature requests.

**Important:**

To help us assist you, prior to contacting Technical Support please prepare a detailed summary of the client and server environment including operating system version, patch level, and configuration.

For details on contacting Technical Support, see <http://www.cloudera.com/content/cloudera/en/products/cloudera-support.html>

## Appendix A Using a Connection String

For some applications, you may need to use a connection string to connect to your data source.

### DSN Connections

The following is an example of a connection string for a connection that uses a DSN:

```
DSN=DataSourceName; Key=Value
```

*DataSourceName* is the DSN that you are using for the connection. *Key* is any connection attribute that is not already specified as a configuration key in the DSN, and *Value* is the value for the attribute. Add key-value pairs to the connection string as needed, separating each pair with a semicolon (;). For information about the connection attributes that are available, see "Driver Configuration Options" on page 47.

### DSN-less Connections

Some client applications provide support for connecting to a data source using a driver without a DSN. The following is an example of a connection string for a DSN-less connection:

```
Driver=DriverNameOrFile; HOST=MyHiveServer; PORT=PortNumber;  
Schema=DefaultSchema; HiveServerType=ServerType
```

The placeholders in the connection string are defined as follows:

- *DriverNameOrFile* is either the symbolic name of the installed driver defined in the `odbcinst.ini` file or the absolute path of the shared object file for the driver. If you use the symbolic name, then you must ensure that the `odbcinst.ini` file is configured to point to the symbolic name to the shared object file. For more information, see "Configuring the `odbcinst.ini` File" on page 29.
- *MyHiveServer* is the IP address or host name of the Hive server.
- *PortNumber* is the number of the port that the Hive server uses.
- *DefaultSchema* is the database schema to use when a schema is not explicitly specified in a query.
- *ServerType* is either 1 (for Hive Server 1) or 2 (for Hive Server 2).

Add key-value pairs to the connection string as needed, separating each pair with a semicolon (;). For information about the connection attributes that are available, see "Driver Configuration Options" on page 47.



## Appendix B Authentication Options

To connect to a Hive server, you must configure the Cloudera ODBC Driver for Apache Hive to use the authentication mechanism that matches the access requirements of the server and provides the necessary credentials. To determine the authentication settings that your Hive server requires, check the server type and then refer to the corresponding section below.

### Hive Server 1

You must use No Authentication as the authentication mechanism. Hive Server 1 instances do not support authentication.

### Hive Server 2

**Note:**

Most default configurations of Hive Server 2 require User Name authentication.

Configuring authentication for a connection to a Hive Server 2 instance involves setting the authentication mechanism, the Thrift transport protocol, and SSL support. To determine the settings that you need to use, check the following three properties in the hive-site.xml file in the Hive server that you are connecting to:

- hive.server2.authentication
- hive.server2.transport.mode
- hive.server2.use.SSL

Use Table 2 to determine the authentication mechanism that you need to configure, based on the hive.server2.authentication value in the hive-site.xml file:

**Table 2. Authentication Mechanisms to Use**

hive.server2.authentication	Authentication Mechanism
NOSASL	No Authentication
KERBEROS	Kerberos
NONE	User Name
LDAP	User Name and Password

Use Table 3 to determine the Thrift transport protocol that you need to configure, based on the `hive.server2.authentication` and `hive.server2.transport.mode` values in the `hive-site.xml` file:

**Table 3. Thrift Transport Protocol Setting to Use**

<code>hive.server2.authentication</code>	<code>hive.server2.transport.mode</code>	Thrift Transport Protocol
NOSASL	binary	Binary
KERBEROS	binary or http	SASL or HTTP
NONE	binary or http	SASL or HTTP
LDAP	binary or http	SASL or HTTP

To determine whether SSL should be enabled or disabled for your connection, check the `hive.server2.use.SSL` value in the `hive-site.xml` file. If the value is true, then you must enable and configure SSL in your connection. If the value is false, then you must disable SSL in your connection.

For detailed instructions on how to configure authentication when using the Windows driver, see "Configuring Authentication" on page 11.

For detailed instructions on how to configure authentication when using a non-Windows driver, see "Configuring Authentication" on page 31.

## Appendix C Configuring Kerberos Authentication for Windows

### Active Directory

The Cloudera ODBC Driver for Apache Hive supports Active Directory Kerberos on Windows. There are two prerequisites for using Active Directory Kerberos on Windows:

- MIT Kerberos is not installed on the client Windows machine.
- The MIT Kerberos Hadoop realm has been configured to trust the Active Directory realm so that users in the Active Directory realm can access services in the MIT Kerberos Hadoop realm.

### MIT Kerberos

#### Downloading and Installing MIT Kerberos for Windows 4.0.1

For information about Kerberos and download links for the installer, see the MIT Kerberos website at <http://web.mit.edu/kerberos/>

#### To download and install MIT Kerberos for Windows 4.0.1:

1. To download the Kerberos installer for 64-bit computers, use the following download link from the MIT Kerberos website: <http://web.mit.edu/kerberos/dist/kfw/4.0/kfw-4.0.1-amd64.msi>

OR

To download the Kerberos installer for 32-bit computers, use the following download link from the MIT Kerberos website: <http://web.mit.edu/kerberos/dist/kfw/4.0/kfw-4.0.1-i386.msi>

**Note:**

The 64-bit installer includes both 32-bit and 64-bit libraries. The 32-bit installer includes 32-bit libraries only.

2. To run the installer, double-click the .msi file that you downloaded in step 1.
3. Follow the instructions in the installer to complete the installation process.
4. When the installation completes, click **Finish**

#### Setting Up the Kerberos Configuration File

Settings for Kerberos are specified through a configuration file. You can set up the configuration file as a .INI file in the default location—the C:\ProgramData\MIT\Kerberos5 directory— or as a .CONF file in a custom location.

Normally, the C:\ProgramData\MIT\Kerberos5 directory is hidden. For information about viewing and using this hidden directory, refer to Microsoft Windows documentation.

**Note:**

For more information on configuring Kerberos, refer to the MIT Kerberos documentation.

**To set up the Kerberos configuration file in the default location:**

1. Obtain a krb5.conf configuration file from your Kerberos administrator.

OR

Obtain the configuration file from the /etc/krb5.conf folder on the computer that is hosting the Hive Server 2 instance.

2. Rename the configuration file from krb5.conf to krb5.ini
3. Copy the krb5.ini file to the C:\ProgramData\MIT\Kerberos5 directory and overwrite the empty sample file.

**To set up the Kerberos configuration file in a custom location:**

1. Obtain a krb5.conf configuration file from your Kerberos administrator.

OR

Obtain the configuration file from the /etc/krb5.conf folder on the computer that is hosting the Hive Server 2 instance.

2. Place the krb5.conf file in an accessible directory and make note of the full path name.
3. Click the **Start** button , then right-click **Computer**, and then click **Properties**
4. Click **Advanced System Settings**
5. In the System Properties dialog box, click the **Advanced** tab and then click **Environment Variables**
6. In the Environment Variables dialog box, under the System variables list, click **New**
7. In the New System Variable dialog box, in the Variable name field, type **KRB5\_CONFIG**
8. In the Variable value field, type the absolute path to the krb5.conf file from step 2.
9. Click **OK** to save the new variable.
10. Ensure that the variable is listed in the System variables list.
11. Click **OK** to close the Environment Variables dialog box, and then click **OK** to close the System Properties dialog box.

## Setting Up the Kerberos Credential Cache File

Kerberos uses a credential cache to store and manage credentials.

**To set up the Kerberos credential cache file:**

1. Create a directory where you want to save the Kerberos credential cache file. For example, create a directory named C:\temp
2. Click the **Start** button , then right-click **Computer**, and then click **Properties**
3. Click **Advanced System Settings**

4. In the System Properties dialog box, click the **Advanced** tab and then click **Environment Variables**
5. In the Environment Variables dialog box, under the System variables list, click **New**
6. In the New System Variable dialog box, in the Variable name field, type **KRB5CCNAME**
7. In the Variable value field, type the path to the folder you created in step 1, and then append the file name krb5cache. For example, if you created the folder C:\temp in step 1, then type C:\temp\krb5cache

**Note:**

krb5cache is a file (not a directory) that is managed by the Kerberos software, and it should not be created by the user. If you receive a permission error when you first use Kerberos, ensure that the krb5cache file does not already exist as a file or a directory.

8. Click **OK** to save the new variable.
9. Ensure that the variable appears in the System variables list.
10. Click **OK** to close the Environment Variables dialog box, and then click **OK** to close the System Properties dialog box.
11. To ensure that Kerberos uses the new settings, restart your computer.

### Obtaining a Ticket for a Kerberos Principal

A principal refers to a user or service that can authenticate to Kerberos. To authenticate to Kerberos, a principal must obtain a ticket by using a password or a keytab file. You can specify a keytab file to use, or use the default keytab file of your Kerberos configuration.

#### To obtain a ticket for a Kerberos principal using a password:

1. Click the **Start** button, then click **All Programs**, and then click the **Kerberos for Windows (64-bit)** or **Kerberos for Windows (32-bit)** program group.
2. Click **MIT Kerberos Ticket Manager**
3. In the MIT Kerberos Ticket Manager, click **Get Ticket**
4. In the Get Ticket dialog box, type your principal name and password, and then click **OK**

If the authentication succeeds, then your ticket information appears in the MIT Kerberos Ticket Manager.

#### To obtain a ticket for a Kerberos principal using a keytab file:

1. Click the Start button, then click All Programs, then click Accessories, and then click Command Prompt
2. In the Command Prompt, type a command using the following syntax:
 

```
kinit -k -t keytab_path principal
```

*keytab\_path* is the full path to the keytab file. For example: C:\mykeytabs\myUser.keytab

*principal* is the Kerberos user principal to use for authentication. For example:  
myUser@EXAMPLE.COM

3. If the cache location KRB5CCNAME is not set or used, then use the -c option of the kinit command to specify the location of the credential cache. In the command, the -c argument must appear last. For example:

```
kinit -k -t C:\mykeytabs\myUser.keytab myUser@EXAMPLE.COM -c  
C:\ProgramData\MIT\krbcache
```

Krbcache is the Kerberos cache file, not a directory.

#### To obtain a ticket for a Kerberos principal using the default keytab file:

**Note:**

For information about configuring a default keytab file for your Kerberos configuration, refer to the MIT Kerberos documentation.

1. Click the Start button , then click All Programs, then click Accessories, and then click Command Prompt
2. In the Command Prompt, type a command using the following syntax:

```
kinit -k principal
```

*principal* is the Kerberos user principal to use for authentication. For example:  
myUser@EXAMPLE.COM

3. If the cache location KRB5CCNAME is not set or used, then use the -c option of the kinit command to specify the location of the credential cache. In the command, the -c argument must appear last. For example:

```
kinit -k -t C:\mykeytabs\myUser.keytab myUser@EXAMPLE.COM -c  
C:\ProgramData\MIT\krbcache
```

Krbcache is the Kerberos cache file, not a directory.

## Appendix D Driver Configuration Options

Appendix D "Driver Configuration Options" on page 47 lists the configuration options available in the Cloudera ODBC Driver for Apache Hive alphabetically by field or button label. Options having only key names—not appearing in the user interface of the driver—are listed alphabetically by key name.

When creating or configuring a connection from a Windows computer, the fields and buttons are available in the Cloudera Hive ODBC Driver Configuration tool and the following dialog boxes:

- Cloudera ODBC Driver for Apache Hive DSN Setup
- HTTP Properties
- SSL Options
- Advanced Options
- Server Side Properties

When using a connection string or configuring a connection from a Linux/Mac OS X/AIX/Debian computer, use the key names provided.

**Note:**

You can pass in configuration options in your connection string or set them in your `odbc.ini` and `cloudera.hiveodbc.ini` files. Configuration options set in a `cloudera.hiveodbc.ini` file apply to all connections, whereas configuration options passed in in the connection string or set in an `odbc.ini` file are specific to a connection. Configuration options passed in using the connection string take precedence over configuration options set in `odbc.ini`. Configuration options set in `odbc.ini` take precedence over configuration options set in `cloudera.hiveodbc.ini`

### Configuration Options Appearing in the User Interface

The following configuration options are accessible via the Windows user interface for the Cloudera ODBC Driver for Apache Hive, or via the key name when using a connection string or configuring a connection from a Linux/Mac OS X/AIX/Debian computer:

- "Allow Common Name Host Name Mismatch" on page 48
- "Allow Self-signed Server Certificate" on page 48
- "Apply Properties with Queries" on page 49
- "Binary Column Length" on page 49
- "Client Certificate File" on page 50
- "Client Private Key File" on page 50
- "Client Private Key Password" on page 50
- "Host FQDN" on page 55
- "HTTP Path" on page 55
- "Mechanism" on page 56
- "Password" on page 56
- "Port" on page 56
- "Realm" on page 57
- "Rows Fetched Per Block" on page 57
- "Save Password (Encrypted)" on page 57
- "Service Discovery Mode" on page 58

- "Convert Key Name to Lower Case" on page 51
- "Data File HDFS Dir" on page 51
- "Database" on page 51
- "Decimal Column Scale" on page 52
- "Default String Column Length" on page 52
- "Delegation UID" on page 52
- "Driver Config Take Precedence" on page 52
- "Enable SSL" on page 53
- "Enable Temporary Table" on page 53
- "Fast SQLPrepare" on page 53
- "Get Tables With Query" on page 54
- "HDFS User" on page 54
- "Hive Server Type" on page 54
- "Host(s)" on page 55
- "Service Name" on page 58
- "Show System Table" on page 58
- "Temp Table TTL" on page 59
- "Thrift Transport" on page 59
- "Trusted Certificates" on page 60
- "Two Way SSL" on page 60
- "Unicode SQL Character Types" on page 61
- "Use Async Exec" on page 61
- "Use Native Query" on page 61
- "User Name" on page 62
- "Web HDFS Host" on page 62
- "Web HDFS Port" on page 62
- "ZooKeeper Namespace" on page 63

### Allow Common Name Host Name Mismatch

Key Name	Default Value	Required
CAIssuedCertNamesMismatch	Clear (0)	No

#### Description

When this option is enabled (1), the driver allows a CA-issued SSL certificate name to not match the host name of the Hive server.

When this option is disabled (0), the CA-issued SSL certificate name must match the host name of the Hive server.

#### Note:

This setting is applicable only when SSL is enabled.

### Allow Self-signed Server Certificate

Key Name	Default Value	Required
AllowSelfSignedServerCert	Clear (0)	No



**Description**

When this option is enabled (1), the driver authenticates the Hive server even if the server is using a self-signed certificate.

When this option is disabled (0), the driver does not allow self-signed certificates from the server.

**Note:**

This setting is applicable only when SSL is enabled.

**Apply Properties with Queries**

Key Name	Default Value	Required
ApplySSPWithQueries	Selected (1)	No

**Description**

When this option is enabled (1), the driver applies each server-side property by executing a `set SSPKey=SSPValue` query when opening a session to the Hive server.

When this option is disabled (0), the driver uses a more efficient method for applying server-side properties that does not involve additional network round-tripping. However, some Hive Server 2 builds are not compatible with the more efficient method.

**Note:**

When connecting to a Hive Server 1 instance, ApplySSPWithQueries is always enabled.

**Binary Column Length**

Key Name	Default Value	Required
BinaryColumnLength	32767	No

**Description**

The maximum data length for BINARY columns.

By default, the columns metadata for Hive does not specify a maximum data length for BINARY columns.

### Client Certificate File

Key Name	Default Value	Required
ClientCert	None	Yes, if two-way SSL verification is enabled.

#### Description

The full path to the PEM file containing the client's SSL certificate.

**Note:**

This setting is applicable only when two-way SSL is enabled.

### Client Private Key File

Key Name	Default Value	Required
ClientPrivateKey	None	Yes, if two-way SSL verification is enabled.

#### Description

The full path to the PEM file containing the client's SSL private key.

If the private key file is protected with a password, then provide the password using the driver configuration option "Client Private Key Password" on page 50.

**Note:**

This setting is applicable only when two-way SSL is enabled.

### Client Private Key Password

Key Name	Default Value	Required
ClientPrivateKeyPassword	None	Yes, if two-way SSL verification is enabled and the client's private key file is protected with a password.

#### Description

The password of the private key file that is specified in the Client Private Key File field (the ClientPrivateKey key).

### Convert Key Name to Lower Case

Key Name	Default Value	Required
LCaseSspKeyName	Selected (1)	No

#### Description

When this option is enabled (1), the driver converts server-side property key names to all lower case characters.

When this option is disabled (0), the driver does not modify the server-side property key names.

### Data File HDFS Dir

Key Name	Default Value	Required
HDFSTempTableDir	/tmp/simba	No

#### Description

The HDFS directory that the driver will use to store the necessary files for supporting the Temporary Table feature.

#### Note:

Due to a problem in Hive (see <https://issues.apache.org/jira/browse/HIVE-4554>), HDFS paths with space characters do not work with versions of Hive prior to 0.12.0.

This option is not applicable when connecting to Hive 0.14 or later.

### Database

Key Name	Default Value	Required
Schema	default	No

#### Description

The name of the database schema to use when a schema is not explicitly specified in a query. You can still issue queries on other schemas by explicitly specifying the schema in the query.

#### Note:

To inspect your databases and determine the appropriate schema to use, type the **show databases** command at the Hive command prompt.

### Decimal Column Scale

Key Name	Default Value	Required
DecimalColumnScale	10	No

#### Description

The maximum number of digits to the right of the decimal point for numeric data types.

### Default String Column Length

Key Name	Default Value	Required
DefaultStringColumnLength	255	No

#### Description

The maximum data length for STRING columns.

By default, the columns metadata for Hive does not specify a maximum data length for STRING columns.

### Delegation UID

Key Name	Default Value	Required
DelegationUID	None	No

#### Description

Use this option to delegate all operations against Hive to a user that is different than the authenticated user for the connection.

**Note:**

This option is applicable only when connecting to a Hive Server 2 instance that supports this feature.

### Driver Config Take Precedence

Key Name	Default Value	Required
DriverConfigTakePrecedence	Clear (0)	No

**Description**

When this option is enabled (1), driver-wide configurations take precedence over connection and DSN settings.

When this option is disabled (0), connection and DSN settings take precedence instead.

**Enable SSL**

Key Name	Default Value	Required
SSL	Clear (0)	No

**Description**

When this option is set to (1), the client verifies the Hive server using SSL.

When this option is set to (0), SSL is disabled.

**Note:**

This option is applicable only when connecting to a Hive server that supports SSL.

**Enable Temporary Table**

Key Name	Default Value	Required
EnableTempTable	Clear (0)	No

**Description**

When this option is enabled (1), the driver supports the creation and use of temporary tables.

When this option is disabled (0), the driver does not support temporary tables.

**Important:**

When connecting to Hive 0.14 or later, the Temporary Tables feature is always enabled and you do not need to configure it in the driver.

**Fast SQLPrepare**

Key Name	Default Value	Required
FastSQLPrepare	Clear (0)	No

**Description**

When this option is enabled (1), the driver defers query execution to SQLExecute.

When this option is disabled (0), the driver does not defer query execution to SQLExecute.

When using Native Query mode, the driver will execute the HiveQL query to retrieve the result set metadata for SQLPrepare. As a result, SQLPrepare might be slow. If the result set metadata is not required after calling SQLPrepare, then enable Fast SQLPrepare.

**Get Tables With Query**

Key Name	Default Value	Required
GetTablesWithQuery	Clear (0)	No

**Description**

When this option is enabled (1), the driver uses the SHOW TABLES query to retrieve the names of the tables in a database.

When this option is disabled (0), the driver uses the GetTables Thrift API call to retrieve the names of the tables in a database.

**Note:**

This option is applicable only when connecting to a Hive Server 2 instance.

**HDFS User**

Key Name	Default Value	Required
HDFSUser	hdfs	No

**Description**

The name of the HDFS user that the driver will use to create the necessary files for supporting the Temporary Tables feature.

This option is not applicable when connecting to Hive 0.14 or later.

**Hive Server Type**

Key Name	Default Value	Required
HiveServerType	Hive Server 2 (2)	No

**Description**

Select **Hive Server 1** or set the key to 1 if you are connecting to a Hive Server 1 instance.

Select **Hive Server 2** or set the key to 2 if you are connecting to a Hive Server 2 instance.

**Note:**

If Service Discovery Mode is enabled, then connections to Hive Server 1 are not supported.

**Host(s)**

Key Name	Default Value	Required
HOST	None	Yes

**Description**

If Service Discovery Mode is disabled, specify the IP address or host name of the Hive server.

If Service Discovery Mode is enabled, specify a comma-separated list of ZooKeeper servers in the following format, where *zk\_host* is the IP address or host name of the ZooKeeper server and *zk\_port* is the number of the port that the ZooKeeper server uses:

*zk\_host1:zk\_port1, zk\_host2:zk\_port2*

**Host FQDN**

Key Name	Default Value	Required
KrbHostFQDN	None	Yes, if the authentication mechanism is Kerberos.

**Description**

The fully qualified domain name of the Hive Server 2 host.

**HTTP Path**

Key Name	Default Value	Required
HTTPPath	None	Yes, if you are using HTTP as the Thrift transport protocol.

**Description**

The partial URL corresponding to the Hive server.

**Note:**

This option is applicable only when you are using HTTP as the Thrift transport protocol. For more information, see "Thrift Transport" on page 59.

**Mechanism**

Key Name	Default Value	Required
AuthMech	No Authentication (0) if you are connecting to Hive Server 1.  User Name (2) if you are connecting to Hive Server 2.	No

**Description**

The authentication mechanism to use.

Select one of the following settings, or set the key to the corresponding number:

- No Authentication (0)
- Kerberos (1)
- User Name (2)
- User Name and Password (3)

**Password**

Key Name	Default Value	Required
PWD	None	Yes, if the authentication mechanism is <b>User Name and Password</b> .

**Description**

The password corresponding to the user name that you provided in the User Name field (the UID key).

**Port**

Key Name	Default Value	Required
PORT	10000	Yes, if Service Discovery Mode is disabled.



**Description**

The number of the TCP port on which the Hive server is listening.

**Realm**

Key Name	Default Value	Required
KrbRealm	Depends on your Kerberos configuration.	No

**Description**

The realm of the Hive Server 2 host.

If your Kerberos configuration already defines the realm of the Hive Server 2 host as the default realm, then you do not need to configure this option.

**Rows Fetched Per Block**

Key Name	Default Value	Required
RowsFetchedPerBlock	10000	No

**Description**

The maximum number of rows that a query returns at a time.

Any positive 32-bit integer is a valid value, but testing has shown that performance gains are marginal beyond the default value of 10000 rows.

**Save Password (Encrypted)**

Key Name	Default Value	Required
N/A	Clear (0)	No

**Description**

This option is available only in the Windows driver. It appears in the Cloudera ODBC Driver for Apache Hive DSN Setup dialog box and the SSL Options dialog box.

When this option is enabled (1), the specified password is saved in the registry.

When this option is disabled (0), the specified password is not saved in the registry.

**Important:**

The password will be obscured (not saved in plain text). However, it is still possible for the encrypted password to be copied and used.

**Service Discovery Mode**

Key Name	Default Value	Required
ServiceDiscoveryMode	No Service Discovery (0)	No

**Description**

When this option is enabled (1), the driver discovers Hive Server 2 services via the ZooKeeper service.

When this option is disabled (0), the driver connects to Hive without using the ZooKeeper service.

**Service Name**

Key Name	Default Value	Required
KrbServiceName	None	Yes, if the authentication mechanism is Kerberos.

**Description**

The Kerberos service principal name of the Hive server.

**Show System Table**

Key Name	Default Value	Required
ShowSystemTable	Clear (0)	No

**Description**

When this option is enabled (1), the driver returns the hive\_system table for catalog function calls such as `SQLTables` and `SQLColumns`.

When this option is disabled (0), the driver does not return the hive\_system table for catalog function calls.

**Temp Table TTL**

Key Name	Default Value	Required
TempTableTTL	10	No

**Description**

The number of minutes a temporary table is guaranteed to exist in Hive after it is created.

This option is not applicable when connecting to Hive 0.14 or later.

**Thrift Transport**

Key Name	Default Value	Required
ThriftTransport	Binary (0) if you are connecting to Hive Server 1.  SASL (1) if you are connecting to Hive Server 2.	No

**Description**

The transport protocol to use in the Thrift layer.

Select one of the following settings, or set the key to the corresponding number:

- Binary (0)
- SASL (1)
- HTTP (2)

## Trusted Certificates

Key Name	Default Value	Required
TrustedCerts	<p>The cacerts.pem file in the lib folder or subfolder within the driver's installation directory.</p> <p>The exact file path varies depending on the version of the driver that is installed. For example, the path for the Windows driver is different from the path for the Mac OS X driver.</p>	No

### Description

The location of the PEM file containing trusted CA certificates for authenticating the Hive server when using SSL.

If this option is not set, then the driver will default to using the trusted CA certificates PEM file installed by the driver.

**Note:**

This setting is applicable only when SSL is enabled.

## Two Way SSL

Key Name	Default Value	Required
TwoWaySSL	Clear (0)	No

### Description

When this option is enabled (1), the client and the Hive server verify each other using SSL. See also the driver configuration options "Client Certificate File" on page 50, "Client Private Key File" on page 50, and "Client Private Key Password" on page 50.

When this option is disabled (0), then the server does not verify the client. Depending on whether one-way SSL is enabled, the client may verify the server. For more information, see "Enable SSL" on page 53.

**Note:**

This option is applicable only when connecting to a Hive server that supports SSL. You must enable SSL before Two Way SSL can be configured. For more information, see "Enable SSL" on page 53.

**Unicode SQL Character Types**

Key Name	Default Value	Required
UseUnicodeSqlCharacterTypes	Clear (0)	No

**Description**

When this option is enabled (1), the driver returns SQL\_WVARCHAR for STRING and VARCHAR columns, and returns SQL\_WCHAR for CHAR columns.

When this option is disabled (0), the driver returns SQL\_VARCHAR for STRING and VARCHAR columns, and returns SQL\_CHAR for CHAR columns.

**Use Async Exec**

Key Name	Default Value	Required
EnableAsyncExec	Clear (0)	No

**Description**

When this option is enabled (1), the driver uses an asynchronous version of the API call against Hive for executing a query.

When this option is disabled (0), the driver executes queries synchronously.

Due to a problem in Hive 0.12.0 (see <https://issues.apache.org/jira/browse/HIVE-5230>), Hive returns generic error messages for errors that occur during query execution. To see the actual error message relevant to the problem, turn off asynchronous query execution and execute the query again.

**Note:**

This option only takes effect when connecting to a Hive cluster running Hive 0.12.0 or higher.

**Use Native Query**

Key Name	Default Value	Required
UseNativeQuery	Clear (0)	No

**Description**

When this option is enabled (1), the driver does not transform the queries emitted by an application, so the native query is used.

When this option is disabled (0), the driver transforms the queries emitted by an application and converts them into an equivalent from in HiveQL.

**Note:**

If the application is Hive-aware and already emits HiveQL, then enable this option to avoid the extra overhead of query transformation.

**User Name**

Key Name	Default Value	Required
UID	For User Name authentication only, the default value is anonymous	No, if the authentication mechanism is User Name. Yes, if the authentication mechanism is <b>User Name and Password</b> .

**Description**

The user name that you use to access Hive Server 2.

**Web HDFS Host**

Key Name	Default Value	Required
WebHDFSHost	The Hive server host.	No

**Description**

The host name or IP address of the machine hosting both the namenode of your Hadoop cluster and the WebHDFS service.

This option is not applicable when connecting to Hive 0.14 or later.

**Web HDFS Port**

Key Name	Default Value	Required
WebHDFSPort	50070	No

**Description**

The WebHDFS port for the namenode.

This option is not applicable when connecting to Hive 0.14 or later.

**ZooKeeper Namespace**

Key Name	Default Value	Required
ZKNamespace	None	Yes, if Service Discovery Mode is enabled.

**Description**

The namespace on ZooKeeper under which Hive Server 2 znodes are added.

**Configuration Options Having Only Key Names**

The following configuration options do not appear in the Windows user interface for the Cloudera ODBC Driver for Apache Hive and are only accessible when using a connection string or configuring a connection from a Linux/Mac OS X/AIX/Debian computer:

- "Driver" on page 63
- "http.header." on page 63
- "SSP\_" on page 64

**Driver**

Default Value	Required
The default value varies depending on the version of the driver that is installed. For example, the value for the Windows driver is different from the value of the Mac OS X driver.	Yes

**Description**

The name of the installed driver (Cloudera ODBC Driver for Apache Hive) or the absolute path of the Cloudera ODBC Driver for Apache Hive shared object file.

**http.header.**

Default Value	Required
None	No

**Description**

Set a custom HTTP header by using the following syntax, where *HeaderKey* is the name of the header to set and *HeaderValue* is the value to assign to the header:

```
http.header.HeaderKey=HeaderValue
```

For example:

```
http.header.AUTHENTICATED_USER=john
```

After the driver applies the header, the `http.header.` prefix is removed from the DSN entry, leaving an entry of *HeaderKey=HeaderValue*

The example above would create the following custom HTTP header:

```
AUTHENTICATED_USER: john
```

**Note:**

The `http.header.` prefix is case-sensitive. This option is applicable only when you are using HTTP as the Thrift transport protocol. For more information, see "Thrift Transport" on page 59.

**SSP\_**

Default Value	Required
None	No

**Description**

Set a server-side property by using the following syntax, where *SSPKey* is the name of the server-side property to set and *SSPValue* is the value to assign to the server-side property:

```
SSP_SSPKey=SSPValue
```

For example:

```
SSP_mapred.queue.names=myQueue
```

After the driver applies the server-side property, the `SSP_` prefix is removed from the DSN entry, leaving an entry of *SSPKey=SSPValue*

**Note:**

The `SSP_` prefix must be upper case.



## Appendix E ODBC API Conformance Level

Table 4 lists the ODBC interfaces that the Cloudera ODBC Driver for Apache Hive implements and the ODBC compliance level of each interface.

ODBC compliance levels are Core, Level 1, and Level 2. These compliance levels are defined in the ODBC Specification published with the Interface SDK from Microsoft.

Interfaces include both the Unicode and non-Unicode versions. For more information, see <http://msdn.microsoft.com/en-us/library/ms716246%28VS.85%29.aspx>

Table 4. ODBC API Conformance Level

Conformance Level	INTERFACES		Conformance Level	INTERFACES
Core	SQLAllocHandle		Core	SQLGetStmtAttr
Core	SQLBindCol		Core	SQLGetTypeInfo
Core	SQLBindParameter		Core	SQLNativeSql
Core	SQLCancel		Core	SQLNumParams
Core	SQLCloseCursor		Core	SQLNumResultCols
Core	SQLColAttribute		Core	SQLParamData
Core	SQLColumns		Core	SQLPrepare
Core	SQLConnect		Core	SQLPutData
Core	SQLCopyDesc		Core	SQLRowCount
Core	SQLDescribeCol		Core	SQLSetConnectAttr
Core	SQLDisconnect		Core	SQLSetCursorName
Core	SQLDriverconnect		Core	SQLSetDescField
Core	SQLEndTran		Core	SQLSetDescRec
Core	SQLExecDirect		Core	SQLSetEnvAttr
Core	SQLExecute		Core	SQLSetStmtAttr
Core	SQLFetch		Core	SQLSpecialColumns

Conformance Level	INTERFACES		Conformance Level	INTERFACES
Core	SQLFetchScroll		Core	SQLStatistics
Core	SQLFreeHandle		Core	SQLTables
Core	SQLFreeStmt		Core	SQLBrowseConnect
Core	SQLGetConnectAttr		Core	SQLPrimaryKeys
Core	SQLGetCursorName		Core	SQLGetInfo
Core	SQLGetData		Level 1	SQLProcedureColumns
Core	SQLGetDescField		Level 1	SQLProcedures
Core	SQLGetDescRec		Level 2	SQLColumnPrivileges
Core	SQLGetDiagField		Level 2	SQLDescribeParam
Core	SQLGetDiagRec		Level 2	SQLForeignKeys
Core	SQLGetEnvAttr		Level 2	SQLTablePrivileges
Core	SQLGetFunctions			