




Simba Drill ODBC Driver with SQL Connector

Installation and Configuration Guide

Simba Technologies Inc.

Version 1.3.22
August 15, 2018



Copyright © 2018 Simba Technologies Inc. All Rights Reserved.

Information in this document is subject to change without notice. Companies, names and data used in examples herein are fictitious unless otherwise noted. No part of this publication, or the software it describes, may be reproduced, transmitted, transcribed, stored in a retrieval system, decompiled, disassembled, reverse-engineered, or translated into any language in any form by any means for any purpose without the express written permission of Simba Technologies Inc.

Trademarks

Simba, the Simba logo, SimbaEngine, and Simba Technologies are registered trademarks of Simba Technologies Inc. in Canada, United States and/or other countries. All other trademarks and/or servicemarks are the property of their respective owners.

Contact Us

Simba Technologies Inc.
938 West 8th Avenue
Vancouver, BC Canada
V5Z 1E5

Tel: +1 (604) 633-0008

Fax: +1 (604) 633-0004

www.simba.com

About This Guide

Purpose

The *Simba Drill ODBC Driver with SQL Connector Installation and Configuration Guide* explains how to install and configure the Simba Drill ODBC Driver with SQL Connector. The guide also provides details related to features of the driver.

Audience

The guide is intended for end users of the Simba Drill ODBC Driver, as well as administrators and developers integrating the driver.

Knowledge Prerequisites

To use the Simba Drill ODBC Driver, the following knowledge is helpful:

- Familiarity with the platform on which you are using the Simba Drill ODBC Driver
- Ability to use the data source to which the Simba Drill ODBC Driver is connecting
- An understanding of the role of ODBC technologies and driver managers in connecting to a data source
- Experience creating and configuring ODBC connections
- Exposure to SQL

Document Conventions

Italics are used when referring to book and document titles.

Bold is used in procedures for graphical user interface elements that a user clicks and text that a user types.

Monospace font indicates commands, source code, or contents of text files.

Note:

A text box with a pencil icon indicates a short note appended to a paragraph.

! Important:

A text box with an exclamation mark indicates an important comment related to the preceding paragraph.

Table of Contents

About the Simba Drill ODBC Driver	6
About Apache Drill	6
About the Driver	6
Windows Driver	9
Windows System Requirements	9
Installing the Driver on Windows	9
Creating a Data Source Name on Windows	10
Configuring Authentication on Windows	12
Configuring TLS (SSL) on a Windows Machine	14
Configuring Logging Options on Windows	15
Using Drill Explorer on Windows	17
Verifying the Driver Version Number on Windows	19
macOS Driver	20
macOS System Requirements	20
Installing the Driver on macOS	20
Using Drill Explorer on macOS	21
Verifying the Driver Version Number on macOS	23
Linux Driver	24
Linux System Requirements	24
Installing the Driver Using the RPM File	24
Using Drill Explorer on Linux	26
Verifying the Driver Version Number on Linux	28
Configuring the ODBC Driver Manager on Non-Windows Machines	29
Specifying ODBC Driver Managers on Non-Windows Machines	29
Specifying the Locations of the Driver Configuration Files	30
Configuring ODBC Connections on a Non-Windows Machine	32
Creating a Data Source Name on a Non-Windows Machine	32
Configuring a DSN-less Connection on a Non-Windows Machine	35
Configuring Authentication on a Non-Windows Machines	37
Configuring TLS (SSL) on a non-Windows Machine	39
Configuring Logging Options on a Non-Windows Machine	39
Testing the Connection on a Non-Windows Machine	41

Using a Connection String	43
DSN Connection String Example	43
DSN-less Connection String Examples	43
Features	46
Catalog and Schema Support	46
Schema Types and File Formats	46
Specifying Column Names	46
Data Types	48
Casting Binary Data	50
Security and Authentication	51
Driver Configuration Options	53
Configuration Options Appearing in the User Interface	53
Configuration Options Having Only Key Names	65
Advanced Properties	66
Third-Party Trademarks	74
Third-Party Licenses	75

About the Simba Drill ODBC Driver

About Apache Drill

Apache Drill is a low latency distributed query engine capable of querying large datasets from multiple data sources using SQL.

The data sources that Drill supports include file systems (local or distributed) and HBase. Drill also integrates seamlessly with the Hive metastore to complement existing Hive environments with low latency queries. Unlike traditional RDBMS or SQL-on-Hadoop solutions that require centralized schema definitions, Drill can query self-describing data as well as complex or multi-structured data that is commonly seen in big data systems. Moreover, Drill does not require a fully structured schema and can support semi-structured or nested data types such as JSON.

Apache Drill processes the data in record batches and discovers the schema during the processing of each record batch. Thus, Drill has the capability to support changing schemas over the lifetime of a query. Drill reconfigures its operators and handles these situations to ensure that data is not lost.

Apache Drill can support multiple Hive, HBase, and DFS data sources, including CSV, TSV, Parquet, and JSON files. Connections between data sources and Drill can be configured through Drill's web UI.

 **Note:**

For information about connecting Apache Drill to data sources, see the Apache Drill documentation: <http://drill.apache.org/docs/>.

About the Driver

The Simba Drill ODBC Driver lets organizations connect their BI tools to Drill. Drill provides an ANSI SQL query layer and also exposes the metadata information through an ANSI SQL standard metadata database called INFORMATION_SCHEMA. The Simba Drill ODBC Driver leverages INFORMATION_SCHEMA to expose Drill's metadata to BI tools as needed.


The driver complies with the ODBC 3.80 data standard, including important functionality such as Unicode and 32- and 64-bit support for high-performance computing environments on all platforms.

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 *Data Access Standards* on the Simba Technologies website: <https://www.simba.com/resources/data-access-standards-glossary>. For complete information about the ODBC specification, see the *ODBC API Reference* from the Microsoft documentation: <https://docs.microsoft.com/en-us/sql/odbc/reference/syntax/odbc-api-reference>.

The Simba Drill ODBC Driver is available for Microsoft® Windows®, Linux, and macOS platforms.

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

 **Note:**

For information about how to use the driver in various BI tools, see the *Simba ODBC Drivers Quick Start Guide for Windows*: http://cdn.simba.com/docs/ODBC_QuickstartGuide/content/quick_start/intro.htm.

About SQLParse Methods

The SOQL_FIRST and SQL_FIRST parse methods may lead to different behavior for similar queries. This occurs when the driver switches between the two modes trying to find a query language that can support the inputted query. Behavior will be consistent when the query language the driver decides to use remains consistent, however changes to the query may cause it to fail in one of the languages. For example, when using SOQL_FIRST mode you issue a query that is executable as SOQL. In a subsequent transaction, you slightly modify the query and it becomes invalid in SOQL, but is valid in SQL. This results in the first query being executed using SOQL and the second query being executed using SQL. This can cause slight differences between the two result sets since the behavior is not the same for all SOQL and SQL queries.

One instance where this occurs is comparisons involving null values, because SOQL and SQL handle comparisons against null differently. SQL returns an unknown state if a comparison operator (such as = or >) is used with null, and the results contain zero rows. However, SOQL allows such a comparison and returns results.

Using the SOQL_FIRST mode, you issue the query `SELECT Name FROM Account WHERE NumberOfEmployees = NULL`. This query is valid SOQL and the returned values contain all non-null values as specified by SOQL. Next you issue the query `SELECT Account.Name FROM Account, Contact WHERE Account.Id = Contact.AccountId AND Account.NumberOfEmployees = NULL`. This query is not valid SOQL but is valid SQL. It returns zero values as

specified by the SQL specification. The first query may have lead you to believe that the second query would also return results, but difference in query language used means second query returns no results.

If you require consistent behavior in these types of instances, use either `SOQL_ONLY` or the `SQL_ONLY` mode.

Windows Driver

Windows System Requirements

Install the driver on client machines where the application is installed. Each machine that you install the driver on must meet the following minimum system requirements:

- One of the following operating systems:
 - Windows 10, 8.1, or 7 SP1
 - Windows Server 2016, 2012, or 2008 R2 SP1
- 75 MB of available disk space
- Visual C++ Redistributable for Visual Studio 2013 installed (with the same bitness as the driver that you are installing).
You can download the installation packages at <https://www.microsoft.com/en-ca/download/details.aspx?id=40784>.

To install the driver, you must have administrator privileges on the machine.

Installing the Driver on Windows

On 64-bit Windows operating systems, you can execute both 32- and 64-bit applications. However, 64-bit applications must use 64-bit drivers, and 32-bit applications must use 32-bit drivers. Make sure that you use a driver whose bitness matches the bitness of the client application:

- `SimbaDrillODBC32.msi` for 32-bit applications
- `SimbaDrillODBC64.msi` for 64-bit applications

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

To install the Simba Drill ODBC Driver on Windows:

1. Depending on the bitness of your client application, double-click to run **SimbaDrillODBC32.msi** or **SimbaDrillODBC64.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**.

7. If you received a license file through email, then copy the license file into the `\lib` subfolder of the installation folder you selected above. You must have Administrator privileges when changing the contents of this folder.


Creating a Data Source Name on Windows

Typically, after installing the Simba Drill ODBC Driver, you need to create a Data Source Name (DSN).

Alternatively, for information about DSN-less connections, see [Using a Connection String](#) on page 43.

To create a Data Source Name on Windows:

1. From the Start menu, go to **ODBC Data Sources**.

 **Note:**

Make sure to select the ODBC Data Source Administrator that has the same bitness as the client application that you are using to connect to Drill.

2. In the ODBC Data Source Administrator, click the **Drivers** tab, and then scroll down as needed to confirm that the Simba Drill ODBC Driver appears in the alphabetical list of ODBC drivers that are installed on your system.
3. Choose one:
 - 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.

 **Note:**

It is recommended that you create a System DSN instead of a User DSN. Some applications load the data using a different user account, and might not be able to detect User DSNs that are created under another user account.

4. Click **Add**.
5. In the Create New Data Source dialog box, select **Simba Drill ODBC Driver** and then click **Finish**. The Simba Drill ODBC Driver DSN Setup dialog box opens.
6. In the **Data Source Name** field, type a name for your DSN.
7. Optionally, in the **Description** field, type relevant details about the DSN.
8. Choose one:
 - To connect to a ZooKeeper cluster, select **ZooKeeper Quorum**, then type a comma-separated list of the servers in your ZooKeeper cluster in the

Quorum field, and then type the unique name of the Drillbit cluster that you want to use in the **Cluster ID** field.

- Or, to connect to a Drill server, select **Direct to Drillbit**, then type the IP address or host name of the Drill server in the field beside the **Direct to Drillbit** option, and then type the port on which the Drill server is listening in the field beside the colon (:).
9. If the server that you are connecting to requires authentication, then use the options in the Authentication area to configure authentication as needed. For more information, see [Configuring Authentication on Windows](#) on page 12.

 **Note:**

If the Authentication Type drop-down list is set to No Authentication, the driver uses a default user name of "anonymous".

10. If operations against Drill are to be done on behalf of a user other than the one associated with the connection, then do one of the following:
- If you configured Plain authentication, then in the **Delegation UID** field, type the name of the user to perform the operations.
 - Or, if you did not configure authentication, then do the following:
 - a. From the **Authentication Type** drop-down list, select **Plain**.
 - b. In the **User** field, type any value. The User option must be set, but the value is not used.
 - c. In the **Delegation UID** field, type the name of the user to perform the operations.
11. In the **Catalog** list, select or type the name of the synthetic catalog that you want to use for the connection. For more information, see [Catalog and Schema Support](#) on page 46 and [Catalog](#) on page 55.
12. In the **Default Schema** list, select or type the name of the default schema for the DSN to use.
13. Optionally, edit the string in the **Advanced Properties** field to set the time zone, timeouts, and schemas to exclude. For more information, see [Advanced Properties](#) on page 66.
14. To disable support for asynchronous queries, select the **Disable Async** check box.
15. To configure logging behavior for the driver, click **Logging Options**. For more information, see [Configuring Logging Options on Windows](#) on page 15.
16. To examine the data source and create SQL views based on the data source, click **Drill Explorer**. For more information, see [Using Drill Explorer on Windows](#) on page 17 and [Features](#) on page 46.
17. 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 Simba Drill ODBC Driver DSN Setup dialog box are correct. Contact your Drill server administrator as needed.

18. To save your settings and close the Simba Drill ODBC Driver DSN Setup dialog box, click **OK**.
19. To close the ODBC Data Source Administrator, click **OK**.

Configuring Authentication on Windows

Some Drill databases require authentication. You can configure the Simba Drill ODBC Driver to authenticate the connection to the database using one of several methods. For more information, see the following sections:

- [Using Your Drill User Name and Password](#) on page 12
- [Using Kerberos](#) on page 13
- [Using MapR-SASL](#) on page 13
- [Using Your Drill Username Only](#) on page 14

 **Note:**

If the Authentication Type drop-down list is set to No Authentication, the driver uses a default user name of "anonymous".

Using Your Drill User Name and Password

You can configure the driver to use your Drill data store credentials to authenticate the connection.

To configure user name and password authentication on Windows:

1. To access authentication options, open the ODBC Data Source Administrator where you created the DSN, select the DSN, and then click **Configure**.
2. From the **Authentication Type** drop-down list, select **Plain**.
3. In the **User** field, type an appropriate user name for accessing the Drill server.
4. In the **Password** field, type the password corresponding to the user name you typed above.
5. To save your settings and close the dialog box, click **OK**.

Using Kerberos

You can configure the driver to use the Kerberos protocol to authenticate the connection.

If you do not enable the Use Only SSPI option, then Kerberos must be installed and configured before you can use this authentication mechanism. For information about how to install and configure Kerberos, see the MIT Kerberos Documentation: <http://web.mit.edu/kerberos/krb5-latest/doc/>.

To configure Kerberos authentication on Windows:

1. To access authentication options, open the ODBC Data Source Administrator where you created the DSN, select the DSN, and then click **Configure**.
2. From the **Authentication Type** drop-down list, select **Kerberos**.
3. Optionally, in the **Host FQDN** field, type the fully qualified domain name (FQDN) of the Drill server host. If you leave this field empty, the driver uses the host name of the server that you are connecting to as the FQDN for Kerberos authentication.
4. Optionally, in the **Service Name** field, type the Kerberos service principal name of the Drill server. If you leave this field empty, the driver uses `drill` as the service principal name.
5. Optionally, do one of the following to specify which mechanism the driver uses by default to handle Kerberos authentication:
 - To use the SSPI plugin by default, select the **Use Only SSPI** check box.
 - To use MIT Kerberos by default and only use the SSPI plugin if the GSSAPI library is not available, clear the **Use Only SSPI** check box.
6. To require that users define a Kerberos host and service name, set the **Kerberos SPN Configuration Required** check box.
7. To save your settings and close the dialog box, click **OK**.

Using MapR-SASL

You can configure the driver to use the MapR-SASL protocol to authenticate the connection.

The `maprlogin` utility must be installed and configured before you can use this authentication mechanism. For more information, see the *MapR Security Guide*: <http://maprdocs.mapr.com/51/SecurityGuide/SecurityOverview.html>.

To configure MapR-SASL authentication on Windows:

1. Use the `maprlogin` utility to acquire a `maprticket`. For more information, see "Logging Into a Cluster with maprlogin" in the *MapR Security Guide*: <http://maprdocs.mapr.com/51/SecurityGuide/LoggingIntoCluster.html>.

2. To access authentication options for the driver, open the ODBC Data Source Administrator where you created the DSN, select the DSN, and then click **Configure**.
3. From the **Authentication Type** drop-down list, select **MapRSASL**.
4. To save your settings and close the dialog box, click **OK**.

Using Your Drill Username Only

This method is intended for servers without any authentication enabled and with (or without) impersonation enabled.

To configure Username only authentication on Windows:

1. To access authentication options for the driver, open the ODBC Data Source Administrator where you created the DSN, select the DSN, and then click **Configure**.
2. From the **Authentication Type** drop-down list, select **Username**.
3. In the **User** field, type an appropriate user name for accessing the Drill server.
4. Optionally, in the **Delegation UID** field enter the details for any relevant impersonation target.
5. To save your settings and close the dialog box, click **OK**.

Configuring TLS (SSL) on a Windows Machine

You can configure the driver to use Transport Layer Security (TLS, formerly called SSL) when connecting to a host.

To configure TLS:

1. To access TLS options, open the ODBC Data Source Administrator where you created the DSN, select the DSN, and then click **Configure**.
2. Click **SSL Options**. The SSL Options dialog opens.
3. Select the **Enable SSL** check box.
4. Optionally, if you do not want to verify that the host in the certificate is the host being connected to, select the **Disable Hostname Verification** check box.
5. Optionally, if you do not want to verify the certificate against the local trust store, select the **Disable Certificate Verification** check box.
6. Optionally, if you want to use the Windows trust store to verify certificates, select the **Use System Truststore** check box.
 - If you are not using the system trust store, click **Browse** next to the **Trusted Certificate** field, and browse to the location of your own trusted certificate.

7. Choose one:
 - a. To set a minimum TLS, from the **Minimum TLS** drop-down list select the minimum version of TLS to use when connecting to your data store.
 - b. Or, to set an exact TLS, select **Exact TLS Version**, then type the TLS version you want used when connecting to your data store in the field. Supported options include `tlsv1`, `tlsv11`, or `tlsv12`.
8. To save your settings and close the SSL dialog box, click **OK**.
9. To save your settings and close the dialog box, click **OK**.

Configuring Logging Options on Windows

To help troubleshoot issues, you can enable logging. In addition to functionality provided in the Simba Drill ODBC Driver, the ODBC Data Source Administrator provides tracing functionality.

! Important:

Only enable logging or tracing long enough to capture an issue. Logging or tracing decreases performance and can consume a large quantity of disk space.

The settings for logging apply to every connection that uses the Simba Drill ODBC Driver, so make sure to disable the feature after you are done using it.

To enable driver logging on Windows:

1. To access logging options, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **Logging Options**.
2. From the **Log Level** drop-down list, select the logging level corresponding to the amount of information that you want to include in log files:

Logging Level	Description
OFF	Disables all logging.
FATAL	Logs severe error events that lead the driver to abort.
ERROR	Logs error events that might allow the driver to continue running.

Logging Level	Description
WARNING	Logs events that might result in an error if action is not taken.
INFO	Logs general information that describes the progress of the driver.
DEBUG	Logs detailed information that is useful for debugging the driver.
TRACE	Logs all driver activity.

3. In the **Log Path** field, specify the full path to the folder where you want to save log files.
4. If requested by Technical Support, type the name of the component for which to log messages in the **Log Namespace** field. Otherwise, do not type a value in the field.
5. Click **OK**.
6. Restart your ODBC application to make sure that the new settings take effect.

The Simba Drill ODBC Driver produces the following log files at the location you specify in the Log Path field, where *[DriverName]* is the name of the driver:

- A *[DriverName].log* file that logs driver activity that is not specific to a connection.
- A *Simba[DriverName]_connection_[Number].log* for each connection made to the database, where *[Number]* is a number that identifies each log file. This file logs driver activity that is specific to the connection.

If you enable the `UseLogPrefix` connection property, the driver prefixes the log file name with the user name associated with the connection and the process ID of the application through which the connection is made. For more information, see [UseLogPrefix](#) on page 66.

To disable driver logging on Windows:

1. Open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **Logging Options**.
2. From the **Log Level** drop-down list, select **LOG_OFF**.
3. Click **OK**.
4. Restart your ODBC application to make sure that the new settings take effect.

Using Drill Explorer on Windows

Drill demonstrates the unique ability to query self-describing data in file systems and NoSQL databases like HBase directly without defining metadata definitions in Hive. Therefore, users need a mechanism for exploring the metadata in these environments using SQL queries and for creating views for visualization using BI tools. Drill Explorer is a standalone application designed for use alongside your Business Intelligence application as part of the ODBC DSN experience, allowing you to do the following:

- Browse data in your Drill data source
- Define and save SQL views


When retrieving data from your Drill data store, you can use the SQL views created through Drill Explorer to select the data to include.

For information about the schemas and file formats that the Simba Drill ODBC Driver supports, see [Schema Types and File Formats](#) on page 46.

To run Drill Explorer with a specific DSN:

1. In the ODBC Data Source Administrator, select the tab corresponding to the type of DSN that you created for connecting to the Drill data source, then select your DSN, and then click **Configure**.
2. In the Simba Drill ODBC Driver DSN Setup dialog box, click **Drill Explorer**.

To run Drill Explorer standalone:

1. Open Drill Explorer:
 - If you are using Windows 7 or earlier, click **Start**  > **All Programs** > **Simba Drill ODBC Driver** > **Drill Explorer**.
 - Or, if you are using Windows 8 or later, click the arrow button at the bottom of the Start screen, and then click **Simba Drill ODBC Driver** > **Drill Explorer**.

Note:

Make sure to select the Drill Explorer from the driver program group that has the same bitness as the client application that you are using to connect to Drill.

2. Choose one:
 - To connect to a DSN, in the ODBC Connection dialog box, on the DSN tab, select the DSN to which you want to connect.
 - To manually enter the connection properties, on the Driver tab, select a driver.

3. Click **Connect**. If you selected a driver, then the Simba Drill ODBC Driver Connection dialog box opens. Type the connection properties, and then click **OK**.

To browse data in your Drill data source:

1. In Drill Explorer, click the **Browse** tab.
2. In the Schemas pane, click to expand the schema you want to browse.
3. Continue expanding branches in the schema as needed, and then select the table or file you want to view. The Metadata pane displays the structure of the selected table or file. The Data Preview pane displays the data that the selected table or file contains.
4. To specify the maximum number of rows that can appear in the Data Preview pane, type a value in the field at the bottom of the Data Preview pane and then click **Sample**.
5. To update the data in Drill Explorer after changes have been made in the data source, click the **Refresh** button at the bottom of the Schemas pane.

! Important:

The Refresh command updates every schema in the Schemas pane. The scope of the updates is not affected by what you have selected.

6. After you select a table or file using the Browse tab, click the **SQL** tab to display the SQL statement that returns the data currently appearing in the Browse tab.

To define and save a SQL view:

1. In Drill Explorer, click the **SQL** tab.
2. In the **View Definition SQL** field, type a SQL statement defining the desired view or modify the SQL statement appearing in the field as needed, and then click **Preview** to test the SQL statement.
3. To save the view, click the **Create As** button. In the Create As dialog box, in the Schema list, select the schema where you want to save the view.
4. In the **View Name** field, type a descriptive name.

! Important:

Do not include spaces in the view name.

5. If you are saving over an existing view, select the **Overwrite Existing Views** check box.
6. Click **Save**. A message appears in the dialog box, indicating whether the view saved successfully. To copy the message, click **Copy**.
7. Close the dialog box.

Verifying the Driver Version Number on Windows

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

To verify the driver version number on Windows:

1. From the Start menu, go to **ODBC Data Sources**.

 **Note:**

Make sure to select the ODBC Data Source Administrator that has the same bitness as the client application that you are using to connect to Drill.

2. Click the **Drivers** tab and then find the Simba Drill ODBC Driver in the list of ODBC drivers that are installed on your system. The version number is displayed in the **Version** column.

macOS Driver

macOS System Requirements

Install the driver on client machines where the application is installed. Each client machine that you install the driver on must meet the following minimum system requirements:

- macOS version 10.11, 10.12, or 10.13
- 60 MB of available disk space
- iODBC 3.52.9, 3.52.10, 3.52.11, or 3.52.12

Installing the Driver on macOS

The Simba Drill ODBC Driver is available for macOS as a `.dmg` file named `SimbaDrillODBC.dmg`. The driver supports both 32- and 64-bit client applications.

To install the Simba Drill ODBC Driver on macOS:

1. Double-click **SimbaDrillODBC.dmg** to mount the disk image.
2. Double-click **SimbaDrillODBC.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**.

 **Note:**

By default, the driver files are installed in the `/Library/simba/drill` directory.

6. To accept the installation location and begin the installation, click **Install**.
7. When the installation completes, click **Close**.
8. If you received a license file through email, then copy the license file into the `/lib` subfolder in the driver installation directory. You must have root privileges when changing the contents of this folder.

For example, if you installed the driver to the default location, you would copy the license file into the `/Library/simba/drill/lib` folder.

Next, configure the environment variables on your machine to make sure that the ODBC driver manager can work with the driver. For more information, see [Configuring the ODBC Driver Manager on Non-Windows Machines](#) on page 29.

Using Drill Explorer on macOS

Drill demonstrates the unique ability to query self-describing data in file systems and NoSQL databases like HBase directly without defining metadata definitions in Hive. Therefore, users need a mechanism for exploring the metadata in these environments using SQL queries and for creating views for visualization using BI tools. Drill Explorer is a standalone application designed for use alongside your Business Intelligence application as part of the ODBC DSN experience, allowing you to do the following:

- Browse data in your Drill data source
- Define and save SQL views

When retrieving data from your Drill data store, you can use the SQL views created through Drill Explorer to select the data to include.

For information about the schemas and file formats that the Simba Drill ODBC Driver supports, see [Schema Types and File Formats](#) on page 46.

! Important:

- Prior to using Drill Explorer, ensure that you have created the necessary ODBC connections as described in [Configuring ODBC Connections on a Non-Windows Machine](#) on page 32.
- The Drill Explorer requires iODBC. Ensure that the `LD_LIBRARY_PATH` environment variable includes the path to the `libiodbcinst.so` file.

To run Drill Explorer:

1. Choose one:
 - On your macOS machine, in the Applications folder, double-click **Drill Explorer**.
 - Or, navigate to the `/Library/simba/drill/DrillExplorer` directory and run **DrillExplorer.app**.
2. In Drill Explorer, click **Connect**.
3. Choose one:
 - If you are connecting through a DSN, then click the **ODBC DSN** tab, select your DSN, and then click **Connect**.
 - Or, if you are using a DSN-less connection, then click the **Advanced** tab, type your connection string in the field, and then click **Connect**. For

information about DSN-less connection strings, see [Using a Connection String](#) on page 43.

4. If your connection is configured to use Plain authentication, then Drill Explorer prompts you to provide your credentials. Enter your user name and password, and then click **OK**.
5. To disconnect from a DSN, select the DSN in the Schemas pane and then click **Disconnect**.

To browse data in your Drill data source:

1. In Drill Explorer, in the Schemas pane, click to expand the schema that you want to browse.
2. Continue expanding branches in the schema as needed, and then select the table or file you want to view.

Note:

When you select an item in the Schemas pane, the Browse tab displays and the SQL tab becomes available. In the Browse tab, the Metadata pane displays the structure of the selected table or file. The Preview pane displays the data that the selected table or file contains. To see the SQL statement that returns the data currently appearing in the Browse tab, click the **SQL** tab.

3. To specify the maximum number of rows that can appear in the Preview pane, type a value in the field at the bottom of the Preview pane, and then click **Sample**.
4. To update the data in Drill Explorer after changes have been made in the data source, in the Schemas pane, right-click the item that you want to update, and then click **Refresh**.

Important:

The Refresh command does not update everything that appears in the Schemas pane. The scope of the Refresh command depends on the item that you right-click to access the command:

- If you right-click a table, view, or file, then Drill Explorer updates everything that is in the same schema or directory as the selected table, view, or file.
- If you right-click a schema or directory, then Drill Explorer updates everything in the selected schema or directory. Other schemas or directories on the same hierarchy level are not affected.

To define and save a SQL view:

1. In Drill Explorer, select an item in the Schemas pane, and then click the **SQL** tab. The SQL tab displays the SQL statement that returns the selected data.
2. In the **View Definition SQL** field, type a SQL statement defining the desired view or modify the SQL statement appearing in the field as needed, and then click **Preview** to test the SQL statement.
3. To save the view, click **Create As**. In the Create As dialog box, in the **Schema** list, select the schema where you want to save the view.
4. In the **View Name** field, type a descriptive name.

! Important:

Do not include spaces in the view name.

5. If you are saving over an existing view, select the **Overwrite Existing Views** check box.
6. Click **Save**. A message appears in the dialog box indicating whether the view saved successfully. To copy the message, click **Copy**.
7. Close the dialog box.

Verifying the Driver Version Number on macOS

If you need to verify the version of the Simba Drill ODBC Driver that is installed on your macOS machine, you can query the version number through the Terminal.

To verify the driver version number on macOS:

- At the Terminal, run the following command:

```
pkgutil --info com.simba.drillodbc
```

The command returns information about the Simba Drill ODBC Driver that is installed on your machine, including the version number.

Linux Driver

Linux System Requirements

Install the driver on client machines where the application is installed. Each client machine that you install the driver on must meet the following minimum system requirements:

- One of the following distributions:
 - Red Hat® Enterprise Linux® (RHEL) 6 or 7
 - CentOS 6 or 7
 - SUSE Linux Enterprise Server (SLES) 11 or 12
 - Debian 8 or 9
 - Ubuntu 14.04, 16.04, or 18.04
- 90 MB of available disk space
- One of the following ODBC driver managers installed:
 - iODBC 3.52.9, 3.52.10, 3.52.11, or 3.52.12
 - unixODBC 2.3.2, 2.3.3, or 2.3.4

To install the driver, you must have root access on the machine.

Installing the Driver Using the RPM File

On 64-bit editions of Linux, you can execute both 32- and 64-bit applications. However, 64-bit applications must use 64-bit drivers, and 32-bit applications must use 32-bit drivers. Make sure that you use a driver whose bitness matches the bitness of the client application:

- `SimbaDrillODBC-32bit-[Version]-[Release].[LinuxDistro].i686.rpm` for the 32-bit driver
- `SimbaDrillODBC-[Version]-[Release].[LinuxDistro].x86_64.rpm` for the 64-bit driver

The placeholders in the file names are defined as follows:

- `[Version]` is the version number of the driver.
- `[Release]` is the release number for this version of the driver.
- `[LinuxDistro]` is one of the following values indicating which Linux distribution the installer is optimized for:

Value	Platform for Installation
e15	CentOS 5 RHEL 5
e16	CentOS 6 RHEL 6
None	SLES 11 or 12

Make sure to install the driver using the RPM that is optimized for your Linux distribution. Otherwise, you may encounter errors when using the driver.

You can install both the 32-bit and 64-bit versions of the driver on the same machine.

To install the Simba Drill ODBC Driver using the RPM File:

1. Log in as the root user.
2. Navigate to the folder containing the RPM package for the driver.
3. Depending on the Linux distribution that you are using, run one of the following commands from the command line, where *[RPMFileName]* is the file name of the RPM package:

- If you are using Red Hat Enterprise Linux or CentOS, run the following command:

```
yum --nogpgcheck localinstall [RPMFileName]
```

- Or, if you are using SUSE Linux Enterprise Server, run the following command:

```
zypper install [RPMFileName]
```

The Simba Drill ODBC Driver files are installed in the `/opt/simba/drill` directory.

4. If you received a license file through email, then copy the license file into the `/opt/simba/drill/lib/32` or `/opt/simba/drill/lib/64` folder, depending on the version of the driver that you installed. You must have root privileges when changing the contents of this folder.

Next, configure the environment variables on your machine to make sure that the ODBC driver manager can work with the driver. For more information, see [Configuring the ODBC Driver Manager on Non-Windows Machines](#) on page 29.

Using Drill Explorer on Linux

Drill demonstrates the unique ability to query self-describing data in file systems and NoSQL databases like HBase directly without defining metadata definitions in Hive. Therefore, users need a mechanism for exploring the metadata in these environments using SQL queries and for creating views for visualization using BI tools. Drill Explorer is a standalone application designed for use alongside your Business Intelligence application as part of the ODBC DSN experience, allowing you to do the following:

- Browse data in your Drill data source
- Define and save SQL views

When retrieving data from your Drill data store, you can use the SQL views created through Drill Explorer to select the data to include.

For information about the schemas and file formats that the Simba Drill ODBC Driver supports, see [Schema Types and File Formats](#) on page 46.

! Important:

- Prior to using Drill Explorer, ensure that you have created the necessary ODBC connections as described in [Configuring ODBC Connections on a Non-Windows Machine](#) on page 32.
- The Drill Explorer requires iODBC. Ensure that the LD_LIBRARY_PATH environment variable includes the path to the `libiodbcinst.so` file.

To run Drill Explorer:

1. Choose one:
 - Either click **Applications**, then click **Other**, and then click **Drill Explorer**.
 - Or navigate to the `opt/simba/drill/DrillExplorer` directory and run **DrillExplorer**.
2. In Drill Explorer, click **Connect**.
3. Choose one:
 - If you are connecting through a DSN, then click the **ODBC DSN** tab, select your DSN, and then click **Connect**.
 - Or, if you are using a DSN-less connection, then click the **Advanced** tab, type your connection string in the field, and then click **Connect**. For information about DSN-less connection strings, see [Using a Connection String](#) on page 43.
4. If your connection requires authentication, type your user name and password.
5. Click **OK**.

6. To disconnect from a DSN, select the DSN in the Schemas pane and then click **Disconnect**

To browse data in your Drill data source:

1. In Drill Explorer, in the Schemas pane, click to expand the schema that you want to browse.
2. Continue expanding branches in the schema as needed, and then select the table or file you want to view.

Note:

When you select an item in the Schemas pane, the Browse tab displays and the SQL tab becomes available. In the Browse tab, the Metadata pane displays the structure of the selected table or file. The Preview pane displays the data that the selected table or file contains. To see the SQL statement that returns the data currently appearing in the Browse tab, click the **SQL** tab.

3. To specify the maximum number of rows that can appear in the Preview pane, type a value in the field at the bottom of the Preview pane, and then click **Sample**.
4. To update the data in Drill Explorer after changes have been made in the data source, in the Schemas pane, right-click the item that you want to update, and then click **Refresh**.

Important:

The Refresh command does not update everything that appears in the Schemas pane. The scope of the Refresh command depends on the item that you right-click to access the command:

- If you right-click a table, view, or file, then Drill Explorer updates everything that is in the same schema or directory as the selected table, view, or file.
- If you right-click a schema or directory, then Drill Explorer updates everything in the selected schema or directory. Other schemas or directories on the same hierarchy level are not affected.

To define and save a SQL view:

1. In Drill Explorer, select an item in the Schemas pane, and then click the **SQL** tab. The SQL tab displays the SQL statement that returns the selected data.
2. In the **View Definition SQL** field, type a SQL statement defining the desired view or modify the SQL statement appearing in the field as needed, and then click **Preview** to test the SQL statement.

3. To save the view, click **Create As**. In the Create As dialog box, in the **Schema** list, select the schema where you want to save the view.
4. In the **View Name** field, type a descriptive name.

! Important:

Do not include spaces in the view name.

5. If you are saving over an existing view, select the **Overwrite Existing Views** check box.
6. Click **Save**. A message appears in the dialog box indicating whether the view saved successfully. To copy the message, click **Copy**.
7. Close the dialog box.

Verifying the Driver Version Number on Linux

If you need to verify the version of the Simba Drill ODBC Driver that is installed on your Linux machine, you can query the version number through the command-line interface if the driver was installed using an RPM file.

To verify the driver version number on Linux:

- Depending on your package manager, at the command prompt, run one of the following commands:

- `yum list | grep SimbaDrillODBC`
- `rpm -qa | grep SimbaDrillODBC`

The command returns information about the Simba Drill ODBC Driver that is installed on your machine, including the version number.

Configuring the ODBC Driver Manager on Non-Windows Machines

To make sure that the ODBC driver manager on your machine is configured to work with the Simba Drill ODBC Driver, do the following:

- Set the library path environment variable to make sure that your machine uses the correct ODBC driver manager. For more information, see [Specifying ODBC Driver Managers on Non-Windows Machines](#) on page 29.
- If the driver configuration files are not stored in the default locations expected by the ODBC driver manager, then set environment variables to make sure that the driver manager locates and uses those files. For more information, see [Specifying the Locations of the Driver Configuration Files](#) on page 30.

After configuring the ODBC driver manager, you can configure a connection and access your data store through the driver. For more information, see [Configuring ODBC Connections on a Non-Windows Machine](#) on page 32.

Specifying ODBC Driver Managers on Non-Windows Machines

You need to make sure that your machine uses the correct ODBC driver manager to load the driver. To do this, set the library path environment variable.

macOS

If you are using a macOS machine, then set the `DYLD_LIBRARY_PATH` environment variable to include the paths to the ODBC driver manager libraries. For example, if the libraries are installed in `/usr/local/lib`, then run the following command to set `DYLD_LIBRARY_PATH` for the current user session:

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

For information about setting an environment variable permanently, refer to the macOS shell documentation.

Linux

If you are using a Linux machine, then set the `LD_LIBRARY_PATH` environment variable to include the paths to the ODBC driver manager libraries. For example, if the libraries are installed in `/usr/local/lib`, then run the following command to set `LD_LIBRARY_PATH` for the current user session:

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

For information about setting an environment variable permanently, refer to the Linux shell documentation.

Specifying the Locations of the Driver Configuration Files

By default, ODBC driver managers are configured to use hidden versions of the `odbc.ini` and `odbcinst.ini` configuration files (named `.odbc.ini` and `.odbcinst.ini`) located in the home directory, as well as the `simba.drillodbc.ini` file in the `lib` subfolder of the driver installation directory. If you store these configuration files elsewhere, then you must set the environment variables described below so that the driver manager can locate the files.

If you are using iODBC, do the following:

- Set `ODBCINI` to the full path and file name of the `odbc.ini` file.
- Set `ODBCINSTINI` to the full path and file name of the `odbcinst.ini` file.
- Set `SIMBADRILLINI` to the full path and file name of the `simba.drillodbc.ini` file.

If you are using unixODBC, do the following:

- Set `ODBCINI` to the full path and file name of the `odbc.ini` file.
- Set `ODBCSYSINI` to the full path of the directory that contains the `odbcinst.ini` file.
- Set `SIMBADRILLINI` to the full path and file name of the `simba.drillodbc.ini` file.

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

For iODBC:

```
export ODBCINI=/usr/local/odbc/odbc.ini
export ODBCINSTINI=/usr/local/odbc/odbcinst.ini
export SIMBADRILLINI=/etc/simba.drillodbc.ini
```

For unixODBC:

```
export ODBCINI=/usr/local/odbc/odbc.ini
export ODBCSYSINI=/usr/local/odbc
```

```
export SIMBADRILLINI=/etc/simba.drillodbc.ini
```

To locate the `simba.drillodbc.ini` file, the driver uses the following search order:

1. If the `SIMBADRILLINI` environment variable is defined, then the driver searches for the file specified by the environment variable.
2. The driver searches the directory that contains the driver library files for a file named `simba.drillodbc.ini`.
3. The driver searches the current working directory of the application for a file named `simba.drillodbc.ini`.
4. The driver searches the home directory for a hidden file named `.simba.drillodbc.ini` (prefixed with a period).
5. The driver searches the `/etc` directory for a file named `simba.drillodbc.ini`.

Configuring ODBC Connections on a Non-Windows Machine

The following sections describe how to configure ODBC connections when using the Simba Drill ODBC Driver on non-Windows platforms:

- [Creating a Data Source Name on a Non-Windows Machine](#) on page 32
- [Configuring a DSN-less Connection on a Non-Windows Machine](#) on page 35
- [Configuring Authentication on a Non-Windows Machines](#) on page 37
- [Configuring Logging Options on a Non-Windows Machine](#) on page 39
- [Testing the Connection on a Non-Windows Machine](#) on page 41

Creating a Data Source Name on a Non-Windows Machine

When connecting to your data store using a DSN, you only need to configure the `odbc.ini` file. Set the properties in the `odbc.ini` file to create a DSN that specifies the connection information for your data store. For information about configuring a DSN-less connection instead, see [Configuring a DSN-less Connection on a Non-Windows Machine](#) on page 35.

If your machine is already configured to use an existing `odbc.ini` file, then update that file by adding the settings described below. Otherwise, copy the `odbc.ini` file from the `Setup` subfolder in the driver installation directory to the home directory, and then update the file as described below.

To create a Data Source Name on a non-Windows machine:

1. In a text editor, open the `odbc.ini` configuration file.

 **Note:**

If you are using a hidden copy of the `odbc.ini` file, you can remove the period (.) from the start of the file name to make the file visible while you are editing it.

2. In the `[ODBC Data Sources]` section, add a new entry by typing a name for the DSN, an equal sign (=), and then the name of the driver.

For example, on a macOS machine:

```
[ODBC Data Sources]
```



```
Sample DSN=Simba Drill ODBC Driver
```

As another example, for a 32-bit driver on a Linux machine:

```
[ODBC Data Sources]
Sample DSN=Simba Drill ODBC Driver 32-bit
```

3. Create a section that has the same name as your DSN, and then specify configuration options as key-value pairs in the section:
 - a. Set the `Driver` property to the full path of the driver library file that matches the bitness of the application.

For example, on a macOS machine:

```
Driver=/Library/simba/drill/lib/libdrillodbc_
sbu.dylib
```

As another example, for a 32-bit driver on a Linux machine:

```
Driver=/opt/simba/drill/lib/32/libdrillodbc_sb32.so
```

- b. Do one of the following:
 - To connect to a Drill server, set the `ConnectionType` property to `Direct`, the `Host` property to the IP address or host name of the server, and the `Port` property to the TCP port that the server uses to listen for client connections.

For example:

```
ConnectionType=Direct
Host=192.168.222.160
Port=31010
```

- Or, to connect to a ZooKeeper cluster, set the `ConnectionType` property to `ZooKeeper`, the `ZKClusterID` property to the name of the ZooKeeper cluster, and the `ZKQuorum` property to a comma-separated list of the servers in the cluster, specifying the host names or IP addresses and port numbers.

For example:

```
ConnectionType=ZooKeeper
ZKClusterID=drill
ZKQuorum=192.168.222.160:31010,
192.168.222.165:31010,192.168.222.231:31010
```

- c. If authentication is required to access the server, then specify the authentication mechanism and your credentials. For more information, see [Configuring Authentication on a Non-Windows Machines](#) on page 37.

**Note:**

If you set `AuthenticationType` to `No Authentication`, the driver uses a default user name of "anonymous".

- d. Optionally, set additional key-value pairs as needed to specify other optional connection settings. For detailed information about all the configuration options supported by the Simba Drill ODBC Driver, see [Driver Configuration Options](#) on page 53.
4. Save the `odbc.ini` configuration file.

**Note:**

If you are storing this file in its default location in the home directory, then prefix the file name with a period (.) so that the file becomes hidden. If you are storing this file in another location, then save it as a non-hidden file (without the prefix), and make sure that the `ODBCINI` environment variable specifies the location. For more information, see [Specifying the Locations of the Driver Configuration Files](#) on page 30.

For example, the following is an `odbc.ini` configuration file for macOS containing a DSN that connects to a Drill server without authentication:

```
[ODBC Data Sources]
Sample DSN=Simba Drill ODBC Driver
[Sample DSN]
Driver=/Library/simba/drill/lib/libdrillodbc_sbu.dylib
ConnectionType=Direct
Host=192.168.222.160
Port=31010
```

As another example, the following is an `odbc.ini` configuration file for a 32-bit driver on a Linux machine, containing a DSN that connects to a Drill server without authentication:

```
[ODBC Data Sources]
Sample DSN=Simba Drill ODBC Driver 32-bit
[Sample DSN]
Driver=/opt/simba/drill/lib/32/libdrillodbc_sb32.so
ConnectionType=Direct
Host=192.168.222.160
```

```
Port=31010
```

You can now use the DSN in an application to connect to the data store.

Configuring a DSN-less Connection on a Non-Windows Machine

To connect to your data store through a DSN-less connection, you need to define the driver in the `odbcinst.ini` file and then provide a DSN-less connection string in your application.

If your machine is already configured to use an existing `odbcinst.ini` file, then update that file by adding the settings described below. Otherwise, copy the `odbcinst.ini` file from the `Setup` subfolder in the driver installation directory to the home directory, and then update the file as described below.

To define a driver on a non-Windows machine:

1. In a text editor, open the `odbcinst.ini` configuration file.

 **Note:**

If you are using a hidden copy of the `odbcinst.ini` file, you can remove the period (.) from the start of the file name to make the file visible while you are editing it.

2. In the `[ODBC Drivers]` section, add a new entry by typing a name for the driver, an equal sign (=), and then `Installed`.

For example:

```
[ODBC Drivers]
Simba Drill ODBC Driver=Installed
```

3. Create a section that has the same name as the driver (as specified in the previous step), and then specify the following configuration options as key-value pairs in the section:
 - a. Set the `Driver` property to the full path of the driver library file that matches the bitness of the application.

For example, on a macOS machine:

```
Driver=/Library/simba/drill/lib/libdrillodbc_
sbu.dylib
```

As another example, for a 32-bit driver on a Linux machine:


```
Driver=/opt/simba/drill/lib/32/libdrillodbc_sb32.so
```

- b. Optionally, set the `Description` property to a description of the driver.

For example:

```
Description=Simba Drill ODBC Driver
```

4. Save the `odbcinst.ini` configuration file.

 **Note:**

If you are storing this file in its default location in the home directory, then prefix the file name with a period (.) so that the file becomes hidden. If you are storing this file in another location, then save it as a non-hidden file (without the prefix), and make sure that the `ODBCINSTINI` or `ODBCSYSINI` environment variable specifies the location. For more information, see [Specifying the Locations of the Driver Configuration Files](#) on page 30.

For example, the following is an `odbcinst.ini` configuration file for macOS:

```
[ODBC Drivers]
Simba Drill ODBC Driver=Installed
[Simba Drill ODBC Driver]
Description=Simba Drill ODBC Driver
Driver=/Library/simba/drill/lib/libdrillodbc_sbu.dylib
```

As another example, the following is an `odbcinst.ini` configuration file for both the 32- and 64-bit drivers on Linux:

```
[ODBC Drivers]
Simba Drill ODBC Driver 32-bit=Installed
Simba Drill ODBC Driver 64-bit=Installed
[Simba Drill ODBC Driver 32-bit]
Description=Simba Drill ODBC Driver (32-bit)
Driver=/opt/simba/drill/lib/32/libdrillodbc_sb32.so
[Simba Drill ODBC Driver 64-bit]
Description=Simba Drill ODBC Driver (64-bit)
Driver=/opt/simba/drill/lib/64/libdrillodbc_sb64.so
```

You can now connect to your data store by providing your application with a connection string where the `Driver` property is set to the driver name specified in the `odbcinst.ini` file, and all the other necessary connection properties are also set.

For more information, see "DSN-less Connection String Examples" in [Using a Connection String](#) on page 43.

For instructions about configuring authentication, see [Configuring Authentication on a Non-Windows Machines](#) on page 37.

For detailed information about all the connection properties that the driver supports, see [Driver Configuration Options](#) on page 53.

Configuring Authentication on a Non-Windows Machines

Some Drill databases require authentication. You can configure the Simba Drill ODBC Driver to authenticate the connection to the database using one of several methods. For more information, see the following sections:

- [Using Your Drill User Name and Password](#) on page 37
- [Using Kerberos](#) on page 38
- [Using MapR-SASL](#) on page 38
- [Using Your Drill Username Only](#) on page 38

Note:

If you set `AuthenticationType` to `No Authentication`, the driver uses a default user name of "anonymous".

You can set the connection properties described below in a connection string or in a DSN (in the `odbc.ini` file). Settings in the connection string take precedence over settings in the DSN.

Using Your Drill User Name and Password

You can configure the driver to use your Drill data store credentials to authenticate the connection.

To configure user name and password authentication on a non-Windows Machine:

1. Set the `AuthenticationType` property to `Plain` or `Basic Authentication`.
2. Set the `UID` property to an appropriate user name for accessing the Drill server.
3. Set the `PWD` property to the password corresponding to the user name you typed above.

Using Kerberos

You can configure the driver to use the Kerberos protocol to authenticate the connection.

Kerberos must be installed and configured before you can use this authentication mechanism. For information about how to install and configure Kerberos, see the MIT Kerberos Documentation: <http://web.mit.edu/kerberos/krb5-latest/doc/>.

To configure Kerberos authentication on a non-Windows Machine:

1. Set the `AuthenticationType` property to `Kerberos`.
2. To require that users set a Kerberos service name and host, set `krbSpnConfigurationsRequired` to 1.
3. Optionally, set the `KrbServiceHost` property to the fully qualified domain name of the Drill server host. If you do not set this property, the driver uses the host name of the server that you are connecting to as the FQDN for Kerberos authentication.
4. Optionally, set the `KrbServiceName` property to the Kerberos service principal name of the Drill server. If you do not set this property, the driver uses `drill` as the service principal name.

Using MapR-SASL

You can configure the driver to use the MapR-SASL protocol to authenticate the connection.

The `maprlogin` utility must be installed and configured before you can use this authentication mechanism. For more information, see the *MapR Security Guide*: <http://maprdocs.mapr.com/51/SecurityGuide/SecurityOverview.html>.

To configure MapR-SASL authentication on a non-Windows Machine:

1. Use the `maprlogin` utility to acquire a `maprticket`. For more information, see "Logging Into a Cluster with `maprlogin`" in the *MapR Security Guide*: <http://maprdocs.mapr.com/51/SecurityGuide/LoggingIntoCluster.html>.
2. Set the `AuthenticationType` property to `MapRSASL`.

Using Your Drill Username Only

This method is intended for servers without any authentication enabled and with (or without) impersonation enabled.

To configure Username only authentication on a non-Windows Machine:

1. Set the `AuthenticationType` property to `Username`.
2. Set the `UID` property to an appropriate user name for accessing the Drill server.
3. Optionally, set the `DelegationUID` to any applicable impersonation target.

Configuring TLS (SSL) on a non-Windows Machine

You can configure the driver to use Transport Layer Security (TLS, formerly called SSL) when connecting to a host.

To configure TLS:

1. Set the `SSL` property to `1`.
2. Optionally, if you do not want to verify the host in the certificate is the host being connected to, set the `disableHostVerification` property to `1`.
3. Optionally, if you do not want to verify the certificate against the local trust store, set the `disableCertificateVerification` property to `1`.
4. Optionally, set the `TrustedCerts` property to the location of the `.pem` file that contains your trusted certificates.
5. Choose one:
 - a. To specify a minimum version of TLS to use, set the `Min_TLS` property to the minimum version of TLS. Supported options include `0` for TLS 1.0, `1` for TLS 1.1, and `2` for TLS 1.2.
 - b. Or, to specify the exact version of TLS to use, set the `UseExactTLSProtocolVersion` property to `1`, then set the `TLSProtocol` property to the version of TLS you want used. Supported options include `tlsv1`, `tlsv11`, or `tlsv12`.

Configuring Logging Options on a Non-Windows Machine

To help troubleshoot issues, you can enable logging in the driver.

! Important:

Only enable logging long enough to capture an issue. Logging decreases performance and can consume a large quantity of disk space.

Logging is configured through driver-wide settings in the `simba.drillodbc.ini` file, which apply to all connections that use the driver.

To enable logging on a non-Windows machine:

1. Open the `simba.drillodbc.ini` configuration file in a text editor.
2. To specify the level of information to include in log files, set the `LogLevel` property to one of the following numbers:

LogLevel Value	Description
0	Disables all logging.
1	Logs severe error events that lead the driver to abort.
2	Logs error events that might allow the driver to continue running.
3	Logs events that might result in an error if action is not taken.
4	Logs general information that describes the progress of the driver.
5	Logs detailed information that is useful for debugging the driver.
6	Logs all driver activity.

3. Set the `LogPath` key to the full path to the folder where you want to save log files.
4. Optionally, to prefix the log file name with the user name and process ID associated with the connection, set the `UseLogPrefix` property to 1.
5. Save the `simba.drillodbc.ini` configuration file.
6. Restart your ODBC application to make sure that the new settings take effect.

The Simba Drill ODBC Driver produces two log files at the location you specify using the `LogPath` key, where `[DriverName]` is the name of the driver:

- A `[DriverName].log` file that logs driver activity that is not specific to a connection.
- A `Simba[DriverName]_connection_[Number].log` for each connection made to the database, where `[Number]` is a number that identifies each log file. This file logs driver activity that is specific to the connection.

If you set the `UseLogPrefix` property to 1, then each file name is prefixed with `[UserName]_[ProcessID]_`, where `[UserName]` is the user name associated with

the connection and `[ProcessID]` is the process ID of the application through which the connection is made. For more information, see [UseLogPrefix](#) on page 66.

To disable logging on a non-Windows machine:

1. Open the `simba.drillodbc.ini` configuration file in a text editor.
2. Set the `LogLevel` key to 0.
3. Save the `simba.drillodbc.ini` configuration file.
4. Restart your ODBC application to make sure that the new settings take effect.

Testing the Connection on a Non-Windows Machine

To test the connection, you can use an ODBC-enabled client application. For a basic connection test, you can also use the test utilities that are packaged with your driver manager installation. For example, the iODBC driver manager includes simple utilities called `iodbctest` and `iodbctestw`. Similarly, the unixODBC driver manager includes simple utilities called `isql` and `iusql`.

Using the iODBC Driver Manager

You can use the `iodbctest` and `iodbctestw` utilities to establish a test connection with your driver. Use `iodbctest` to test how your driver works with an ANSI application, or use `iodbctestw` to test how your driver works with a Unicode application.

Note:

There are 32-bit and 64-bit installations of the iODBC driver manager available. If you have only one or the other installed, then the appropriate version of `iodbctest` (or `iodbctestw`) is available. However, if you have both 32- and 64-bit versions installed, then you need to make sure that you are running the version from the correct installation directory.

For more information about using the iODBC driver manager, see <http://www.iodbc.org>.

To test your connection using the iODBC driver manager:

1. Run **`iodbctest`** or **`iodbctestw`**.
2. Optionally, if you do not remember the DSN, then type a question mark (?) to see a list of available DSNs.
3. Type the connection string for connecting to your data store, and then press ENTER. For more information, see [Using a Connection String](#) on page 43.

If the connection is successful, then the `SQL>` prompt appears.

Using the unixODBC Driver Manager

You can use the `isql` and `iusql` utilities to establish a test connection with your driver and your DSN. `isql` and `iusql` can only be used to test connections that use a DSN. Use `isql` to test how your driver works with an ANSI application, or use `iusql` to test how your driver works with a Unicode application.

 **Note:**

There are 32-bit and 64-bit installations of the unixODBC driver manager available. If you have only one or the other installed, then the appropriate version of `isql` (or `iusql`) is available. However, if you have both 32- and 64-bit versions installed, then you need to make sure that you are running the version from the correct installation directory.

For more information about using the unixODBC driver manager, see <http://www.unixodbc.org>.

To test your connection using the unixODBC driver manager:

➤ Run `isql` or `iusql` by using the corresponding syntax:

- `isql [DataSourceName]`
- `iusql [DataSourceName]`

`[DataSourceName]` is the DSN that you are using for the connection.

If the connection is successful, then the `SQL>` prompt appears.

 **Note:**

For information about the available options, run `isql` or `iusql` without providing a DSN.

Using a Connection String

For some applications, you might need to use a connection string to connect to your data source. For detailed information about how to use a connection string in an ODBC application, refer to the documentation for the application that you are using.

The connection strings in the following sections are examples showing the minimum set of connection attributes that you must specify to successfully connect to the data source. Depending on the configuration of the data source and the type of connection you are working with, you might need to specify additional connection attributes. For detailed information about all the attributes that you can use in the connection string, see [Driver Configuration Options](#) on page 53.

DSN Connection String Example

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

```
DSN= [DataSourceName]
```

[DataSourceName] is the DSN that you are using for the connection.

You can set additional configuration options by appending key-value pairs to the connection string. Configuration options that are passed in using a connection string take precedence over configuration options that are set in the DSN.

DSN-less Connection String Examples

Some applications provide support for connecting to a data source using a driver without a DSN. To connect to a data source without using a DSN, use a connection string instead.

The placeholders in the examples are defined as follows, in alphabetical order:

- *[ClusterName]* is the name of the ZooKeeper cluster to which you are connecting.
- *[PortNumber]* is the number of the TCP port that the Drill server uses to listen for client connections.
- *[Server]* is the IP address or host name of the Drill server to which you are connecting.
- *[YourPassword]* is the password corresponding to your user name.
- *[YourUserName]* is the user name that you use to access the Drill server.

Connecting to a Drillbit

The following is the format of a DSN-less connection string for a Drillbit that does not require authentication:

```
Driver=Simba Drill ODBC Driver;ConnectionType=Direct;  
Host=[Server];Port=[PortNumber]
```

For example:

```
Driver=Simba Drill ODBC Driver;ConnectionType=Direct;  
Host=192.168.222.160;Port=31010
```

Connecting to a ZooKeeper Cluster

The following is the format of a DSN-less connection string for a ZooKeeper cluster that does not require authentication:

```
Driver=Simba Drill ODBC Driver;ConnectionType=ZooKeeper;  
ZKQuorum=[Server1]:[PortNumber1], [Server2]:[PortNumber2],  
[Server3]:[PortNumber3];ZKClusterID=[ClusterName]
```

For example:

```
Driver=Simba Drill ODBC Driver;ConnectionType=ZooKeeper;  
ZKQuorum=192.168.222.160:31010, 192.168.222.165:31010,  
192.168.222.231:31010;ZKClusterID=drill;
```

Connecting to Drill Using Plain Authentication

The following is the format of a DSN-less connection string for a Drillbit that requires Plain authentication:

```
Driver=Simba Drill ODBC Driver;ConnectionType=Direct;  
Host=[Server];Port=[PortNumber];AuthenticationType=Plain;  
UID=[YourUserName];PWD=[YourPassword]
```

For example:

```
Driver=Simba Drill ODBC Driver;ConnectionType=Direct;  
Host=192.168.222.160;Port=31010;  
AuthenticationType=Plain;UID=simba;PWD=simba123
```

Connecting to Drill Using Kerberos Authentication

The following is the format of a DSN-less connection string for a Drillbit that requires Kerberos authentication:

```
Driver=Simba Drill ODBC Driver;ConnectionType=Direct;  
Host=[Server];Port=[PortNumber];AuthenticationType=Kerberos
```

For example:

```
Driver=Simba Drill ODBC Driver;ConnectionType=Direct;  
Host=192.168.222.160;Port=31010;AuthenticationType=Kerberos
```

Connecting to Drill Using MapR-SASL Authentication

The following is the format of a DSN-less connection string for a Drillbit that requires MapR-SASL authentication:

```
Driver=Simba Drill ODBC Driver;ConnectionType=Direct;  
Host=[Server];Port=[PortNumber];AuthenticationType=MapRSASL
```

For example:

```
Driver=Simba Drill ODBC Driver;ConnectionType=Direct;  
Host=192.168.222.160;Port=31010;AuthenticationType=MapRSASL
```

Features

For more information on the features of the Simba Drill ODBC Driver, see the following:

- [Catalog and Schema Support](#) on page 46
- [Schema Types and File Formats](#) on page 46
- [Specifying Column Names](#) on page 46
- [Data Types](#) on page 48
- [Casting Binary Data](#) on page 50
- [Security and Authentication](#) on page 51

Catalog and Schema Support

The Simba Drill ODBC Driver supports both catalogs and schemas to make it easy for the driver to work with various ODBC applications. Since Drill only organizes tables into schemas/databases, the driver provides a synthetic catalog named DRILL under which all of the schemas/databases are organized. The driver also maps the ODBC schema to the Drill catalog.

Schema Types and File Formats

The Simba Drill ODBC Driver supports the following schema types:

- HBase
- Distributed File System (DFS), supporting the following file formats:
 - Parquet
 - JavaScript Object Notation (JSON)
 - Comma-Separated Values (CSV)
 - Tab-Separated Values (TSV)
- Hive

Specifying Column Names

Use of the SQL asterisk (*) selector to return all columns in a table is limited. Modify the queries generated by Drill Explorer to specify columns to return based on the following syntax guidelines, depending on the schema type or file format.

HBase

The Simba Drill ODBC Driver presents an HBase column-family as an ODBC column. An HBase column is derived using the lookup scheme *ColumnFamily*['*Column*'] where *ColumnFamily* is an HBase column-family and *Column* is an HBase column contained in *ColumnFamily*. You must specify individual HBase columns, not column-families, so you need to have some familiarity with the data. For example, the following query may return incorrect results:

```
SELECT * FROM `hbase`.`students`
```

When the query is modified to specify individual columns, then the query results are correct:

```
SELECT CAST(account['name'] AS varchar(20)) FROM  
`hbase`.`students`
```

Parquet

When a query is used as a subquery, such as when you use a SQL view that Drill Explorer creates to connect to a data source using Tableau, you must specify column names in the query. For example, Tableau may not connect to your Drill data source successfully using the SQL view that Drill Explorer creates using the following query:

```
SELECT * FROM `dfs`.`default`.`./opt/drill/test.parquet`
```

When the query is modified to specify column names, Tableau connects to your Drill data source successfully using the SQL view:

```
SELECT CAST(column1 AS varchar(20)) FROM  
`dfs`.`default`.`./opt/drill/test.parquet`
```

JSON

When a query is used as a subquery, such as when you use a SQL view that Drill Explorer creates to connect to a data source using Tableau, you must specify column names in the query. For example, Tableau may not connect to your Drill data source successfully using the SQL view that Drill Explorer creates using the following query:

```
SELECT * FROM `dfs`.`default`.`./opt/drill/interval.json`
```

When the query is modified to specify column names, Tableau connects to your Drill data source successfully using the SQL view:

```
SELECT column1 FROM
`dfs`.`default`../opt/drill/interval.json`
```

To query nested elements, use the following syntax, where menu is a child of column1:

```
SELECT column1['menu'] FROM
`dfs`.`default`../opt/drill/interval.json`
```

You can query elements that are multiple levels deep. Continuing the example, if menuitem is a child of menu, then use the following syntax:

```
SELECT column1['menu']['menuitem'] FROM
`dfs`.`default`../opt/drill/interval.json`
```

CSV or TSV

Due to a limitation of Drill, you must specify columns using indices, requiring some familiarity with the data. For example, the following query may return incorrect results:

```
SELECT * FROM `dfs`.`default`../root/Test/Table.csv`
```

When the query is modified to specify columns using indices, the query results are correct:

```
SELECT columns[0], columns[2] FROM
`dfs`.`default`../root/Test/Table.csv`
```

Hive

When using a Hive schema, you can use queries that Drill Explorer generates without making any modifications.

Data Types

The Simba Drill ODBC Driver supports many common data formats, converting between Drill data types and SQL data types.

The table below lists the supported data type mappings.

Drill Type	SQL Type
ANY	SQL_VARCHAR

Drill Type	SQL Type
ARRAY	SQL_VARCHAR
BIGINT	SQL_BIGINT
BIT	SQL_BOOLEAN
CHARACTER	SQL_CHAR
CHARACTER VARYING	SQL_VARCHAR
DATE	SQL_DATE
DATETIME	SQL_DATE
FLOAT	SQL_REAL
DECIMAL9	SQL_DECIMAL
DECIMAL18	SQL_DECIMAL
DECIMAL28DENSE	SQL_DECIMAL
DECIMAL38DENSE	SQL_DECIMAL
DECIMAL38SPARSE	SQL_DECIMAL
FIXEDBINARY	SQL_BINARY
FIXEDCHAR	SQL_CHAR
FIXED16CHAR	SQL_WCHAR
FLOAT4	SQL_FLOAT
FLOAT8	SQL_DOUBLE
INT	SQL_INT
MAP	SQL_VARCHAR

Drill Type	SQL Type
MONEY	SQL_DECIMAL
NULL	SQL_VARCHAR
SMALLINT	SQL_SMALLINT
TIMESTAMP	SQL_TIMESTAMP
TIME WITH TIME ZONE	SQL_TYPE_TIMESTAMP
TIMESTAMP WITH TIME	SQL_TYPE_TIMESTAMP
TINYINT	SQL_TINYINT
UNIT1	SQL_TINYINT
UINT2	SQL_SMALLINT
UINT4	SQL_INTEGER
UINT8	SQL_BIGINT
UNION	SQL_VARCHAR
VAR16CHAR	SQL_WVARCHAR
VARCHAR	SQL_VARCHAR
VARDECIMAL	SQL_DECIMAL

Casting Binary Data

Given that Drill can work with self-describing data such as HBase and file systems without central metadata definitions, there are scenarios where the file formats do not have defined data types for the data. HBase, for example, always treats data as binary. Drill provides auxiliary functions to cast (or interpret) the data as certain data types. The following are a few examples of how to cast data as a specific data type. For more information about the SQL queries in Drill, refer to the Apache Drill documentation located at <http://drill.apache.org/docs/>.

HBase and Parquet store data in binary format. In SQL statements, you need to cast binary data to another format explicitly to view the data. For example, the following query displays results from an HBase database in binary format:

```
SELECT account['name'] FROM `hbase`.`students`
```

The following query displays the same results in string format:

```
SELECT CAST(account['name'] AS varchar(20)) FROM
`hbase`.`students`
```

The following query displays results from a Parquet file in binary format:

```
SELECT column1 FROM
`dfs`.`default`.`./opt/drill/test.parquet`
```

The following query displays the same results in string format:

```
SELECT CAST(column1 AS varchar(20)) FROM
`dfs`.`default`.`./opt/drill/test.parquet`
```

You can also cast the data as other data types, such as integer or date formats, as needed.

Security and Authentication

To protect data from unauthorized access, some Drill data stores require connections to be authenticated with user credentials. The Simba Drill ODBC Driver provides a mechanism that enables you to authenticate your connection using your Drill user name and password, the Kerberos protocol, or the MapR-SASL protocol. For detailed configuration instructions, see [Configuring Authentication on Windows](#) on page 12 or [Configuring Authentication on a Non-Windows Machines](#) on page 37.

Additionally, the driver supports the following types of SSL connections:

- No identity verification
- One-way authentication

Additionally, the driver supports SSL connections with one-way authentication. If the server has an SSL-enabled socket, then you can configure the driver to connect to it.

It is recommended that you enable SSL whenever you connect to a server that is configured to support it. SSL encryption protects data and credentials when they are transferred over the network, and provides stronger security than authentication alone. For detailed configuration instructions, see [Configuring TLS \(SSL\) on a Windows](#)

[Machine](#) on page 14 or [Configuring TLS \(SSL\) on a non-Windows Machine](#) on page 39.

Driver Configuration Options

Driver Configuration Options lists the configuration options available in the Simba Drill ODBC Driver alphabetically by field or button label. Options having only key names, that is, 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 Simba Drill ODBC Driver DSN Setup dialog box. When using a connection string or configuring a connection from a Linux or macOS computer, use the key names provided.

Configuration Options Appearing in the User Interface

The following configuration options are accessible via the Windows user interface for the Simba Drill ODBC Driver, or via the key name when using a connection string or configuring a connection from a Linux or macOS computer:

- [Advanced Properties](#) on page 54
- [Authentication Type](#) on page 54
- [Catalog](#) on page 55
- [Cluster ID](#) on page 55
- [Connection Type](#) on page 56
- [Default Schema](#) on page 56
- [Delegation UID](#) on page 56
- [Disable Async](#) on page 57
- [Disable Certificate Verification](#) on page 57
- [Disable Hostname Verification](#) on page 58
- [Enable SSL](#) on page 58
- [Host](#) on page 58
- [Host FQDN](#) on page 59
- [Kerberos SPN Configuration Required](#) on page 59
- [Log Level](#) on page 59
- [Log Path](#) on page 60
- [Minimum TLS](#) on page 61
- [Password](#) on page 61
- [Port](#) on page 62
- [Quorum](#) on page 62
- [Service Name](#) on page 62
- [TLS Protocol](#) on page 62
- [TrustedCertificates](#) on page 63
- [Use Only SSPI](#) on page 63
- [User](#) on page 64
- [Use Exact TLS Version](#) on page 64
- [Use System Trust Store](#) on page 64

Advanced Properties

Key Name	Default Value	Required
Advanced Properties	<pre>CastAnyToVarchar=true; HandshakeTimeout=5; QueryTimeout=180; TimestampTZDisplayTimezone=local; ExcludedSchemas=sys, INFORMATION_ SCHEMA; NumberOfPrefetchBuffers=5</pre>	No

Description

Advanced properties for configuring the driver. Separate advanced properties by using a semicolon (;), and then surround all the advanced properties in a connection string by using braces ({ }).

For detailed information about supported advanced properties, see [Advanced Properties](#) on page 66.

Note:

You can set custom connection properties by specifying them as advanced properties. If you specify a property that the driver does not explicitly support, the driver still accepts the property, and passes it to the server for processing.

Authentication Type

Key Name	Default Value	Required
AuthenticationType	None	Yes

Description

This option specifies how the driver authenticates the connection to Drill.

- **No Authentication** (`No Authentication`): The driver does not authenticate the connection to Drill.
- **Plain** (`Plain or Basic Authentication`): The driver authenticates the connection using a user name and a password.

- **Kerberos** (`Kerberos`): The driver authenticates the connection using the Kerberos protocol.
- **MapR-SASL** (`MapRSASL`): The driver authenticates the connection using the MapR-SASL protocol.
- **Username only** (`Username`): The driver does not authenticate the connection to Drill, and the server has impersonation enabled or disabled.

Catalog

Key Name	Default Value	Required
Catalog	(See description below.)	No

Description

The name of the synthetic catalog under which all of the schemas/databases are organized. This catalog name is used as the value for `SQL_DATABASE_NAME` or `CURRENT CATALOG`.

If a value is not specified, the default is determined by one of the following behaviors:

- If the server supports metadata or `GetMetadataWithQueries` is set to `false`, the driver defaults to querying the server for the catalog name.
- If the server does not support metadata or `GetMetadataWithQueries` is set to `true`, the driver defaults to the default catalog name specified in the driver's `.did` file (typically, `DRILL`).

Cluster ID

Key Name	Default Value	Required
<code>ZKClusterID</code>	<code>drillbits1</code>	Yes, if the connection type is ZooKeeper Quorum.

Description

The name of the ZooKeeper cluster that the driver connects to.

Connection Type

Key Name	Default Value	Required
ConnectionType	Direct to Drillbit (Direct)	Yes

Description

This option specifies whether the driver connects to a single server or a ZooKeeper cluster.

- Direct to Drillbit (Direct): The driver connects to a single Drill server.
- ZooKeeper Quorum (ZooKeeper): The driver connects to a ZooKeeper cluster.

Default Schema

Key Name	Default Value	Required
Schema	None	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.

Delegation UID

Key Name	Default Value	Required
DelegationUID	None	No

Description

If a value is specified for this setting, the driver delegates all operations against Drill to the specified user, rather than to the authenticated user for the connection.

Note:

This option is applicable only when Plain authentication is enabled. For more information, see [Configuring Authentication on Windows](#) on page 12 or [Configuring Authentication on a Non-Windows Machines](#) on page 37.

Disable Async

Key Name	Default Value	Required
DisableAsync	Clear (0)	No

Description

This option specifies whether the driver supports asynchronous queries.

- Enabled (1): The driver does not support asynchronous queries.
- Disabled (0): The driver supports asynchronous queries.

! Important:

This option is not supported in connection strings or DSNs. Instead, it must be set as a driver-wide property in the `simba.drillodbc.ini` file. Settings in that file apply to all connections that use the driver.

Disable Certificate Verification

Key Name	Default Value	Required
DisableCertificateVerification	0	No

Description

This property specifies whether the driver verifies the host certificate against the trust store. Accepted values are:

- Disabled (0): The driver verifies the certificate against the trust store.
- Enabled (1): The driver does not verify the certificate against the trust store.

Disable Hostname Verification

Key Name	Default Value	Required
DisableHostVerification	0	No

Description

This property specifies whether the driver verifies that the host in the certificate is the host being connected to. Accepted values are

- Disabled (0): The driver verifies the certificate against the host being connected to.
- Enabled (1): The driver does not verify the certificate against the host.

Enable SSL

Key Name	Default Value	Required
SSL	Clear (0)	No

Description

This option specifies whether the client uses an SSL encrypted connection to communicate with the Drill server.

- Enabled (1): The client communicates with the Drill server using SSL.
- Disabled (0): SSL is disabled.

SSL is configured independently of authentication. When authentication and SSL are both enabled, the driver performs the specified authentication method over an SSL connection.

Host

Key Name	Default Value	Required
Host	localhost	Yes, if the connection type is Direct to Drillbit.

Description

The IP address or host name of the Drill server.

Host FQDN

Key Name	Default Value	Required
KrbServiceHost	The host name of the server that you are connecting to.	No

Description

The fully qualified domain name of the Drill server host.

Kerberos SPN Configuration Required

Key Name	Default Value	Required
krbSpnConfigurationsRequired	0	No

Description

This property specifies if the driver requires Kerberos Service Host and Kerberos Service Name to be specified, and not accept empty values for these options. Accepted values are

- 0: The driver treats Kerberos Service Host and Service Name as optional.
- 1: The driver treats Kerberos Service Host and Service Name as required.

Log Level

Key Name	Default Value	Required
LogLevel	OFF (0)	No

Description

Use this property to enable or disable logging in the driver and to specify the amount of detail included in log files.

! Important:

- Only enable logging long enough to capture an issue. Logging decreases performance and can consume a large quantity of disk space.
- The settings for logging apply to every connection that uses the Simba Drill ODBC Driver, so make sure to disable the feature after you are done using it.
- This option is not supported in connection strings. To configure logging for the Windows driver, you must use the Logging Options dialog box. To configure logging for a non-Windows driver, you must use the `simba.drillodbc.ini` file.

Set the property to one of the following values:

- OFF (0): Disable all logging.
- FATAL (1): Logs severe error events that lead the driver to abort.
- ERROR (2): Logs error events that might allow the driver to continue running.
- WARNING (3): Logs events that might result in an error if action is not taken.
- INFO (4): Logs general information that describes the progress of the driver.
- DEBUG (5): Logs detailed information that is useful for debugging the driver.
- TRACE (6): Logs all driver activity.

When logging is enabled, the driver produces two log files at the location you specify in the Log Path (`LogPath`) property, where `[DriverName]` is the name of the driver:

- A `[DriverName].log` file that logs driver activity that is not specific to a connection.
- A `[DriverName]_connection_[Number].log` for each connection made to the database, where `[Number]` is a number that identifies each log file. This file logs driver activity that is specific to the connection.

If you enable the `UseLogPrefix` connection property, the driver prefixes the log file name with the user name associated with the connection and the process ID of the application through which the connection is made. For more information, see [UseLogPrefix](#) on page 66.

Log Path

Key Name	Default Value	Required
<code>LogPath</code>	None	Yes, if logging is enabled.

Description

The full path to the folder where the driver saves log files when logging is enabled.

! Important:

This option is not supported in connection strings. To configure logging for the Windows driver, you must use the Logging Options dialog box. To configure logging for a non-Windows driver, you must use the `simba.drillodbc.ini` file.

Minimum TLS

Key Name	Default Value	Required
Min_TLS	2	No

Description

The minimum version of TLS/SSL that the driver allows the data store to use for encrypting connections. For example, if TLS 1.1 is specified, TLS 1.0 cannot be used to encrypt connections.

- TLS 1.0 (0): The connection must use at least TLS 1.0.
- TLS 1.1 (1): The connection must use at least TLS 1.1.
- TLS 1.2 (2): The connection must use at least TLS 1.2.

Password

Key Name	Default Value	Required
PWD	None	Yes, if Authentication Type is set to Plain.

Description

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

Port

Key Name	Default Value	Required
Port	31010	Yes, if the connection type is Direct to Drillbit.

Description

The TCP port that the Drill server uses to listen for client connections.

Quorum

Key Name	Default Value	Required
ZKQuorum	None	Yes, if the connection type is ZooKeeper Quorum.

Description

The server or servers in your ZooKeeper cluster. Separate multiple servers using a comma (,).

Service Name

Key Name	Default Value	Required
KrbServiceName	drill	No

Description

The Kerberos service principal name of the Drill server.

TLS Protocol

Key Name	Default Value	Required
TLSProtocol	Empty, which defaults to tlsv12.	No

Description

This property specifies the TLS protocol version used. Accepted values are:

- `tlsv1`
- `tlsv11`
- `tlsv12`

TrustedCertificates

Key Name	Default Value	Required
<code>TrustedCerts</code>	The <code>cacerts.pem</code> file in the <code>\lib</code> subfolder within the driver's installation directory. 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 macOS driver.	No

Description

The full path of the `.pem` file containing trusted CA certificates, for verifying the server.

If this option is not set, then the driver defaults to using the trusted CA certificates `.pem` file installed by the driver.

Use Only SSPI

Key Name	Default Value	Required
<code>UseOnlySSPI</code>	Clear (0)	No

Description

This option specifies how the driver handles Kerberos authentication: either with the SSPI plugin or with MIT Kerberos.

- Enabled (1): The driver handles Kerberos authentication by using the SSPI plugin instead of MIT Kerberos by default.
- Disabled (0): The driver uses MIT Kerberos to handle Kerberos authentication, and only uses the SSPI plugin if the GSSAPI library is not available.

! Important:

This option is available only in the Windows driver.

User

Key Name	Default Value	Required
UID	None	Yes, if Authentication Type is set to Plain.

Description

The user name that you use to access the Drill server.

Use Exact TLS Version

Key Name	Default Value	Required
UseExactTLSProtocolVersion	Clear (0)	No

Description

This option specifies whether to use the TLS Protocol or Minimum TLS configuration options to determine the accepted level of TLS when connecting to the data store. See [Minimum TLS](#) on page 61 and [TLS Protocol](#) on page 62 for details.

- Disabled (0): The driver uses the Minimum TLS configuration option.
- Enabled (1): The driver uses the TLS Protocol configuration option.

Use System Trust Store

Key Name	Default Value	Required
UseSystemTrustStore	Clear (0)	No

Description

This option specifies whether to use a CA certificate from the system trust store, or from a specified PEM file.

- Enabled (1): The driver verifies the connection using a certificate in the system trust store.
- Disabled (0): The driver verifies the connection using a specified `.pem` file. For information about specifying a `.pem` file, see [.pem](#).



Note:

This option is only available on Windows.

Configuration Options Having Only Key Names

The `Driver` configuration option does not appear in the Windows user interface for the Simba Drill ODBC Driver. It is accessible only when you use a connection string or configure a connection on macOS or Linux.

The `UseLogPrefix` property must be configured as a Windows Registry key value, or as a driver-wide property in the `simba.drillodbc.ini` file for macOS or Linux.

- [UseLogPrefix](#) on page 66

Driver

Key Name	Default Value	Required
<code>Driver</code>	Simba Drill ODBC Driver when installed on Windows, or the absolute path of the driver shared object file when installed on a non-Windows machine.	Yes

Description

On Windows, the name of the installed driver (`Simba Drill ODBC Driver`).

On other platforms, the name of the installed driver as specified in `odbcinst.ini`, or the absolute path of the driver shared object file.

UseLogPrefix

Key Name	Default Value	Required
UseLogPrefix	0	No

Description

This option specifies whether the driver includes a prefix in the names of log files so that the files can be distinguished by user and application.

! Important:

To configure this option for the Windows driver, you create a value for it in one of the following registry keys:

- For a 32-bit driver installed on a 64-bit machine: **HKEY_LOCAL_MACHINE\SOFTWAREWow6432Node\Simba\Simba Drill ODBC Driver\Driver**
- Otherwise: **HKEY_LOCAL_MACHINE\SOFTWARE\Simba\Simba Drill ODBC Driver\Driver**

Use `UseLogPrefix` as the value name, and either 0 or 1 as the value data.

To configure this option for a non-Windows driver, you must use the `simba.drillodbc.ini` file.

Set the property to one of the following values:

- 1: The driver prefixes log file names with the user name and process ID associated with the connection that is being logged.

For example, if you are connecting as a user named "jdoe" and using the driver in an application with process ID 7836, the generated log files would be named `jdoe_7836_Simba[DriverName]_driver.log` and `jdoe_7836_Simba[DriverName]_connection_[Number].log`, where *[Number]* is a number that identifies each connection-specific log file.

- 0: The driver does not include the prefix in log file names.

Advanced Properties

Advanced Properties lists in alphabetical order and describes advanced properties you can set when using the Simba Drill ODBC Driver.

Separate advanced properties using a semicolon (;).

For example, the following Advanced Properties string excludes the schemas named `test` and `abc`, sets the timeout to 30 seconds, and sets the time zone to Coordinated Universal Time:

```
HandshakeTimeout=30;QueryTimeout=30;TimestampTZDisplayTimezone=utc;ExcludedSchemas=test,abc
```


Note:

You can set custom connection properties by specifying them as advanced properties. If you specify a property that the driver does not explicitly support, the driver still accepts the property, and passes it to the server for processing.

CastAnyToVarchar

Property Name	Default Value	Required
CastAnyToVarchar	true	No

Description

When this property is set to true, if `SQLColumns` returns columns of type ANY, then the driver casts the columns to VARCHAR.

When this property is set to false, the driver does not change the returned columns.

ConvertToCast

Property Name	Default Value	Required
ConvertToCast	false	No

Description

When activated this enables the Cast Query Translation function. This function helps optimize your queries for Power BI

DefaultDecimalScale

Property Name	Default Value	Required
DefaultDecimalScale	10	No

Description

The default scale for DECIMAL columns that are returned through SQLColumns. The driver uses this value only if SQLColumns does not return a scale value or a numeric precision value that the driver can use to determine the scale.

EnforceSaslEncrypt

Key Name	Default Value	Required
EnforceSaslEncrypt	0	No

Description

This property specifies whether the driver enforces encryption when connecting to Drill using the Simple Authentication and Security Layer (SASL) protocol. To use this option the Authentication Type must be set to `Kerberos`.

- 1: The driver only connects to Drillbit with SASL encryption support. Recommended for applications that require strong security.
- 0: The driver negotiates SASL encryption based only on the server-side response and will not enforce its own check. This means the connection may not be encrypted depending on the data source server settings.

ExceptionBufferSize

Key Name	Default Value	Required
ExceptionBufferSize	384	No

Description

If error truncation is enabled, this property specifies the length, in characters, of the displayed error messages. The full error messages are still recorded in the driver log. For more information about error truncation, see [ForceErrorTruncation](#) on page 69

ExcludedSchemas

Property Name	Default Value	Required
ExcludedSchemas	sys, INFORMATION_ SCHEMA	No

Description

A list of schemas that do not appear in client applications such as Drill Explorer, Tableau, and Excel. Separate schemas in the list using a comma (,).

This property should not be used at the same time as `IncludedSchemas`. If both `IncludedSchemas` and `ExcludedSchemas` are specified, `IncludedSchemas` takes precedence and `ExcludedSchemas` is ignored.

ForceErrorTruncation

Key Name	Default Value	Required
ForceErrorTruncation	false when the driver is installed on Windows, or true when the driver is installed on a non-Windows machine.	No

Description

This property specifies whether the driver truncates displayed error messages. The full error messages are still recorded in the driver log.

Note:

If you are using the unixODBC driver manager, this property should be set to `true`.

- `false`: The driver displays the full error message.
- `true`: The driver truncates the displayed error messages to the length specified in the `ExceptionBufferSize` property. For more information, see [ExceptionBufferSize](#) on page 68.

GetMetadataWithQueries

Key Name	Default Value	Required
GetMetadataWithQueries	None	No

Description

This property specifies whether the driver uses queries or native API calls when retrieving metadata, preparing statements, or executing statements. The driver can only use native API calls if it detects that the server is running a version of Drill that supports those calls.

- If this property is set to `True`, the driver uses queries to retrieve metadata, prepare statements, and execute statements.
- If this property is set to `False`, the driver instead uses native API calls when it is connected to a server that supports native API calls.

HandshakeTimeout

Property Name	Default Value	Required
HandshakeTimeout	5	No

Description

An integer value representing the number of seconds that the driver waits before stopping an attempt to connect to a data source.

If this property is set to `0`, the driver does not stop connection attempts.

HeartBeatFreqSec

Property Name	Default Value	Required
HeartBeatFreqSec	15	No

Description

The number of seconds of inactivity before the Drill client sends a heartbeat to the server in order to check the status of the connection.

If this property is set to 0, the Drill client does not send any heartbeats.

IncludedSchemas

Property Name	Default Value	Required
IncludedSchemas	None	No

Description

A list of schemas that appear in client applications such as Drill Explorer, Tableau, and Excel. Separate schemas in the list using a comma (,). If this option is set, then schemas that are not in this list are not queried by the driver.

This property should not be used at the same time as `ExcludedSchemas`. If both `IncludedSchemas` and `ExcludedSchemas` are specified, `IncludedSchemas` takes precedence and `ExcludedSchemas` is ignored.

NonZeroNameMaxLen

Property Name	Default Value	Required
NonZeroNameMaxLen	false	No

Description

The Simba Drill ODBC Driver does not enforce a maximum length for schema, catalog, table, and column names. By default, the driver reports these lengths as 0, which typically indicates that there is no maximum length. However, some BI tools interpret a length of 0 literally. To make sure that the names are displayed correctly in your BI tool, set the `NonZeroNameMaxLen` property.

When this property is set to `true`, the driver reports an appropriate non-zero value for the maximum length of schema, catalog, table, and column names.

When this property is set to `false`, the driver reports 0 as the maximum length of schema, catalog, table, and column names.

NumberOfPrefetchBuffers

Property Name	Default Value	Required
NumberOfPrefetchBuffers	5	No

Description

The size of the record batch queue in the driver.

Note:

When set to a value below 1, the value defaults to 1.

QueryTimeout

Property Name	Default Value	Required
QueryTimeout	180	No

Description

The number of seconds that the driver waits before stopping a query.

If this property is set to 0, the driver does not stop queries.

RemoveQryCatalog

Property Name	Default Value	Required
RemoveQryCatalog	true	No

Description

When this property is set to `true`, the driver removes catalog names from queries. Enable this option if you are working with a server that does not accept queries containing catalog names.

When this property is set to `false`, the driver does not remove catalog names before executing queries.

SaslPluginDir

Property Name	Default Value	Required
SaslPluginDir	None	No

Description

Allows you to override the default SASL plugin directory with a custom directory path.

StringColumnLength

Key Name	Default Value	Required
StringColumnLength	65535	No

Description

The maximum column length that the driver reports for string columns.

TimestampTZDisplayTimezone

Property Name	Default Value	Required
TimestampTZDisplayTimezone	local	No

Description

When set to `local`, timestamps use the time zone of the user.

When set to `utc`, timestamps appear in Coordinated Universal Time (UTC).

Third-Party Trademarks

Linux is the registered trademark of Linus Torvalds in Canada, United States and/or other countries.

Mac, macOS, Mac OS, and OS X are trademarks or registered trademarks of Apple, Inc. or its subsidiaries in Canada, United States and/or other countries.

Microsoft, MSDN, Windows, Windows Server, Windows Vista, and the Windows start button are trademarks or registered trademarks of Microsoft Corporation or its subsidiaries in Canada, United States and/or other countries.

Red Hat, Red Hat Enterprise Linux, and CentOS are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in Canada, United States and/or other countries.

SUSE is a trademark or registered trademark of SUSE LLC or its subsidiaries in Canada, United States and/or other countries.

Apache Drill, Apache, and Drill are trademarks or registered trademarks of The Apache Software Foundation or its subsidiaries in Canada, United States and/or other countries.

All other trademarks are trademarks of their respective owners.

Third-Party Licenses

The licenses for the third-party libraries that are included in this product are listed below.

Boost Software License - Version 1.0 - August 17th, 2003

Permission is hereby granted, free of charge, to any person or organization obtaining a copy of the software and accompanying documentation covered by this license (the "Software") to use, reproduce, display, distribute, execute, and transmit the Software, and to prepare derivative works of the Software, and to permit third-parties to whom the Software is furnished to do so, all subject to the following:

The copyright notices in the Software and this entire statement, including the above license grant, this restriction and the following disclaimer, must be included in all copies of the Software, in whole or in part, and all derivative works of the Software, unless such copies or derivative works are solely in the form of machine-executable object code generated by a source language processor.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR ANYONE DISTRIBUTING THE SOFTWARE BE LIABLE FOR ANY DAMAGES OR OTHER LIABILITY, WHETHER IN CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

cURL License

COPYRIGHT AND PERMISSION NOTICE

Copyright (c) 1996 - 2015, Daniel Stenberg, daniel@haxx.se.

All rights reserved.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR

OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

Cyrus SASL License

Copyright (c) 1994-2012 Carnegie Mellon University. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name "Carnegie Mellon University" must not be used to endorse or promote products derived from this software without prior written permission. For permission or any other legal details, please contact

Office of Technology Transfer
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213-3890
(412) 268-4387, fax: (412) 268-7395
tech-transfer@andrew.cmu.edu

4. Redistributions of any form whatsoever must retain the following acknowledgment:

"This product includes software developed by Computing Services at Carnegie Mellon University (<http://www.cmu.edu/computing/>)."

CARNEGIE MELLON UNIVERSITY DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL CARNEGIE MELLON UNIVERSITY BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

dtoa License

The author of this software is David M. Gay.

Copyright (c) 1991, 2000, 2001 by Lucent Technologies.

Permission to use, copy, modify, and distribute this software for any purpose without fee is hereby granted, provided that this entire notice is included in all copies of any software which is or includes a copy or modification of this software and in all copies of the supporting documentation for such software.

THIS SOFTWARE IS BEING PROVIDED "AS IS", WITHOUT ANY EXPRESS OR IMPLIED WARRANTY. IN PARTICULAR, NEITHER THE AUTHOR NOR LUCENT MAKES ANY REPRESENTATION OR WARRANTY OF ANY KIND CONCERNING THE MERCHANTABILITY OF THIS SOFTWARE OR ITS FITNESS FOR ANY PARTICULAR PURPOSE.

Expat License

Copyright (c) 1998, 1999, 2000 Thai Open Source Software Center Ltd

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NOINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Google Protocol Buffers License

This license applies to all parts of Protocol Buffers except the following:

- Atomicops support for generic gcc, located in `src/google/protobuf/stubs/atomicops_internals_generic_gcc.h`. This file is copyrighted by Red Hat Inc.
- Atomicops support for AIX/POWER, located in `src/google/protobuf/stubs/atomicops_internals_aix.h`. This file is copyrighted by Bloomberg Finance LP.

Copyright 2014, Google Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Code generated by the Protocol Buffer compiler is owned by the owner of the input file used when generating it. This code is not standalone and requires a support library to be linked with it. This support library is itself covered by the above license.

ICU License - ICU 1.8.1 and later

COPYRIGHT AND PERMISSION NOTICE

Copyright (c) 1995-2014 International Business Machines Corporation and others

All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

All trademarks and registered trademarks mentioned herein are the property of their respective owners.

JSON in Java License

Copyright © 2015 Douglas Crockford

The license includes this restriction: "The software shall be used for good, not evil." If your conscience cannot live with that, then choose a different package.

OpenSSL License

Copyright (c) 1998-2016 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"

4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.

5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License

Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)

All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

"This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)"

The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related :-).

4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement:

"This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The licence and distribution terms for any publically available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution licence [including the GNU Public Licence.]

RapidJSON License

Tencent is pleased to support the open source community by making RapidJSON available.

Copyright (C) 2015 THL A29 Limited, a Tencent company, and Milo Yip. All rights reserved.

If you have downloaded a copy of the RapidJSON binary from Tencent, please note that the RapidJSON binary is licensed under the MIT License.

If you have downloaded a copy of the RapidJSON source code from Tencent, please note that RapidJSON source code is licensed under the MIT License, except for the third-party components listed below which are subject to different license terms. Your integration of RapidJSON into your own projects may require compliance with the MIT License, as well as the other licenses applicable to the third-party components included within RapidJSON.

A copy of the MIT License is included in this file.

Other dependencies and licenses:

Open Source Software Licensed Under the BSD License:

The msinttypes r29
Copyright (c) 2006-2013 Alexander Chemeris
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS AND CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Open Source Software Licensed Under the JSON License:

json.org
Copyright (c) 2002 JSON.org
All Rights Reserved.

JSON_checker
Copyright (c) 2002 JSON.org
All Rights Reserved.

Terms of the JSON License:

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

The Software shall be used for Good, not Evil.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Terms of the MIT License:

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Stringencoders License

Copyright 2005, 2006, 2007

Nick Galbreath -- nickg [at] modp [dot] com

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of the modp.com nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This is the standard "new" BSD license:

<http://www.opensource.org/licenses/bsd-license.php>

y2038 License

Copyright (c) 2007-2010 Michael G Schwern

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

zlib License

Copyright (C) 1995-2013 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly

Mark Adler

jloup@gzip.orgmadler@alumni.caltech.edu

Apache License, Version 2.0

The following notice is included in compliance with the Apache License, Version 2.0 and is applicable to all software licensed under the Apache License, Version 2.0.

Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. **Grant of Copyright License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. **Grant of Patent License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
4. **Redistribution.** You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
 - (a) You must give any other recipients of the Work or Derivative Works a

copy of this License; and

- (b) You must cause any modified files to carry prominent notices stating that You changed the files; and
- (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
- (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. **Submission of Contributions.** Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. **Trademarks.** This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. **Disclaimer of Warranty.** Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND,

either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. **Accepting Warranty or Additional Liability.** While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

This product includes software that is licensed under the Apache License, Version 2.0 (listed below):

Apache Drill

Copyright © 2012-2014 The Apache Software Foundation

Apache Hadoop Common

Copyright © 2014 The Apache Software Foundation

Apache ZooKeeper

Copyright © 2010 The Apache Software Foundation

Java Native Access

Copyright © 2007-2014 Todd Fast, Timothy Wall, Liang Chen

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.