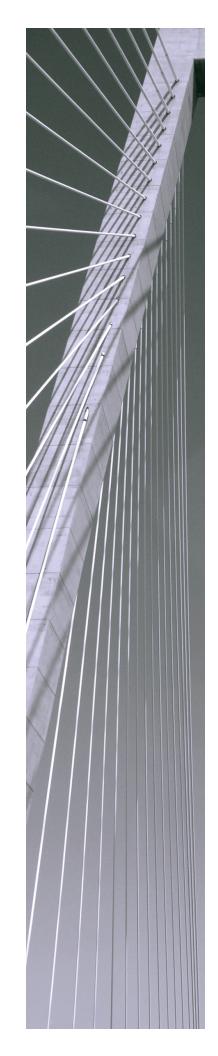
# MAGNITUDE

# Simba Spark ODBC Driver with SQL Connector

# Installation and Configuration Guide

Simba Technologies Inc.

Version 2.6.5 December 21, 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 Spark ODBC Driver with SQL Connector Installation and Configuration Guide explains how to install and configure the Simba Spark 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 Spark ODBC Driver, as well as administrators and developers integrating the driver.

### **Knowledge Prerequisites**

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

- Familiarity with the platform on which you are using the Simba Spark ODBC
   Driver
- Ability to use the data source to which the Simba Spark 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 Spark ODBC Driver	7
Windows Driver	
Windows System Requirements	8
Installing the Driver on Windows	8
Creating a Data Source Name on Windows	
Configuring a DSN-less Connection on Windows	12
Configuring Authentication on Windows	14
Configuring Advanced Options on Windows	
Configuring HTTP Options on Windows	20
Configuring SSL Verification on Windows	21
Configuring Server-Side Properties on Windows	
Configuring Logging Options on Windows	24
Configuring Kerberos Authentication for Windows	
Verifying the Driver Version Number on Windows	
macOS Driver	
macOS System Requirements	
Installing the Driver on macOS	
Verifying the Driver Version Number on macOS	33
Linux Driver	
Linux System Requirements	
Installing the Driver Using the RPM File	
Installing the Driver Using the Tarball Package	
Verifying the Driver Version Number on Linux	
	20
AIX Driver	
AIX System Requirements	
Installing the Driver on AIX	
Solaris Driver	
Solaris System Requirements	
Installing the Driver on Solaris	
Verifying the Driver Version Number on Solaris	41
Configuring the ODBC Driver Manager on Non-Windows Machines	42
Specifying ODBC Driver Managers on Non-Windows Machines	42

Specifying the Locations of the Driver Configuration Files	43
Configuring ODBC Connections on a Non-Windows Machine	45
Creating a Data Source Name on a Non-Windows Machine	45
Configuring a DSN-less Connection on a Non-Windows Machine	
Configuring Service Discovery Mode on a Non-Windows Machine	
Configuring Authentication on a Non-Windows Machine	51
Configuring SSL Verification on a Non-Windows Machine	54
Configuring Server-Side Properties on a Non-Windows Machine	55
Configuring Logging Options on a Non-Windows Machine	56
Setting Driver-Wide Configuration Options on a Non-Windows Machine	58
Testing the Connection on a Non-Windows Machine	59
Authentication Mechanisms	61
Shark Server	
Using a Connection String	
DSN Connection String Example	
DSN-less Connection String Examples	63
Features	67
SQL Connector for HiveQL	67
Data Types	68
Timestamp Function Support	69
Catalog and Schema Support	69
spark_system Table	69
Server-Side Properties	70
Get Tables With Query	70
Active Directory	70
Write-back	71
Dynamic Service Discovery using Datastax AOSS	71
Security and Authentication	71
Driver Configuration Options	73
Configuration Options Appearing in the User Interface	
Configuration Options Having Only Key Names	
Third-Party Trademarks	
Third-Party Licenses	110

### About the Simba Spark ODBC Driver

The Simba Spark ODBC Driver is used for direct SQL and HiveQL access to Apache Hadoop / Spark distributions, enabling Business Intelligence (BI), analytics, and reporting on Hadoop-based data. The driver efficiently transforms an application's SQL query into the equivalent form in HiveQL, which is a subset of SQL-92. If an application is Spark-aware, then the driver is configurable to pass the query through to the database for processing. The driver interrogates Spark to obtain schema information to present to a SQL-based application. Queries, including joins, are translated from SQL to HiveQL. For more information about the differences between HiveQL and SQL, see SQL Connector for HiveQL on page 67.

The Simba Spark ODBC Driver complies with the ODBC 3.80 data standard and adds important functionality such as Unicode and 32- and 64-bit support for high-performance computing environments.

ODBC is one of the most established and widely supported APIs for connecting to and working with databases. At the heart of the technology is the ODBC driver, which connects an application to the database. For more information about ODBC, see *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 Spark ODBC Driver is available for Microsoft® Windows®, Linux, Solaris, AIX, and macOS platforms.

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

#### 🖉 Note:

For basic configuration instructions that allow you to quickly set up the Windows driver so that you can evaluate and use it, see the *Simba ODBC Drivers Quick Start Guide for Windows*. The Quick Start Guide also explains how to use the driver in various applications.

### Windows Driver

### Windows System Requirements

The Simba Spark ODBC Driver supports Apache Spark versions 1.6 through 2.3.

Install the driver on client machines where the application is installed. Before installing the driver, make sure that you have the following:

- Administrator rights on your machine.
- A machine that meets the following system requirements:
  - One of the following operating systems:
    - Windows 10, 8.1, or 7 SP1
    - Windows Server 2016, 2012, or 2008 R2 SP1
  - 100 MB of available disk space

Before the driver can be used, the Visual C++ Redistributable for Visual Studio 2013 with the same bitness as the driver must also be installed. If you obtained the driver from the Simba website, then your installation of the driver automatically includes this dependency. Otherwise, you must install the redistributable manually. You can download the installation packages for the redistributable at https://www.microsoft.com/en-ca/download/details.aspx?id=40784.

### 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:

- SimbaSparkODBC32.msi for 32-bit applications
- SimbaSparkODBC64.msi for 64-bit applications

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

#### To install the Simba Spark ODBC Driver on Windows:

- 1. Depending on the bitness of your client application, double-click to run SimbaSparkODBC32.msi or SimbaSparkODBC64.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 Spark ODBC Driver, you need to create a Data Source Name (DSN). A DSN is a data structure that stores connection information so that it can be used by the driver to connect to Spark.

Alternatively, you can specify connection settings in a connection string or as driverwide settings. Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over driver-wide settings.

The following instructions describe how to create a DSN. For information about specifying settings in a connection string, see Using a Connection String on page 63. For information about driver-wide settings, see Configuring a DSN-less Connection on Windows on page 12.

#### 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 Spark.

- 2. In the ODBC Data Source Administrator, click the **Drivers** tab, and then scroll down as needed to confirm that the Simba Spark 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.

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 Spark ODBC Driver** and then click **Finish**. The Simba Spark 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. In the **Spark Server Type** list, select the appropriate server type for the version of Spark that you are running:
  - If you are running Shark 0.8.1 or earlier, then select **SharkServer**.
  - If you are running Shark 0.9, or Spark 1.1 or later, then select **SparkThriftServer**.
- 9. Choose one:
  - To connect to Spark without using the DataStax AOSS service, in the Service Discovery Mode drop-down list, select No Service Discovery.
  - Or, to discover Spark services via the DataStax AOSS service, in the Service Discovery Mode drop-down list, select AOSS.
- 10. Choose one:
  - If you selected **No Service Discovery** above, then in the **Host(s)** field, type the IP address or host name of the Spark server.
  - Or, if you selected AOSS above, then in the Host(s) field, type a commaseparated list of AOSS endpoints. Use the following format, where [AOSS\_ Endpoint] is the IP address or host name of the AOSS endpoint, and [AOSS\_Port] is the number of the TCP port that the AOSS endpoint uses to listen for client connections:

[AOSS\_Endpoint1]: [AOSS\_Port1], [AOSS\_Endpoint2]:
[AOSS\_Port2]

- 11. In the **Port** field, type the number of the TCP port that the Spark server uses to listen for client connections.
- 12. In the **Database** field, type 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. To inspect your databases and determine the appropriate schema to use, type the show databases command at the Spark command prompt.

13. In the Authentication area, configure authentication as needed. For more information, see Configuring Authentication on Windows on page 14.

#### Note:

Shark Server does not support authentication. Most default configurations of Spark Thrift Server require User Name authentication. To verify the authentication mechanism that you need to use for your connection, check the configuration of your Hadoop / Spark distribution. For more information, see Authentication Mechanisms on page 61.

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

#### Note:

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

15. In the **Thrift Transport** drop-down list, select the transport protocol to use in the Thrift layer.

#### Note:

For information about how to determine which Thrift transport protocols your Spark server supports, see Authentication Mechanisms on page 61.

- 16. If the Thrift Transport option is set to HTTP, then to configure HTTP options such as custom headers, click **HTTP Options**. For more information, see Configuring HTTP Options on Windows on page 20.
- 17. To configure client-server verification over SSL, click **SSL Options**. For more information, see Configuring SSL Verification on Windows on page 21.
- 18. To configure advanced driver options, click **Advanced Options**. For more information, see Configuring Advanced Options on Windows on page 18.
- 19. To configure server-side properties, click **Advanced Options** and then click **Server Side Properties**. For more information, see Configuring Server-Side Properties on Windows on page 23.

- 20. To configure logging behavior for the driver, click **Logging Options**. For more information, see Configuring Logging Options on Windows on page 24.
- 21. To test the connection, click **Test**. Review the results as needed, and then click **OK**.

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

- 22. To save your settings and close the Simba Spark ODBC Driver DSN Setup dialog box, click **OK**.
- 23. To close the ODBC Data Source Administrator, click **OK**.

### Configuring a DSN-less Connection on Windows

Some client applications provide support for connecting to a data source using a driver without a Data Source Name (DSN). To configure a DSN-less connection, you can use a connection string or the Simba Spark ODBC Driver Configuration tool that is installed with the Simba Spark ODBC Driver. Settings in a connection string apply only when you connect to Spark using that particular string, while settings in the driver configuration tool apply to every connection that uses the Simba Spark ODBC Driver.

The following section explains how to use the driver configuration tool. For information about using connection strings, see Using a Connection String on page 63.

#### Note:

Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over driver-wide settings.

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

- 1. Choose one:
  - If you are using Windows 7 or earlier, click Start <a>Simba Spark ODBC Driver 2.6 > Driver Configuration.</a>
  - Or, if you are using Windows 8 or later, click the arrow button at the bottom of the Start screen, and then click Simba Spark ODBC Driver 2.6 > Driver Configuration.

Make sure to select the Driver Configuration Tool that has the same bitness as the client application that you are using to connect to Spark.

2. If you are prompted for administrator permission to make modifications to the machine, click **OK**.

#### Note:

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

- 3. In the Spark Server Type list, select the appropriate server type for the version of Spark that you are running:
  - If you are running Shark 0.8.1 or earlier, then select **SharkServer**.
  - If you are running Shark 0.9, Spark 1.1 or later, then select **SparkThriftServer**.
- 4. Choose one:
  - To connect to Spark without using the DataStax AOSS service, in the Service Discovery Mode drop-down list, select No Service Discovery.
  - Or, to discover Spark services via the DataStax AOSS service, in the Service Discovery Mode drop-down list, select **AOSS**.
- 5. In the Authentication area, configure authentication as needed. For more information, see Configuring Authentication on Windows on page 14.

#### Note:

Shark Server does not support authentication. Most default configurations of Spark Thrift Server require User Name authentication. To verify the authentication mechanism that you need to use for your connection, check the configuration of your Hadoop / Spark distribution. For more information, see Authentication Mechanisms on page 61.

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

#### Note:

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

7. In the **Thrift Transport** drop-down list, select the transport protocol to use in the Thrift layer.

For information about how to determine which Thrift transport protocols your Spark server supports, see Authentication Mechanisms on page 61.

- 8. If the Thrift Transport option is set to HTTP, then to configure HTTP options such as custom headers, click **HTTP Options**. For more information, see Configuring HTTP Options on Windows on page 20.
- 9. To configure client-server verification over SSL, click **SSL Options**. For more information, see Configuring SSL Verification on Windows on page 21.
- 10. To configure advanced options, click **Advanced Options**. For more information, see Configuring Advanced Options on Windows on page 18.
- 11. To configure server-side properties, click **Advanced Options** and then click **Server Side Properties**. For more information, see Configuring Server-Side Properties on Windows on page 23.
- 12. To save your settings and close the Simba Spark ODBC Driver Configuration tool, click **OK**.

### **Configuring Authentication on Windows**

Some Spark Thrift Server instances are configured to require authentication for access. To connect to a Spark server, you must configure the Simba Spark ODBC Driver to use the authentication mechanism that matches the access requirements of the server and provides the necessary credentials.

For information about how to determine the type of authentication your Spark server requires, see Authentication Mechanisms on page 61.

You can specify authentication settings in a DSN, in a connection string, or as driverwide settings. Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over driver-wide settings.

### Using No Authentication

When connecting to a Spark server of type Shark Server, you must use No Authentication. When you use No Authentication, Binary is the only Thrift transport protocol that is supported.

#### To configure a connection without authentication:

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

- Or, to access authentication options for a DSN-less connection, open the Simba Spark ODBC Driver Configuration tool.
- 2. In the **Mechanism** drop-down list, select **No Authentication**.
- 3. If the Spark server is configured to use SSL, then click **SSL Options** to configure SSL for the connection. For more information, see Configuring SSL Verification on Windows on page 21.
- 4. To save your settings and close the dialog box, click **OK**.

#### Using Kerberos

If the Use Only SSPI advanced option is disabled, then Kerberos must be installed and configured before you can use this authentication mechanism. For information about configuring Kerberos on your machine, see Configuring Kerberos Authentication for Windows on page 26. For information about setting the Use Only SSPI advanced option, see Configuring Advanced Options on Windows on page 18.

#### Note:

This authentication mechanism is available only for Spark Thrift Server on non-HDInsight distributions.

#### To configure Kerberos authentication:

- 1. Choose one:
  - To access authentication options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, and then click **Configure**.
  - Or, to access authentication options for a DSN-less connection, open the Simba Spark ODBC Driver Configuration tool.
- 2. In the Mechanism drop-down list, select Kerberos.
- 3. Choose one:
  - To use the default realm defined in your Kerberos setup, leave the **Realm** field empty.
  - Or, if your Kerberos setup does not define a default realm or if the realm of your Spark Thrift Server host is not the default, then, in the **Realm** field, type the Kerberos realm of the Spark Thrift Server.
- 4. In the **Host FQDN** field, type the fully qualified domain name of the Spark Thrift Server host.

#### Note:

To use the Spark server host name as the fully qualified domain name for Kerberos authentication, in the **Host FQDN** field, type **\_HOST**.

- 5. In the **Service Name** field, type the service name of the Spark server.
- 6. To allow the driver to pass your credentials directly to the server for use in authentication, select **Delegate Kerberos Credentials**.
- 7. In the **Thrift Transport** drop-down list, select the transport protocol to use in the Thrift layer.

#### Important:

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

- 8. If the Spark server is configured to use SSL, then click **SSL Options** to configure SSL for the connection. For more information, see Configuring SSL Verification on Windows on page 21.
- 9. To save your settings and close the dialog box, click **OK**.

#### Using User Name

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

#### Note:

This authentication mechanism is available only for Spark Thrift Server on non-HDInsight distributions. Most default configurations of Spark Thrift Server require User Name authentication.

#### To configure User Name authentication:

- 1. Choose one:
  - To access authentication options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, and then click Configure.
  - Or, to access authentication options for a DSN-less connection, open the Simba Spark ODBC Driver Configuration tool.
- 2. In the Mechanism drop-down list, select User Name.
- 3. In the **User Name** field, type an appropriate user name for accessing the Spark server.
- 4. To save your settings and close the dialog box, click **OK**.

#### Using User Name And Password

This authentication mechanism requires a user name and a password.

#### To configure User Name And Password authentication:

- 1. Choose one:
  - To access authentication options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, and then click **Configure**.
  - Or, to access authentication options for a DSN-less connection, open the Simba Spark ODBC Driver Configuration tool.
- 2. In the Mechanism drop-down list, select User Name And Password.
- 3. In the **User Name** field, type an appropriate user name for accessing the Spark server.
- 4. In the **Password** field, type the password corresponding to the user name you typed above.
- 5. To save the password, select the **Save Password (Encrypted)** check box.

#### Important:

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

- 6. In the **Thrift Transport** drop-down list, select the transport protocol to use in the Thrift layer.
- If the Spark server is configured to use SSL, then click SSL Options to configure SSL for the connection. For more information, see Configuring SSL Verification on Windows on page 21.
- 8. To save your settings and close the dialog box, click **OK**.

#### Using Windows Azure HDInsight Emulator

This authentication mechanism is available only for Spark Thrift Server instances running on Windows Azure HDInsight Emulator.

## To configure a connection to a Spark server on Windows Azure HDInsight Emulator:

- 1. Choose one:
  - To access authentication options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, and then click **Configure**.
  - Or, to access authentication options for a DSN-less connection, open the Simba Spark ODBC Driver Configuration tool.
- 2. In the Mechanism drop-down list, select Windows Azure HDInsight Emulator.
- 3. In the **User Name** field, type an appropriate user name for accessing the Spark server.

- 4. In the **Password** field, type the password corresponding to the user name you specified above.
- 5. Click **HTTP Options**, and in the **HTTP Path** field, type the partial URL corresponding to the Spark server. Click **OK** to save your HTTP settings and close the dialog box.
- 6. To save your settings and close the dialog box, click **OK**.

#### Using Windows Azure HDInsight Service

This authentication mechanism is available only for Spark Thrift Server on HDInsight distributions.

# To configure a connection to a Spark server on Windows Azure HDInsight Service:

- 1. Choose one:
  - To access authentication options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, and then click **Configure**.
  - Or, to access authentication options for a DSN-less connection, open the Simba Spark ODBC Driver Configuration tool.
- 2. In the Mechanism drop-down list, select Windows Azure HDInsight Service.
- 3. In the **User Name** field, type an appropriate user name for accessing the Spark server.
- 4. In the **Password** field, type the password corresponding to the user name you typed abpve.
- 5. Click **HTTP Options**, and in the **HTTP Path** field, type the partial URL corresponding to the Spark server. Click **OK** to save your HTTP settings and close the dialog box.

#### Note:

If necessary, you can create custom HTTP headers. For more information, see Configuring HTTP Options on Windows on page 20.

- 6. Click **SSL Options** and configure SSL settings as needed. For more information, see Configuring SSL Verification on Windows on page 21.
- 7. Click **OK** to save your SSL configuration and close the dialog box, and then click **OK** to save your authentication settings and close the dialog box.

### **Configuring Advanced Options on Windows**

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

The following instructions describe how to configure advanced options in a DSN and in the driver configuration tool. You can specify the connection settings described below in a DSN, in a connection string, or as driver-wide settings. Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over driver-wide settings.

#### To configure advanced options on Windows:

- 1. Choose one:
  - To access advanced options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **Advanced Options**.
  - Or, to access advanced options for a DSN-less connection, open the Simba Spark ODBC Driver Configuration tool, and then click **Advanced Options**.
- 2. To disable the SQL Connector feature, select the Use Native Query check box.

#### Important:

- When this option is enabled, the driver cannot execute parameterized queries.
- By default, the driver applies transformations to the queries emitted by an application to convert the queries into an equivalent form in HiveQL. If the application is Spark-aware and already emits HiveQL, then turning off the translation avoids the additional overhead of query transformation.
- 3. To defer query execution to SQLExecute, select the **Fast SQLPrepare** check box.
- 4. To allow driver-wide configurations to take precedence over connection and DSN settings, select the **Driver Config Take Precedence** check box.
- 5. To use the asynchronous version of the API call against Spark for executing a query, select the **Use Async Exec** check box.
- 6. To retrieve table names from the database by using the SHOW TABLES query, select the **Get Tables With Query** check box.

#### Note:

This option is applicable only when connecting to Spark Thrift Server.

- To enable the driver to return SQL\_WVARCHAR instead of SQL\_VARCHAR for STRING and VARCHAR columns, and SQL\_WCHAR instead of SQL\_CHAR for CHAR columns, select the Unicode SQL Character Types check box.
- 8. To enable the driver to return the spark\_system table for catalog function calls such as SQLTables and SQLColumns, select the **Show System Table** check box.

- 9. To have the driver automatically attempt to reconnect to the server if communications are lost, select **Enable Auto Reconnect**.
- 10. To specify which mechanism the driver uses by default to handle Kerberos authentication, do one of the following:
  - 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.
- 11. To enable the driver to automatically open a new session when the existing session is no longer valid, select the **Invalid Session Auto Recover** check box.

This option is applicable only when connecting to Spark Thrift Server.

- 12. In the **Rows Fetched Per Block** field, type the number of rows to be fetched per block.
- 13. In the **Default String Column Length** field, type the maximum data length for STRING columns.
- 14. In the **Binary Column Length** field, type the maximum data length for BINARY columns.
- 15. In the **Decimal Column Scale** field, type the maximum number of digits to the right of the decimal point for numeric data types.
- 16. In the **Async Exec Poll Interval (ms)** field, type the time in milliseconds between each poll for the query execution status.

Note:

This option is applicable only to HDInsight clusters.

17. In the **Socket Timeout** field, type the number of seconds that an operation can remain idle before it is closed.

#### Note:

This option is applicable only when asynchronous query execution is being used against Spark Thrift Server instances.

18. To save your settings and close the Advanced Options dialog box, click **OK**.

### **Configuring HTTP Options on Windows**

You can configure options such as custom headers when using the HTTP transport protocol in the Thrift layer. For information about how to determine if your Spark server supports the HTTP transport protocol, see Authentication Mechanisms on page 61.

The following instructions describe how to configure HTTP options in a DSN and in the driver configuration tool. You can specify the connection settings described below in a DSN, in a connection string, or as driver-wide settings. Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over driver-wide settings.

#### To configure HTTP options on Windows:

- 1. Choose one:
  - If you are configuring HTTP for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click Configure, and then make sure that the Thrift Transport option is set to HTTP.
  - Or, if you are configuring HTTP for a DSN-less connection, open the Simba Spark ODBC Driver Configuration tool and then make sure that the Thrift Transport option is set to **HTTP**.
- 2. To access HTTP options, click HTTP Options.

#### Note:

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

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

### **Configuring SSL Verification on Windows**

If you are connecting to a Spark server that has Secure Sockets Layer (SSL) enabled, you can configure the driver to connect to an SSL-enabled socket. When using SSL to connect to a server, the driver supports identity verification between the client (the driver itself) and the server.

The following instructions describe how to configure SSL in a DSN and in the driver configuration tool. You can specify the connection settings described below in a DSN, in a connection string, or as driver-wide settings. Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over driver-wide settings.

#### To configure SSL verification on Windows:

- 1. Choose one:
  - To access SSL options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **SSL Options**.
  - Or, to access advanced options for a DSN-less connection, open the Simba Spark ODBC Driver Configuration tool, and then click **SSL Options**.
- 2. Select the Enable SSL check box.
- 3. To allow authentication using self-signed certificates that have not been added to the list of trusted certificates, select the **Allow Self-signed Server Certificate** check box.
- 4. To allow the common name of a CA-issued SSL certificate to not match the host name of the Spark server, select the **Allow Common Name Host Name Mismatch** check box.
- 5. To specify the CA certificates that you want to use to verify the server, do one of the following:
  - To verify the server using the trusted CA certificates from a specific .pem file, specify the full path to the file in the **Trusted Certificates** field and clear the **Use System Trust Store** check box.
  - Or, to use the trusted CA certificates . pem file that is installed with the driver, leave the **Trusted Certificates** field empty, and clear the **Use System Trust Store** check box.
  - Or, to use the Windows trust store, select the **Use System Trust Store** check box.

#### Important:

- If you are using the Windows trust store, make sure to import the trusted CA certificates into the trust store.
- If the trusted CA supports certificate revocation, select the **Check Certificate Revocation** check box.
- 6. To configure two-way SSL verification, select the **Two Way SSL** check box and then do the following:
  - a. In the **Client Certificate File** field, specify the full path of the PEM file containing the client's certificate.
  - b. In the **Client Private Key File** field, specify the full path of the file containing the client's private key.
  - c. If the private key file is protected with a password, type the password in the Client Private Key Password field. To save the password, select the Save Password (Encrypted) check box.

#### Important:

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

- 7. From the **Minimum TLS** drop-down list, select the minimum version of TLS to use when connecting to your data store.
- 8. To save your settings and close the SSL Options dialog box, click **OK**.

### **Configuring Server-Side Properties on Windows**

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

The following instructions describe how to configure server-side properties in a DSN and in the driver configuration tool. You can specify the connection settings described below in a DSN, in a connection string, or as driver-wide settings. Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over driver-wide settings.

#### To configure server-side properties on Windows:

- 1. Choose one:
  - To configure server-side properties for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click Configure, then click Advanced Options, and then click Server Side Properties.
  - Or, to configure server-side properties for a DSN-less connection, open the Simba Spark ODBC Driver Configuration tool, then click Advanced Options, and then click Server Side Properties.
- 2. To create a server-side property, click **Add**, then type appropriate values in the **Key** and **Value** fields, and then click **OK**.

#### Note:

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

- 3. To edit a server-side property, select the property from the list, then click **Edit**, then update the **Key** and **Value** fields as needed, and then click **OK**.
- 4. To delete a server-side property, select the property from the list, and then click **Remove**. In the confirmation dialog box, click **Yes**.
- 5. To change the method that the driver uses to apply server-side properties, do one of the following:

- To configure the driver to apply each server-side property by executing a query when opening a session to the Spark server, select the **Apply Server Side Properties With Queries** check box.
- Or, to configure the driver to use a more efficient method for applying server-side properties that does not involve additional network round-tripping, clear the **Apply Server Side Properties With Queries** check box.

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

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

### **Configuring Logging Options on Windows**

To help troubleshoot issues, you can enable logging. In addition to functionality provided in the Simba Spark 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 Spark 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.

Logging Level	Description
FATAL	Logs severe error events that lead the driver to abort.
ERROR	Logs error events that might allow the driver to continue running.
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. In the Max Number Files field, type the maximum number of log files to keep.

After the maximum number of log files is reached, each time an additional file is created, the driver deletes the oldest log file.

5. In the **Max File Size** field, type the maximum size of each log file in megabytes (MB).

#### Note:

After the maximum file size is reached, the driver creates a new file and continues logging.

- 6. Click **OK**.
- 7. Restart your ODBC application to make sure that the new settings take effect.

The Simba Spark ODBC Driver produces the following log files at the location you specify in the Log Path field:

• A simbasparkodbcdriver.log file that logs driver activity that is not specific to a connection.

• A simbasparkodbcdriver\_connection\_[Number].log file 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.

#### 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.

### **Configuring Kerberos Authentication for Windows**

#### **Active Directory**

The Simba Spark ODBC Driver supports Active Directory Kerberos on Windows. There are two prerequisites for using Active Directory Kerberos on Windows:

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

#### **MIT Kerberos**

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

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

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

- 1. Download the appropriate Kerberos installer:
  - For a 64-bit machine, use the following download link from the MIT Kerberos website: http://web.mit.edu/kerberos/dist/kfw/4.0/kfw-4.0.1amd64.msi.
  - For a 32-bit machine, use the following download link from the MIT Kerberos website: http://web.mit.edu/kerberos/dist/kfw/4.0/kfw-4.0.1i386.msi.

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

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

#### Setting Up the Kerberos Configuration File

Settings for Kerberos are specified through a configuration file. You can set up the configuration file as an .ini file in the default location, which is the C:\ProgramData\MIT\Kerberos5 directory, or as a .conf file in a custom location.

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

#### Note:

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

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

- 1. Obtain a krb5.conf configuration file. You can obtain this file from your Kerberos administrator, or from the /etc/krb5.conf folder on the machine that is hosting the Spark Thrift Server instance.
- 2. Rename the configuration file from krb5.conf to krb5.ini.
- 3. Copy the krb5.ini file to the C:\ProgramData\MIT\Kerberos5 directory and overwrite the empty sample file.

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

- 1. Obtain a krb5.conf configuration file. You can obtain this file from your Kerberos administrator, or from the /etc/krb5.conf folder on the machine that is hosting the Spark Thrift Server instance.
- 2. Place the krb5.conf file in an accessible directory and make note of the full path name.
- 3. Open the System window:
  - If you are using Windows 7 or earlier, click Start <sup>(1)</sup>, then right-click Computer, and then click Properties.

- Or, if you are using Windows 8 or later, right-click **This PC** on the Start screen, and then click **Properties**.
- 4. Click Advanced System Settings.
- 5. In the System Properties dialog box, click the **Advanced** tab and then click **Environment Variables**.
- 6. In the Environment Variables dialog box, under the System Variables list, click **New**.
- 7. In the New System Variable dialog box, in the **Variable Name** field, type **KRB5\_ CONFIG**.
- 8. In the Variable Value field, type the full path to the krb5.conf file.
- 9. Click **OK** to save the new variable.
- 10. Make sure that the variable is listed in the System Variables list.
- 11. Click **OK** to close the Environment Variables dialog box, and then click **OK** to close the System Properties dialog box.

#### Setting Up the Kerberos Credential Cache File

Kerberos uses a credential cache to store and manage credentials.

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

- 1. Create a directory where you want to save the Kerberos credential cache file. For example, create a directory named C:\temp.
- 2. Open the System window:
  - If you are using Windows 7 or earlier, click Start <sup>1</sup>, then right-click Computer, and then click Properties.
  - Or, if you are using Windows 8 or later, right-click **This PC** on the Start screen, and then click **Properties**.
- 3. Click Advanced System Settings.
- 4. In the System Properties dialog box, click the **Advanced** tab and then click **Environment Variables**.
- 5. In the Environment Variables dialog box, under the System Variables list, click **New**.
- 6. In the New System Variable dialog box, in the **Variable Name** field, type **KRB5CCNAME**.
- 7. In the Variable Value field, type the path to the folder you created above, and then append the file name krb5cache. For example, if you created the folder C:\temp, then type C:\temp\krb5cache.

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

- 8. Click **OK** to save the new variable.
- 9. Make sure that the variable appears in the System Variables list.
- 10. Click **OK** to close the Environment Variables dialog box, and then click **OK** to close the System Properties dialog box.
- 11. To make sure that Kerberos uses the new settings, restart your machine.

#### Obtaining a Ticket for a Kerberos Principal

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

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

- 1. Open MIT Kerberos Ticket Manager.
- 2. In MIT Kerberos Ticket Manager, click Get Ticket.
- 3. In the Get Ticket dialog box, type your principal name and password, and then click **OK**.

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

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

- 1. Open a command prompt:
  - If you are using Windows 7 or earlier, click Start 

     then click All
     Programs, then click Accessories, and then click Command Prompt.
  - If you are using Windows 8 or later, click the arrow button at the bottom of the Start screen, then find the Windows System program group, and then click Command Prompt.
- 2. In the Command Prompt, type a command using the following syntax:

```
kinit -k -t [KeytabPath] [Principal]
```

[KeytabPath] is the full path to the keytab file. For example:

```
C:\mykeytabs\myUser.keytab.
```

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

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

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

Krbcache is the Kerberos cache file, not a directory.

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

#### Note:

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

- 1. Open a command prompt:
  - If you are using Windows 7 or earlier, click Start 
     , then click All Programs, then click Accessories, and then click Command Prompt.
  - If you are using Windows 8 or later, click the arrow button at the bottom of the Start screen, then find the Windows System program group, and then click Command Prompt.
- 2. In the Command Prompt, type a command using the following syntax:

```
kinit -k [principal]
```

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

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

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

Krbcache is the Kerberos cache file, not a directory.

### Verifying the Driver Version Number on Windows

If you need to verify the version of the Simba Spark 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 Spark.

2. Click the **Drivers** tab and then find the Simba Spark 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

The Simba Spark ODBC Driver supports Apache Spark versions 1.6 through 2.3.

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 10.12, 10.13, or 10.14
- 100 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 Spark ODBC Driver is available for macOS as a .dmg file named SimbaSparkODBC.dmg. The driver supports both 32- and 64-bit client applications.

To install the Simba Spark ODBC Driver on macOS:

- 1. Double-click **SimbaSparkODBC.dmg** to mount the disk image.
- 2. Double-click SimbaSparkODBC.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/spark 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/spark/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 42.

### Verifying the Driver Version Number on macOS

If you need to verify the version of the Simba Spark 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.sparkodbc
```

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

### Linux Driver

The Linux driver is available as an RPM file and as a tarball package.

### **Linux System Requirements**

The Simba Spark ODBC Driver supports Apache Spark versions 1.6 through 2.3.

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
- 150 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
- All of the following libsasl libraries installed:
  - cyrus-sasl-2.1.22-7 or later
  - cyrus-sasl-gssapi-2.1.22-7 or later
  - cyrus-sasl-plain-2.1.22-7 or later

#### Note:

If the package manager in your Linux distribution cannot resolve the dependencies automatically when installing the driver, then download and manually install the packages.

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:

- SimbaSparkODBC-32bit-[Version]-[Release].i686.rpm for the 32bit driver
- SimbaSparkODBC-[Version]-[Release].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.

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

#### To install the Simba Spark 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 Spark ODBC Driver files are installed in the /opt/simba/spark directory.

#### Note:

If the package manager in your Linux distribution cannot resolve the libsasl dependencies automatically when installing the driver, then download and manually install the packages.

4. If you received a license file through email, then copy the license file into the /opt/simba/spark/lib/32 or /opt/simba/spark/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 42.

### Installing the Driver Using the Tarball Package

The Simba Spark ODBC Driver is available as a tarball package named SimbaSparkODBC-[Version].[Release]-Linux.tar.gz, where [Version] is the version number of the driver and [Release] is the release number for this version of the driver. The package contains both the 32-bit and 64-bit versions of the driver.

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

#### To install the driver using the tarball package:

- 1. Log in as the root user, and then navigate to the folder containing the tarball package.
- 2. Run the following command to extract the package and install the driver:

tar --directory=/opt -zxvf [TarballName]

Where *[TarballName]* is the name of the tarball package containing the driver.

The Simba Spark ODBC Driver files are installed in the opt/simba/spark directory.

3. If you received a license file through email, then copy the license file into the opt/simba/spark/lib/32 or opt/simba/spark/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 42.

### Verifying the Driver Version Number on Linux

If you need to verify the version of the Simba Spark 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 SimbaSparkODBC
  - rpm -qa | grep SimbaSparkODBC

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

## AIX Driver

## **AIX System Requirements**

The Simba Spark ODBC Driver supports Apache Spark versions 1.6 through 2.3.

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:

- IBM AIX 5.3, 6.1, or 7.1
- 150 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 on AIX

On 64-bit editions of AIX, 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 the version of the driver that matches the bitness of the client application:

- SimbaSparkODBC-32bit-[Version]-[Release].ppc.rpm for the 32-bit driver
- SimbaSparkODBC-[Version]-[Release].ppc.rpm for the 64-bit driver

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

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

#### To install the Simba Spark ODBC Driver on AIX:

- 1. Log in as the root user, and then navigate to the folder containing the RPM package for the driver.
- 2. Run the following command from the command line, where [RPMFileName] is the file name of the RPM package:

rpm --install [RPMFileName]

The Simba Spark ODBC Driver files are installed in the /opt/simba/spark directory.

3. If you received a license file via email, then copy the license file into the /opt/simba/spark/lib/32 or /opt/simba/spark/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 42.

# Verifying the Driver Version Number on AIX

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

#### To verify the driver version number on AIX:

> At the command prompt, run the following command:

rpm -qa | grep SimbaSparkODBC

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

## Solaris Driver

## **Solaris System Requirements**

The Simba Spark ODBC Driver supports Apache Spark versions 1.6 through 2.3.

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:

- Solaris 10 or later (sparc and sparc64 editions are supported)
- 150 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 on Solaris**

The Simba Spark ODBC Driver is available for Solaris as a tarball package named Simba Spark ODBC Driver\_Solaris-gcc\_[Version].[Release]\_Solaris.tar.gz, where [Version] is the version number of the driver and [Release] is the release number for this version of the driver. The package contains both the 32-bit and 64-bit versions of the driver.

On sparc64 editions of Solaris, you can execute both sparc and sparc64 applications. However, sparc64 applications must use 64-bit drivers, and sparc applications must use 32-bit drivers. Make sure that you use the version of the driver that matches the bitness of the client application. You can install both versions of the driver on the same machine.

#### To install the Simba Spark ODBC Driver on Solaris:

- 1. Log in as the root user, and then navigate to the folder containing the tarball package.
- 2. Run the following command to extract the package and install the driver:

```
tar --directory=/opt -zxvf [TarballName]
```

Where [TarballName] is the name of the tarball package containing the driver.

The Simba Spark ODBC Driver files are installed in the /opt/simba/spark/ directory.

3. If you received a license file via email, then copy the license file into the /opt/simba/spark/lib/32 or /opt/simba/spark/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 42.

# Verifying the Driver Version Number on Solaris

If you need to verify the version of the Simba Spark ODBC Driver that is installed on your Solaris machine, you can query the version number through the command-line interface.

#### To verify the driver version number on Solaris:

> At the command prompt, run the following command:

rpm -qa | grep SimbaSparkODBC

The command returns information about the Simba Spark 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 Spark 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 42.
- 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 43.

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 45.

## 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 or AIX

If you are using a Linux or AIX 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 or AIX shell documentation.

#### Solaris

If you are using a Solaris machine, then set the LD\_LIBRARY\_PATH environment variable to include the paths to the ODBC driver manager libraries and the third-party libraries that are installed with the driver. For example, if the driver manager libraries are installed in /usr/local/lib and the 32-bit driver is installed in /opt/simba/spark, 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:/opt/simba/spark/lib/32
```

For information about setting an environment variable permanently, refer to the Solaris 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.sparkodbc.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 SIMBASPARKINI to the full path and file name of the simba.sparkodbc.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 SIMBASPARKINI to the full path and file name of the simba.sparkodbc.ini file.

For example, if your odbc.ini and odbcinst.ini files are located in /usr/local/odbc and your simba.sparkodbc.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 SIMBASPARKINI=/etc/simba.sparkodbc.ini
```

For unixODBC:

```
export ODBCINI=/usr/local/odbc/odbc.ini
export ODBCSYSINI=/usr/local/odbc
export SIMBASPARKINI=/etc/simba.sparkodbc.ini
```

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

- 1. If the SIMBASPARKINI 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.sparkodbc.ini.
- 3. The driver searches the current working directory of the application for a file named simba.sparkodbc.ini.
- 4. The driver searches the home directory for a hidden file named .simba.sparkodbc.ini (prefixed with a period).
- 5. The driver searches the /etc directory for a file named simba.sparkodbc.ini.

# Configuring ODBC Connections on a Non-Windows Machine

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

- Creating a Data Source Name on a Non-Windows Machine on page 45
- Configuring a DSN-less Connection on a Non-Windows Machine on page 48
- Configuring Authentication on a Non-Windows Machine on page 51
- Configuring SSL Verification on a Non-Windows Machine on page 54
- Configuring Server-Side Properties on a Non-Windows Machine on page 55
- Configuring Logging Options on a Non-Windows Machine on page 56
- Setting Driver-Wide Configuration Options on a Non-Windows Machine on page
   58
- Testing the Connection on a Non-Windows Machine on page 59

# Creating a Data Source Name on a Non-Windows Machine

Typically, after installing the Simba Spark ODBC Driver, you need to create a Data Source Name (DSN). A DSN is a data structure that stores connection information so that it can be used by the driver to connect to Spark.

You can specify connection settings in a DSN (in the odbc.ini file), in a connection string, or as driver-wide settings (in the simba.sparkodbc.ini file). Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over driver-wide settings.

The following instructions describe how to create a DSN by specifying connection settings in the odbc.ini file. 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.

For information about specifying settings in a connection string, see Configuring a DSN-less Connection on a Non-Windows Machine on page 48 and Using a Connection String on page 63. For information about driver-wide settings, see Setting Driver-Wide Configuration Options on a Non-Windows Machine on page 58.

#### 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 Spark ODBC Driver
```

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

```
[ODBC Data Sources]
Sample DSN=Simba Spark 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/spark/lib/libsparkodbc_
sbu.dylib
```

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

Driver=/opt/simba/spark/lib/32/libsparkodbc sb32.so

- b. Set the SparkServerType property to one of the following values:
  - If you are running Shark 0.8.1 or earlier, set the property to 1.
  - If you are running Shark 0.9 or Spark 1.1 or later, set the property to 3.

For example:

```
SparkServerType=1
```

c. Set the Host property to the IP address or host name of the server. For example: Host=192.168.222.160

d. Set the Port property to the number of the TCP port that the server uses to listen for client connections.

For example:

Port=10000

- e. 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 Machine on page 51.
- f. If you want to connect to the server through SSL, then enable SSL and specify the certificate information. For more information, see Configuring SSL Verification on a Non-Windows Machine on page 54.
- g. If you want to configure server-side properties, then set them as key-value pairs using a special syntax. For more information, see Configuring Server-Side Properties on a Non-Windows Machine on page 55.
- h. 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 Spark ODBC Driver, see Driver Configuration Options on page 73.
- 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 43.

For example, the following is an odbc.ini configuration file for macOS containing a DSN that connects to a Spark Thrift Server instance and authenticates the connection using a user name and password:

```
[ODBC Data Sources]
Sample DSN=Simba Spark ODBC Driver
[Sample DSN]
Driver=/Library/simba/spark/lib/libsparkodbc_sbu.dylib
SparkServerType=3
UID=username
PWD=userpassword
```

As another example, the following is an odbc.ini configuration file for a 32-bit driver on a Linux/AIX/Solaris machine, containing a DSN that connects to a SparkThrift Server instance:

```
[ODBC Data Sources]
Sample DSN=Simba Spark ODBC Driver 32-bit
[Sample DSN]
Driver=/opt/simba/spark/lib/32/libsparkodbc_sb32.so
SparkServerType=3
Host=192.168.222.160
Port=10000
```

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 Spark 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/spark/lib/libsparkodbc_
sbu.dylib
```

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

Driver=/opt/simba/spark/lib/32/libsparkodbc sb32.so

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

For example:

Description=Simba Spark 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 43.

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

```
[ODBC Drivers]
Simba Spark ODBC Driver=Installed
[Simba Spark ODBC Driver]
Description=Simba Spark ODBC Driver
Driver=/Library/simba/spark/lib/libsparkodbc_sbu.dylib
```

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

```
[ODBC Drivers]
Simba Spark ODBC Driver 32-bit=Installed
Simba Spark ODBC Driver 64-bit=Installed
```

```
[Simba Spark ODBC Driver 32-bit]
Description=Simba Spark ODBC Driver (32-bit)
Driver=/opt/simba/spark/lib/32/libsparkodbc_sb32.so
[Simba Spark ODBC Driver 64-bit]
Description=Simba Spark ODBC Driver (64-bit)
Driver=/opt/simba/spark/lib/64/libsparkodbc 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 63.

For instructions about configuring specific connection features, see the following:

- Configuring Authentication on a Non-Windows Machine on page 51
- Configuring SSL Verification on a Non-Windows Machine on page 54
- Configuring Server-Side Properties on a Non-Windows Machine on page 55

For detailed information about all the connection properties that the driver supports, see Driver Configuration Options on page 73.

## Configuring Service Discovery Mode on a Non-Windows Machine

You can configure the Simba Spark ODBC Driver to discover Spark services through DataStax AOSS.

You can set the connection properties described below in a connection string, in a DSN (in the odbc.ini file), or as a driver-wide setting (in the simba.sparkodbc.ini file). Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over driver-wide settings.

#### To enable Service Discovery Mode on a non-Windows machine:

- 1. Set the ServiceDiscoveryMode connection attribute to AOSS.
- 2. Set the Host connection attribute to a comma-separated list of AOSS endpoints. Use the following format, where [AOSS\_Endpoint] is the IP address or host name of the AOSS endpoint, and [AOSS\_Port] is the number of the TCP port that the AOSS end point uses to listen for client connections:

```
[AOSS_Endpoint1]:[AOSS_Port1],[AOSS_Endpoint2]:[AOSS_
Port2]
```

## **Configuring Authentication on a Non-Windows Machine**

Some Spark Thrift Server instances are configured to require authentication for access. To connect to a Spark server, you must configure the Simba Spark ODBC Driver to use the authentication mechanism that matches the access requirements of the server and provides the necessary credentials.

For information about how to determine the type of authentication your Spark server requires, see Authentication Mechanisms on page 61.

You can set the connection properties for authentication in a connection string, in a DSN (in the odbc.ini file), or as a driver-wide setting (in the simba.sparkodbc.ini file). Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over driver-wide settings.

Depending on the authentication mechanism you use, there might be additional connection attributes that you must define. For more information about the attributes involved in configuring authentication, see Driver Configuration Options on page 73.

## Using No Authentication

When connecting to a Spark server of type Shark Server, you must use No Authentication. When you use No Authentication, Binary is the only Thrift transport protocol that is supported.

#### To configure a connection without authentication:

- 1. Set the AuthMech connection attribute to 0.
- 2. If the Spark server is configured to use SSL, then configure SSL for the connection. For more information, see Configuring SSL Verification on a Non-Windows Machine on page 54.

## Using Kerberos

Kerberos must be installed and configured before you can use this authentication mechanism. For more information, refer to the MIT Kerberos Documentation: http://web.mit.edu/kerberos/krb5-latest/doc/.

#### To configure Kerberos authentication:

- 1. Set the AuthMech connection attribute to 1.
- 2. Choose one:
  - To use the default realm defined in your Kerberos setup, do not set the KrbRealm attribute.
  - Or, if your Kerberos setup does not define a default realm or if the realm of your Spark server is not the default, then set the appropriate realm using the KrbRealm attribute.
- 3. Set the KrbHostFQDN attribute to the fully qualified domain name of the Spark Thrift Server host.

### Note:

To use the Spark server host name as the fully qualified domain name for Kerberos authentication, set KrbHostFQDN to HOST.

- 4. Set the KrbServiceName attribute to the service name of the Spark Thrift Server.
- 5. To allow the driver to pass your credentials directly to the server for use in authentication, set DelegateKrbCreds to 1.
- 6. Set the ThriftTransport connection attribute to the transport protocol to use in the Thrift layer.

#### Important:

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

7. If the Spark server is configured to use SSL, then configure SSL for the connection. For more information, see Configuring SSL Verification on a Non-Windows Machine on page 54.

### Using User Name

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

This authentication mechanism is available only for Spark Thrift Server on non-HDInsight distributions. Most default configurations of require User Name authentication. When you use User Name authentication, SSL is not supported and SASL is the only Thrift transport protocol available.

#### To configure User Name authentication:

- 1. Set the AuthMech connection attribute to 2.
- 2. Set the UID attribute to an appropriate user name for accessing the Spark server.

## Using User Name And Password

This authentication mechanism requires a user name and a password.

This authentication mechanism is available only for Spark Thrift Server on non-HDInsight distributions.

#### To configure User Name And Password authentication:

- 1. Set the AuthMech connection attribute to 3.
- 2. Set the UID attribute to an appropriate user name for accessing the Spark server.
- 3. Set the PWD attribute to the password corresponding to the user name you provided above.
- 4. Set the ThriftTransport connection attribute to the transport protocol to use in the Thrift layer.
- 5. If the Spark server is configured to use SSL, then configure SSL for the connection. For more information, see Configuring SSL Verification on a Non-Windows Machine on page 54.

## Using Windows Azure HDInsight Emulator

This authentication mechanism is available only for Spark Thrift Server instances running on Windows Azure HDInsight Emulator. When you use this authentication mechanism, SSL is not supported and HTTP is the only Thrift transport protocol available.

# To configure a connection to a Spark server on Windows Azure HDInsight Emulator:

- 1. Set the AuthMech connection attribute to 5.
- 2. Set the HTTPPath attribute to the partial URL corresponding to the Spark server.
- 3. Set the UID attribute to an appropriate user name for accessing the Spark server.
- 4. Set the PWD attribute to the password corresponding to the user name you provided above.
- 5. If necessary, you can create custom HTTP headers. For more information, see http.header. on page 105.

## Using Windows Azure HDInsight Service

This authentication mechanism is available only for Spark Thrift Server on HDInsight distributions. When you use this authentication mechanism, you must enable SSL, and HTTP is the only Thrift transport protocol available.

# To configure a connection to a Spark server on Windows Azure HDInsight Service:

- 1. Set the AuthMech connection attribute to 6.
- 2. Set the HTTPPath attribute to the partial URL corresponding to the Spark server.
- 3. Set the UID attribute to an appropriate user name for accessing the Spark server.
- 4. Set the PWD attribute to the password corresponding to the user name you typed above.
- 5. If necessary, you can create custom HTTP headers. For more information, see http.header. on page 105.
- 6. Configure SSL settings as needed. For more information, see Configuring SSL Verification on a Non-Windows Machine on page 54.
- 7. Choose one:
  - To configure the driver to load SSL certificates from a specific file, set the TrustedCerts attribute to the path of the file.
  - Or, to use the trusted CA certificates PEM file that is installed with the driver, do not specify a value for the TrustedCerts attribute.

## **Configuring SSL Verification on a Non-Windows** Machine

If you are connecting to a Spark server that has Secure Sockets Layer (SSL) enabled, you can configure the driver to connect to an SSL-enabled socket. When using SSL to connect to a server, the driver supports identity verification between the client (the driver itself) and the server.

You can set the connection properties described below in a connection string, in a DSN (in the odbc.ini file), or as a driver-wide setting (in the simba.sparkodbc.ini file). Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over driver-wide settings.

#### To configure SSL verification on a non-Windows machine:

- 1. To enable SSL connections, set the SSL attribute to 1.
- 2. To allow authentication using self-signed certificates that have not been added to the list of trusted certificates, set the AllowSelfSignedServerCert attribute to 1.
- 3. To allow the common name of a CA-issued SSL certificate to not match the host name of the Spark server, set the CAIssuedCertNamesMismatch attribute to 1.
- 4. Choose one:
  - To configure the driver to load SSL certificates from a specific .pem file when verifying the server, set the TrustedCerts attribute to the full path of the .pem file.
  - Or, to use the trusted CA certificates .pem file that is installed with the driver, do not specify a value for the TrustedCerts attribute.
- 5. To configure two-way SSL verification, set the TwoWaySSL attribute to 1 and then do the following:
  - a. Set the ClientCert attribute to the full path of the .pem file containing the client's certificate.
  - b. Set the ClientPrivateKey attribute to the full path of the file containing the client's private key.
  - c. If the private key file is protected with a password, set the ClientPrivateKeyPassword attribute to the password.
- 6. To specify the minimum version of TLS to use, set the Min\_TLS property to the minimum version of TLS. Supported options include 1.0 for TLS 1.0, 1.1 for TLS 1.1, and 1.2 for TLS 1.2.

## Configuring Server-Side Properties on a Non-Windows Machine

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

You can set the connection properties described below in a connection string, in a DSN (in the odbc.ini file), or as a driver-wide setting (in the simba.sparkodbc.ini file). Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over driver-wide settings.

#### To configure server-side properties on a non-Windows machine:

1. To set a server-side property, use the syntax SSP\_[SSPKey] = [SSPValue], where [SSPKey] is the name of the server-side property and [SSPValue] is the

value to specify for that property.

#### Note:

- When setting a server-side property in a connection string, it is recommended that you enclose the value in braces ({ }) to make sure that special characters can be properly escaped.
- For a list of all Hadoop and Spark server-side properties that your implementation supports, type set -v at the Spark CLI command line. You can also execute the set -v query after connecting using the driver.
- 2. To change the method that the driver uses to apply server-side properties, do one of the following:
  - To configure the driver to apply each server-side property by executing a query when opening a session to the Spark server, set the ApplySSPWithQueries property to 1.
  - Or, to configure the driver to use a more efficient method for applying server-side properties that does not involve additional network round-tripping, set the ApplySSPWithQueries property to 0.

### Note:

The more efficient method is not available for Shark Server, and it might not be compatible with some Spark Thrift Server builds. If the server-side properties do not take effect when the ApplySSPWithQueries property is set to 0, then set it to 1.

3. To disable the driver's default behavior of converting server-side property key names to all lower-case characters, set the LCaseSspKeyName property to 0.

## **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.

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

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

#### To enable logging on a non-Windows machine:

- 1. Open the simba.sparkodbc.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  ${\tt LogPath}$  key to the full path to the folder where you want to save log files.
- 4. Set the LogFileCount key to the maximum number of log files to keep.

#### Note:

After the maximum number of log files is reached, each time an additional file is created, the driver deletes the oldest log file.

5. Set the LogFileSize key to the maximum size of each log file in megabytes (MB).

#### Note:

After the maximum file size is reached, the driver creates a new file and continues logging.

- 6. Save the simba.sparkodbc.ini configuration file.
- 7. Restart your ODBC application to make sure that the new settings take effect.

The Simba Spark ODBC Driver produces the following log files at the location you specify using the LogPath key:

- A simbasparkodbcdriver.log file that logs driver activity that is not specific to a connection.
- A simbasparkodbcdriver\_connection\_[Number].log file 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.

#### To disable logging on a non-Windows machine:

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

# Setting Driver-Wide Configuration Options on a Non-Windows Machine

When you specify connection settings in a DSN or connection string, those settings apply only when you connect to Spark using that particular DSN or string. As an alternative, you can specify settings that apply to every connection that uses the Simba Spark ODBC Driver by configuring them in the simba.sparkodbc.ini file.

#### Note:

Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over driver-wide settings.

#### To set driver-wide configuration options on a non-Windows machine:

- 1. In a text editor, open the simba.sparkodbc.ini configuration file.
- 2. In the [Driver] section, specify configuration options as key-value pairs. Start a new line for each key-value pair.

For example, to enable User Name authentication using "simba" as the user name, type the following:

```
AuthMech=2
UID=simba
```

For detailed information about all the configuration options supported by the

driver, see Driver Configuration Options on page 73.

3. Save the simba.sparkodbc.ini configuration file.

## 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 .

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.

## Authentication Mechanisms

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

# **Shark Server**

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

# Spark Thrift Server on an HDInsight Distribution

If you are connecting to HDInsight Emulator running on Windows Azure, then you must use the Windows Azure HDInsight Emulator mechanism.

If you are connecting to HDInsight Service running on Windows Azure, then you must use the Windows Azure HDInsight Service mechanism.

## Spark Thrift Server on a non-HDInsight Distribution

Note:

Most default configurations of Spark Thrift Server on non-HDInsight distributions require User Name authentication.

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

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

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

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

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

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

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

For detailed instructions on how to configure authentication when using the Windows driver, see Configuring Authentication on Windows on page 14.

For detailed instructions on how to configure authentication when using a non-Windows driver, see Configuring Authentication on a Non-Windows Machine on page 51.

## 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 73.

## **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:

- [AOSS\_Endpoint] is the IP address or host name of the DataStax AOSS endpoint.
- [AOSS\_Port] is the number of the TCP port that the DataStax AOSS endpoint uses to listen for client connections.
- [DomainName] is the fully qualified domain name of the Spark server host.
- [PortNumber] is the number of the TCP port that the Spark server uses to listen for client connections.
- [Realm] is the Kerberos realm of the Spark server host.

- [Server] is the IP address or host name of the Spark server to which you are connecting.
- [ServerURL] is the partial URL corresponding to the Spark server.
- [ServiceName] is the Kerberos service principal name of the Spark server.
- [YourPassword] is the password corresponding to your user name.
- [YourUserName] is the user name that you use to access the Spark server.

### Connecting to a Shark Server Instance

The following is the format of a DSN-less connection string that connects to a Shark Server instance:

```
Driver=Simba Spark ODBC Driver;SparkServerType=1;
Host=[Server];Port=[PortNumber];
```

#### For example:

```
Driver=Simba Spark ODBC Driver;SparkServerType=1;
Host=192.168.222.160;Port=10000;
```

## Connecting to a Standard Spark Thrift Server Instance

The following is the format of a DSN-less connection string for a standard connection to a Spark Thrift Server instance. By default, the driver is configured to connect to a Spark Thrift Server instance. Most default configurations of Spark Thrift Server require User Name authentication. When configured to provide User Name authentication, the driver uses **anonymous** as the user name by default.

```
Driver=Simba Spark ODBC Driver;Host=[Server];
Port=[PortNumber];AuthMech=2;
```

For example:

```
Driver=Simba Spark ODBC Driver;Host=192.168.222.160;
Port=10000;AuthMech=2;
```

# Connecting to a Spark Thrift Server Instance Without Authentication

The following is the format of a DSN-less connection string that for a Spark Thrift Server instance that does not require authentication.

```
Driver=Simba Spark ODBC Driver;Host=[Server];
Port=[PortNumber];AuthMech=0;
```

#### For example:

```
Driver=Simba Spark ODBC Driver;Host=192.168.222.160;
Port=10000;AuthMech=0;
```

# Connecting to a Spark Server that Requires Kerberos Authentication

The following is the format of a DSN-less connection string that connects to a Spark Thrift Server instance requiring Kerberos authentication. By default, the driver is configured to connect to a Spark Thrift Server instance.

```
Driver=Simba Spark ODBC Driver;Host=[Server];
Port=[PortNumber];AuthMech=1;KrbRealm=[Realm];
KrbHostFQDN=[DomainName];KrbServiceName=[ServiceName];
```

#### For example:

```
Driver=Simba Spark ODBC Driver;Host=192.168.222.160;
Port=10000;AuthMech=1;KrbRealm=SIMBA;
KrbHostFQDN=localhost.localdomain;KrbServiceName=spark;
```

# Connecting to a Spark Server that Requires User Name And Password Authentication

The following is the format of a DSN-less connection string that connects to a Spark Thrift Server instance requiring User Name and Password authentication. By default, the driver is configured to connect to a Spark Thrift Server instance.

```
Driver=Simba Spark ODBC Driver;Host=[Server];
Port=[PortNumber];AuthMech=3;UID=[YourUserName];
PWD=[YourPassword];
```

#### For example:

```
Driver=Simba Spark ODBC Driver;Host=192.168.222.160;
Port=10000;AuthMech=3;UID=simba;PWD=simba;
```

## Connecting to a Spark Server on Windows Azure HDInsight Emulator

The following is the format of a DSN-less connection string that connects to a Spark Thrift Server instance running on Windows Azure HDInsight Emulator. By default, the

#### driver is configured to connect to a Spark Thrift Server instance.

```
Driver=Simba Spark ODBC Driver;Host=[Server];
Port=[PortNumber];AuthMech=5;UID=[YourUserName];
PWD=[YourPassword];HTTPPath=[ServerURL];
```

#### For example:

```
Driver=Simba Spark ODBC Driver;Host=192.168.222.160;
Port=10000;AuthMech=5;UID=simba;PWD=simba;
HTTPPath=gateway/sandbox/spark;
```

# Connecting to a Spark Server on Windows Azure HDInsight Service

The following is the format of a DSN-less connection string that connects to a Spark Thrift Server instance running on Windows Azure HDInsight Service. By default, the driver is configured to connect to a Spark Thrift Server instance.

```
Driver=Simba Spark ODBC Driver;Host=[Server];
Port=[PortNumber];AuthMech=6;UID=[YourUserName];
PWD=[YourPassword];HTTPPath=[ServerURL];
```

#### For example:

```
Driver=Simba Spark ODBC Driver;Host=192.168.222.160;
Port=10000;AuthMech=6;UID=simba;PWD=simba;
HTTPPath=gateway/sandbox/spark;
```

#### Connecting using Dynamic Service Discovery

The following is the format of a DSN-less connection string that discovers Spark services via the DataStax AOSS service.

```
Driver=Simba Spark ODBC Driver;ServiceDiscoveryMode=AOSS;
Host=[AOSS_Endpoint1]:[AOSS_Port1], [AOSS_Endpoint2]:[AOSS_
Port2];
```

#### For example:

```
Driver=Simba Spark ODBC Driver;ServiceDiscoveryMode=AOSS;
Host=192.168.222.160:9077, 192.168.222.165:9077;
```

## Features

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

- SQL Connector for HiveQL on page 67
- Data Types on page 68
- Timestamp Function Support on page 69
- Catalog and Schema Support on page 69
- spark\_system Table on page 69
- Server-Side Properties on page 70
- Get Tables With Query on page 70
- Active Directory on page 70
- Write-back on page 71
- Dynamic Service Discovery using Datastax AOSS on page 71
- Security and Authentication on page 71

## **SQL Connector for HiveQL**

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

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

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

# **Data Types**

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

The following table lists the supported data type mappings.

Spark Type	SQL Type
BIGINT	SQL_BIGINT
BINARY	SQL_VARBINARY
BOOLEAN	SQL_BIT
CHAR(n)	SQL_CHAR
DATE	SQL_TYPE_DATE
DECIMAL	SQL_DECIMAL
DECIMAL(p,s)	SQL_DECIMAL
DOUBLE	SQL_DOUBLE
FLOAT	SQL_REAL
INT	SQL_INTEGER
SMALLINT	SQL_SMALLINT
STRING	SQL_VARCHAR
TIMESTAMP	SQL_TYPE_TIMESTAMP
TINYINT	SQL_TINYINT
VARCHAR(n)	SQL_VARCHAR

#### Note:

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

## **Timestamp Function Support**

The Simba Spark ODBC Driver supports the following ODBC functions for working with data of type TIMESTAMP:

- **TIMESTAMPADD**: You can call this function to increment a TIMESTAMP value by a specified interval of time.
- **TIMESTAMPDIFF**: You can call this function to calculate the interval of time between two specified TIMESTAMP values.

The types of time intervals that are supported for these functions might vary depending on the Spark server version that you are connecting to. To return a list of the intervals supported for TIMESTAMPADD, call the SQLGetInfo catalog function using SQL\_ TIMEDATE\_ADD\_INTERVALS as the argument. Similarly, to return a list of the intervals supported for TIMESTAMPDIFF, call SQLGetInfo using SQL\_TIMEDATE\_ DIFF\_INTERVALS as the argument.

#### Note:

The SQL\_TSI\_FRAC\_SECOND interval is not supported by Spark.

# **Catalog and Schema Support**

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

## spark\_system Table

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

```
SELECT * FROM SPARK.spark_system.spark_system WHERE envkey
LIKE '%spark%'
```

The above query returns all of the Spark system environment entries whose key contains the word "spark". A special query, set -v, is executed to fetch system environment information. Some versions of Spark do not support this query. For

versions of Spark that do not support querying system environment information, the driver returns an empty result set.

## **Server-Side Properties**

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

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

For more information about setting server-side properties when using the Windows driver, see Configuring Server-Side Properties on Windows on page 23. For information about setting server-side properties when using the driver on a non-Windows platform, see Configuring Server-Side Properties on a Non-Windows Machine on page 55.

## **Get Tables With Query**

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

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

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

# **Active Directory**

The Simba Spark ODBC Driver supports Active Directory Kerberos on Windows. There are two prerequisites for using Active Directory Kerberos on Windows:

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

## Write-back

The Simba Spark ODBC Driver supports translation for the following syntax when connecting to a Spark Thrift Server instance that is running Spark 1.3 or later:

- INSERT
- CREATE
- DROP

Spark does not support UPDATE or DELETE syntax.

If the statement contains non-standard SQL-92 syntax, then the driver is unable to translate the statement to SQL and instead falls back to using HiveQL.

## **Dynamic Service Discovery using Datastax AOSS**

The Simba Spark ODBC Driver can be configured to discover services via the DataStax AOSS service.

For information about configuring this feature in the Windows driver, see Creating a Data Source Name on Windows on page 9 or Configuring a DSN-less Connection on Windows on page 12. For information about configuring this feature when using the driver on a non-Windows platform, see Configuring Service Discovery Mode on a Non-Windows Machine on page 50.

## Security and Authentication

To protect data from unauthorized access, some Spark data stores require connections to be authenticated with user credentials or the SSL protocol. The Simba Spark ODBC Driver provides full support for these authentication protocols.

### Note:

In this documentation, "SSL" refers to both TLS (Transport Layer Security) and SSL (Secure Sockets Layer). The driver supports TLS 1.0, 1.1, and 1.2. The SSL version used for the connection is the highest version that is supported by both the driver and the server.

The driver provides mechanisms that enable you to authenticate your connection using the Kerberos protocol, your Spark user name only, or your Spark user name and password. You can also authenticate a connection to an HDInsight distribution on Windows Azure. You must use the authentication mechanism that matches the security requirements of the Spark server. For information about determining the appropriate authentication mechanism to use based on the Spark server configuration, see Authentication Mechanisms on page 61. For detailed driver configuration instructions, see Configuring Authentication on Windows on page 14 or Configuring Authentication on a Non-Windows Machine on page 51.

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

- No identity verification
- One-way authentication
- Two-way authentication

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 SSL Verification on Windows on page 21 or Configuring SSL Verification on a Non-Windows Machine on page 54.

# **Driver Configuration Options**

Driver Configuration Options lists the configuration options available in the Simba Spark 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 machine, the fields and buttons are available in the Simba Spark ODBC Driver Configuration tool and the following dialog boxes:

- Simba Spark ODBC Driver DSN Setup
- Advanced Options
- Server Side Properties
- SSL Options
- HTTP Properties

When using a connection string or configuring a connection from a Linux/macOS/AIX/Solaris machine, use the key names provided.

#### Note:

If you are using the driver on a non-Windows machine, you can set driver configuration properties in a connection string, in a DSN (in the odbc.ini file), or as a driver-wide setting (in the simba.sparkodbc.ini file). Settings in the connection string take precedence over settings in the DSN, and settings in the DSN take precedence over driver-wide settings.

# **Configuration Options Appearing in the User Interface**

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

- Allow Common Name Host Name Mismatch on page 75
- Allow Common Name Host Name Mismatch (for AOSS) on page 75
- Allow Self-Signed Server

- Max File Size on page 89
- Max Number Files on page 89
- Mechanism on page 90
- Minimum TLS Version on page 90
- Minimum TLS Version (for AOSS)

Certificate on page 76

- Allow Self-Signed Server Certificate (for AOSS) on page 76
- Apply Properties with Queries on page 77
- Async Exec Poll Interval on page
  78
- Binary Column Length on page 78
- Check Certificate Revocation on page 78
- Check Certificate Revocation (for AOSS) on page 79
- Client Certificate File on page 80
- Client Private Key File on page 80
- Client Private Key Password on page 81
- Convert Key Name to Lower Case
   on page 81
- Database on page 81
- Decimal Column Scale on page
   82
- Default String Column Length on page 82
- Delegate Kerberos Credentials on page 82
- Delegation UID on page 83
- Driver Config Take Precedence on page 83
- Enable Auto Reconnect on page 83
- Enable SSL on page 84
- Enable SSL (for AOSS) on page 84
- Fast SQLPrepare on page 85
- Get Tables With Query on page 85
- Host(s) on page 86
- Host FQDN on page 86
- HTTP Path on page 87

on page 91

- Password on page 92
- Password (for AOSS) on page 92
- Port on page 93
- Realm on page 93
- Rows Fetched Per Block on page
   93
- Save Password (Encrypted) on page 94
- Service Discovery Mode on page
   94
- Service Name on page 95
- Show System Table on page 95
- Socket Timeout on page 95
- Spark Server Type on page 96
- Status Endpoint Authentication on page 96
- Thrift Transport on page 97
- Trusted Certificates on page 97
- Trusted Certificates (for AOSS) on page 98
- Two Way SSL on page 99
- Unicode SQL Character Types on page 99
- Use Async Exec on page 100
- Use Native Query on page 100
- Use Only SSPI on page 100
- Use System Trust Store on page
  101
- Use System Trust Store (for AOSS) on page 101
- User Name on page 102
- User Name (for AOSS) on page
  103

- Invalid Session Auto Recover on page 87
- Log Level on page 88
- Log Path on page 89

## Allow Common Name Host Name Mismatch

Key Name	Default Value	Required
AllowHostNameCNMismatch	Clear (0)	No

### Description

This option specifies whether a CA-issued SSL certificate name must match the host name of the Spark server.

- Enabled (1): The driver allows a CA-issued SSL certificate name to not match the host name of the Spark server.
- Disabled (0): The CA-issued SSL certificate name must match the host name of the Spark server.

#### Note:

This setting is applicable only when SSL is enabled.

## Allow Common Name Host Name Mismatch (for AOSS)

Key Name	Default Value	Required
AOSS_ AllowHostNameCNMismatch	The value being used for the non-AOSS version of this property, which defaults to Clear (0).	No

### Description

This property specifies whether the driver requires the name of the CA-issued SSL certificate to match the host name of the AOSS endpoint.

- Enabled (1): The driver allows the names to mismatch.
- Disabled (0): The driver requires the names to match.

#### Note:

This property is applicable only when AOSS dynamic service discovery is enabled, and SSL is enabled for AOSS connections.

## Allow Self-Signed Server Certificate

Key Name	Default Value	Required
AllowSelfSigned ServerCert	Clear (0)	No

## Description

This option specifies whether the driver allows a connection to a Spark server that uses a self-signed certificate, even if this certificate is not in the list of trusted certificates. This list is contained in the Trusted Certificates file, or in the system Trust Store if the system Trust Store is used instead of a file.

- Enabled (1): The driver authenticates the Spark server even if the server is using a self-signed certificate that has not been added to the list of trusted certificates.
- Disabled (0): The driver does not allow self-signed certificates from the server unless they have already been added to the list of trusted certificates.

#### Note:

This setting is applicable only when SSL is enabled.

## Allow Self-Signed Server Certificate (for AOSS)

Key Name	Default Value	Required
AOSS_ AllowSelfSignedServerCert	The value being used for the non- AOSS version of this property, which defaults to Clear (0).	No

This property specifies whether the driver allows the AOSS endpoint to use self-signed SSL certificates.

• Enabled (1): The driver allows self-signed certificates.

#### Important:

When this property is enabled, SSL verification is disabled. The driver does not verify the endpoint's certificate against the trust store, and does not verify if the endpoint's host name matches the common name or subject alternative names in the endpoint's certificate.

• Disabled (0): The driver does not allow self-signed certificates.

#### Note:

This property is applicable only when AOSS dynamic service discovery is enabled, and SSL is enabled for AOSS connections.

## Apply Properties with Queries

Key Name	Default Value	Required
ApplySSPWithQueries	Selected (1)	No

### Description

This option specifies how the driver applies server-side properties.

- Enabled (1): The driver applies each server-side property by executing a set SSPKey=SSPValue query when opening a session to the Spark server.
- Disabled (0): The driver uses a more efficient method for applying server-side properties that does not involve additional network round-tripping. However, some Spark Thrift Server builds are not compatible with the more efficient method.

#### Note:

When connecting to a Shark Server instance, this option is always enabled.

## Async Exec Poll Interval

Key Name	Default Value	Required
AsyncExecPollInterval	100	No

### Description

The time in milliseconds between each poll for the query execution status.

"Asynchronous execution" refers to the fact that the RPC call used to execute a query against Spark is asynchronous. It does not mean that ODBC asynchronous operations are supported.

#### Note:

This option is applicable only to HDInsight clusters.

## Binary Column Length

Key Name	Default Value	Required
BinaryColumnLength	32767	No

### Description

The maximum data length for BINARY columns.

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

## **Check Certificate Revocation**

Key Name	Default Value	Required
CheckCertRevocation	Selected (1)	No

#### Description

This option specifies whether the driver checks to see if a certificate has been revoked while retrieving a certificate chain from the Windows Trust Store.

This option is only applicable if you are using a CA certificate from the Windows Trust Store (see Use System Trust Store on page 101).

- Enabled (1): The driver checks for certificate revocation while retrieving a certificate chain from the Windows Trust Store.
- Disabled (0): The driver does not check for certificate revocation while retrieving a certificate chain from the Windows Trust Store.

#### Note:

This option is only available on Windows.

## Check Certificate Revocation (for AOSS)

Key Name	Default Value	Required
AOSS_ CheckCertRevocation	The value being used for the non-AOSS version of this property, which defaults to Selected (1).	No

#### Description

This property specifies whether, when connecting to an AOSS endpoint with SSL enabled, the driver checks to see if a certificate has been revoked while retrieving a certificate chain from the Windows Trust Store.

- Enabled (1): The driver checks for certificate revocation.
- Disabled (0): The driver does not check for certificate revocation.

#### Note:

This property is applicable only when all of the following is true:

- AOSS dynamic service discovery is enabled.
- SSL is enabled for AOSS connections.
- The driver is configured to use a CA certificate from the Windows Trust Store (see Use System Trust Store (for AOSS) on page 101).

## **Client Certificate File**

Key Name	Default Value	Required
ClientCert	None	No

#### Description

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

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

## Client Private Key File

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

## Description

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

If the private key file is protected with a password, then provide the password using the driver configuration option Client Private Key Password on page 81.

## Note:

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

## Client Private Key Password

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

### Description

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

## Convert Key Name to Lower Case

Key Name	Default Value	Required
LCaseSspKeyName	Selected (1)	No

### Description

This option specifies whether the driver converts server-side property key names to all lower-case characters.

- Enabled (1): The driver converts server-side property key names to all lowercase characters.
- Disabled (0): The driver does not modify the server-side property key names.

## Database

Key Name	Default Value	Required
Schema	default	No

#### Description

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

schema in the query.

#### Note:

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

## Decimal Column Scale

Key Name	Default Value	Required
DecimalColumnScale	10	No

#### Description

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

## Default String Column Length

Key Name	Default Value	Required
Default StringColumnLength	255	No

#### Description

The maximum number of characters that can be contained in STRING columns.

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

## **Delegate Kerberos Credentials**

Key Name	Default Value	Required
DelegateKrbCreds	0	No

#### Description

This option specifies whether your Kerberos credentials are forwarded to the server and used for authentication.

#### Note:

This option is only applicable when Authentication Mechanism is set to Kerberos (AuthMech=1).

## **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 Spark to the specified user, rather than to the authenticated user for the connection.

#### Note:

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

## Driver Config Take Precedence

Key Name	Default Value	Required
DriverConfigTakePrecedence	Clear(0)	No

#### Description

This option specifies whether driver-wide configuration settings take precedence over connection and DSN settings.

- Enabled (1): Driver-wide configurations take precedence over connection and DSN settings.
- Disabled (0): Connection and DSN settings take precedence instead.

## Enable Auto Reconnect

Key Name	Default Value	Required
AutoReconnect	Selected (1)	Yes

This option specifies whether the driver attempts to automatically reconnect to the server when a communication link error occurs.

- Enabled (1): The driver attempts to reconnect.
- Disabled (0): The driver does not attempt to reconnect.

### 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 Spark server.

- Enabled (1): The client communicates with the Spark 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.

#### Note:

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

## Enable SSL (for AOSS)

Key Name	Default Value	Required
AOSS_SSL	The value being used for the non-AOSS version of this property, which defaults to Clear (0).	No

This property specifies whether the driver communicates with AOSS endpoints through SSL-enabled sockets.

- Enabled (1): The driver connects to SSL-enabled sockets.
- Disabled (0): The driver does not connect to SSL-enabled sockets.

#### Note:

This property is applicable only when AOSS dynamic service discovery is enabled.

## Fast SQLPrepare

Key Name	Default Value	Required
FastSQLPrepare	Clear (0)	No

### Description

This option specifies whether the driver defers query execution to SQLExecute.

- Enabled (1): The driver defers query execution to SQLExecute.
- Disabled (0): The driver does not defer query execution to SQLExecute.

#### Note:

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

## Get Tables With Query

Key Name	Default Value	Required
GetTablesWithQuery	Selected (1)	No

#### Description

This option specifies whether the driver uses the SHOW TABLES query or the GetTables Thrift API call to retrieve table names from the database.

- Enabled (1): The driver uses the SHOW TABLES query to retrieve table names.
- Disabled (0): The driver uses the GetTables Thrift API call to retrieve table names.

#### Note:

This option is applicable only when connecting to a Spark Thrift Server instance.

### Host(s)

Key Name	Default Value	Required
Host	None	Yes

#### Description

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

If Service Discovery Mode is enabled, a comma-separated list of AOSS endpoints in the following format, where [AOSS\_Endpoint] is the IP address or host name of the AOSS endpoint, and [AOSS\_Port] is the number of the TCP port that the AOSS endpoint uses to listen for client connections:

```
[AOSS_Endpoint1]: [AOSS_Port1], [AOSS_Endpoint2]: [AOSS_Port2]
```

## Host FQDN

Key Name	Default Value	Required
KrbHostFQDN	_HOST	No

### Description

The fully qualified domain name of the Spark Thrift Server host.

When the value of Host FQDN is \_HOST, the driver uses the Spark server host name as the fully qualified domain name for Kerberos authentication.

## HTTP Path

Key Name	Default Value	Required
HTTPPath	/spark if using Windows Azure HDInsight Service (6). / if using non-Windows Azure HDInsight Service with Thrift Transport set to HTTP (2).	No

## Description

The partial URL corresponding to the Spark server.

The driver forms the HTTP address to connect to by appending the HTTP Path value to the host and port specified in the DSN or connection string. For example, to connect to the HTTP address

http://localhost:10002/gateway/sandbox/spark/version, you would set HTTP Path to /gateway/sandbox/spark/version.

## Invalid Session Auto Recover

Key Name	Default Value	Required
InvalidSessionAutoRecover	Selected (1)	No

## Description

This option specifies whether the driver automatically opens a new session when the existing session is no longer valid.

- Enabled (1): The driver automatically opens a new session when the existing session is no longer valid.
- Disabled (0): The driver does not automatically open new sessions.

#### Note:

This option is applicable only when connecting to Spark Thrift Server.

## Log Level

Key Name	Default Value	Required
LogLevel	<b>OFF (</b> 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.

ł	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 Spark 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.sparkodbc.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 the following log files at the location you specify in the Log Path (LogPath) property:

- A simbasparkodbcdriver.log file that logs driver activity that is not specific to a connection.
- A simbasparkodbcdriver\_connection\_[Number].log file 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.

## Log Path

Key Name	Default Value	Required
LogPath	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.sparkodbc.ini file.

### Max File Size

Key Name	Default Value	Required
LogFileSize	20	No

### Description

The maximum size of each log file in megabytes (MB). After the maximum file size is reached, the driver creates a new file and continues logging.

#### 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.sparkodbc.ini file.

### Max Number Files

Key Name	Default Value	Required
LogFileCount	50	No

The maximum number of log files to keep. After the maximum number of log files is reached, each time an additional file is created, the driver deletes the oldest log file.

#### 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.sparkodbc.ini file.

### Mechanism

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

## Description

The authentication mechanism to use.

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

- No Authentication (0)
- Kerberos (1)
- User Name (2)
- User Name And Password (3)
- Windows Azure HDInsight Emulator (5)
- Windows Azure HDInsight Service (6)

## Minimum TLS Version

Key Name	Default Value	Required
Min_TLS	TLS 1.2 (1.2)	No

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

## Minimum TLS Version (for AOSS)

Key Name	Default Value	Required
AOSS_Min_TLS	The value being used for the non-AOSS version of this property, which defaults to TLS 1.2 (1.2).	No

#### Description

The minimum version of TLS/SSL that the driver allows the AOSS endpoint 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 (1.0): The connection must use at least TLS 1.0.
- TLS 1.1 (1.1): The connection must use at least TLS 1.1.
- TLS 1.2 (1.2): The connection must use at least TLS 1.2.

#### Note:

This property is applicable only when AOSS dynamic service discovery is enabled, and SSL is enabled for AOSS connections.

## Password

Key Name	Default Value	Required
PWD	None	Yes, if the authentication mechanism is User Name And Password (3), Windows Azure HDInsight Emulator (5), or Windows Azure HDInsight Service (6).

## Description

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

### Password (for AOSS)

Key Name	Default Value	Required
AOSS_PWD	The value being used for the non-AOSS version of this property, which does not have a default value.	No

### Description

The password corresponding to the user name that you provided using the property User Name (for AOSS) on page 103.

#### Note:

This property is applicable only when AOSS dynamic service discovery is enabled.

## Port

Key Name	Default Value	Required
Port	<ul> <li>non-HDInsight clusters: 10000</li> </ul>	Yes
	Windows Azure     HDInsight     Emulator: 10001	
	Windows Azure     HDInsight Service:     443	

### Description

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

### Realm

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

### Description

The realm of the Spark Thrift Server host.

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

## **Rows Fetched Per Block**

Key Name	Default Value	Required
RowsFetchedPerBlock	10000	No

### Description

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

Valid values for this setting include any positive 32-bit integer. However, testing has shown that performance gains are marginal beyond the default value of 10000 rows.

## Save Password (Encrypted)

Key Name	Default Value	Required
N/A	Selected	No

#### Description

This option specifies whether the password is saved in the registry.

- Enabled: The password is saved in the registry.
- Disabled: The password is not saved in the registry.

This option is available only in the Windows driver. It appears in the Simba Spark ODBC Driver DSN Setup dialog box and the SSL Options dialog box.

#### Important:

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

## Service Discovery Mode

Key Name	Default Value	Required
ServiceDiscoveryMode	No Service Discovery	No

### Description

This option specifies whether the driver uses the DataStax AOSS service.

- AOSS: The driver discovers Spark services via the AOSS service.
- No Service Discovery: The driver connects to Spark without using a discovery service.

## Service Name

Key Name	Default Value	Required
KrbServiceName	spark	No

#### Description

The Kerberos service principal name of the Spark server.

## Show System Table

Key Name	Default Value	Required
ShowSystemTable	Clear (0)	No

### Description

This option specifies whether the driver returns the spark\_system table for catalog function calls such as SQLTables and SQLColumns.

- Enabled (1): The driver returns the spark\_system table for catalog function calls such as SQLTables and SQLColumns.
- Disabled (0): The driver does not return the spark\_system table for catalog function calls.

## Socket Timeout

Key Name	Default Value	Required
SocketTimeout	60	No

#### Description

The number of seconds that an operation can remain idle before it is closed.

#### Note:

This option is applicable only when asynchronous query execution is being used against Spark Thrift Server instances.

## Spark Server Type

Key Name	Default Value	Required
SparkServerType	Spark Thrift Server (3)	No

#### Description

This option specifies the type of Spark server.

#### Note:

The Shark Server 2 option is provided only for backwards compatibility with previous applications. If the driver will connect to Shark 0.9, or Spark 1.1 or later, then set Spark Thrift Server (3).

- Shark Server (1): The driver connects to a Shark Server instance.
- Shark Server 2 (2): The driver connects to a Shark Server 2 instance.
- Spark Thrift Server (3): The driver connects to a Spark Thrift Server instance.

## Status Endpoint Authentication

Key Name	Default Value	Required
AOSS_AuthMech	0	No

### Description

This property specifies whether the driver authenticates AOSS connections.

- No Authentication (0): The driver does not authenticate connections to AOSS endpoints.
- User Name and Password (3): The driver authenticates connections to AOSS endpoints using a user name and password.

#### Note:

This property is applicable only when AOSS dynamic service discovery is enabled.

## Thrift Transport

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

## Description

The transport protocol to use in the Thrift layer.

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

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

#### Note:

For information about how to determine which Thrift transport protocols your Spark server supports, see Authentication Mechanisms on page 61.

## **Trusted Certificates**

Key Name	Default Value	Required
TrustedCerts	The cacerts.pem file in the \lib 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

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

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

## Note:

This setting is applicable only when SSL is enabled.

## Trusted Certificates (for AOSS)

Key Name	Default Value	Required
AOSS_TrustedCerts	The value being used for the non-AOSS version of this property, which defaults to the path to the cacerts.pem file in the \lib subfolder within the driver's installation directory.	No

#### Description

The full path of the pem file containing trusted CA certificates, for verifying the AOSS endpoint when using SSL.

### Note:

This property is applicable only when AOSS dynamic service discovery is enabled, and SSL is enabled for AOSS connections.

#### ! Important:

If you are connecting from a Windows machine and the Use System Trust Store option is enabled, the driver uses the certificates from the Windows trust store instead of your specified .pem file. For more information, see Use System Trust Store on page 101.

## Two Way SSL

Key Name	Default Value	Required
TwoWaySSL	Clear (0)	No

### Description

This option specifies whether two-way SSL is enabled.

- Enabled (1): The client and the Spark server verify each other using SSL. See also the driver configuration options Client Certificate File on page 80, Client Private Key File on page 80, and Client Private Key Password on page 81.
- Disabled (0): The server does not verify the client. Depending on whether oneway SSL is enabled, the client might verify the server. For more information, see Enable SSL on page 84.

#### Note:

This option is applicable only when connecting to a Spark server that supports SSL. You must enable SSL before Two Way SSL can be configured. For more information, see Enable SSL on page 84.

## Unicode SQL Character Types

Key Name	Default Value	Required
UseUnicodeSqlCharacterTypes	Clear (0)	No

## Description

This option specifies the SQL types to be returned for string data types.

- Enabled (1): The driver returns SQL\_WVARCHAR for STRING and VARCHAR columns, and returns SQL\_WCHAR for CHAR columns.
- Disabled (0): The driver returns SQL\_VARCHAR for STRING and VARCHAR columns, and returns SQL\_CHAR for CHAR columns.

## Use Async Exec

Key Name	Default Value	Required
EnableAsyncExec	Clear (0)	No

#### Description

This option specifies whether to execute queries synchronously or asynchronously.

- Enabled (1): The driver uses an asynchronous version of the API call against Spark for executing a query.
- Disabled (0): The driver executes queries synchronously.

## Use Native Query

Key Name	Default Value	Required
UseNativeQuery	Clear (0)	No

#### Description

This option specifies whether the driver uses native HiveQL queries, or converts the queries emitted by an application into an equivalent form in HiveQL. If the application is Spark-aware and already emits HiveQL, then enable this option to avoid the extra overhead of query transformation.

- Enabled (1): The driver does not transform the queries emitted by an application, and executes HiveQL queries directly.
- Disabled (0): The driver transforms the queries emitted by an application and converts them into an equivalent form in HiveQL.

#### Important:

When this option is enabled, the driver cannot execute parameterized queries.

## Use Only SSPI

Key Name	Default Value	Required
UseOnlySSPI	Clear (0)	No

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.

### 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 Trusted Certificates on page 97.

### Note:

This option is only available on Windows.

## Use System Trust Store (for AOSS)

Key Name	Default Value	Required
AOSS_ UseSystemTrustStore	The value being used for the non-AOSS version of this property, which defaults to Clear (0).	No

This property specifies whether, when connecting to an AOSS endpoint with SSL enabled, the driver uses 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 Trusted Certificates (for AOSS) on page 98.

#### Note:

This property is available only on Windows, and applicable only when AOSS dynamic service discovery is enabled.

#### User Name

Key Name	Default Value	Required
UID	For User Name (2) authentication only, the default value is anonymous	Yes, if the authentication mechanism is User Name And Password (3), Windows Azure HDInsight Emulator (5), or Windows Azure HDInsight Service (6). No, if the authentication mechanism is User Name (2).

### Description

The user name that you use to access Spark Thrift Server.

## User Name (for AOSS)

Key Name	Default Value	Required
AOSS_UID	The value being used for the non-AOSS version of the User Name property (UID), which does not have a default value for User Name And Password authentication.	No

## Description

The user name that you use to access the AOSS endpoints.

#### Note:

This property is applicable only when AOSS dynamic service discovery is enabled.

# **Configuration Options Having Only Key Names**

The following configuration options do not appear in the Windows user interface for the Simba Spark ODBC Driver. They are accessible only when you use a connection string or configure a connection from a Linux/macOS/AIX/Solaris machine:

- ADUserNameCase on page 104
- Driver on page 104
- ForceSynchronousExec on page 105
- http.header. on page 105
- ProxyHost on page 106
- ProxyPort on page 106
- ProxyPWD on page 106
- ProxyUID on page 107
- SSP\_ on page 107
- UseProxy on page 107

## ADUserNameCase

Key Name	Default Value	Required
ADUserNameCase	Unchanged	No

#### Description

This option controls whether the driver changes the user name part of an AD Kerberos UPN to all upper-case or all lower-case. The following values are supported:

- Upper: Change the user name to all upper-case.
- Lower: Change the user name to all lower-case.
- Unchanged: Do not modify the user name.

#### Note:

This option is applicable only when using Active Directory Kerberos from a Windows client machine to authenticate.

## Driver

Key Name	Default Value	Required
Driver	Simba Spark 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 Spark 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.

## ForceSynchronousExec

Key Name	Default Value	Required
ForceSynchronousExec	0	No

### Description

When this option is enabled (1), the driver is forced to execute queries synchronously when connected to an HDInsight cluster.

When this option is disabled (0), the driver is able to execute queries asynchronously when connected to an HDInsight cluster.

#### Note:

This option is applicable only to HDInsight clusters.

### http.header.

Key Name	Default Value	Required
http.header	None	No

### Description

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

http.header.[HeaderKey]=[HeaderValue]

For example:

http.header.AUTHENTICATED\_USER=john

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

The example above would create the following custom HTTP header:

AUTHENTICATED\_USER: john

#### Note:

The http.header. prefix is case-sensitive. This option is applicable only when you are using HTTP as the Thrift transport protocol. For more information, see Thrift Transport on page 97.

## ProxyHost

Key Name	Default Value	Required
ProxyHost	None	Yes, if connecting through a proxy server.

### Description

The host name or IP address of a proxy server that you want to connect through.

## ProxyPort

Key Name	Default Value	Required
ProxyPort	None	Yes, if connecting through a proxy server.

### Description

The number of the port that the proxy server uses to listen for client connections.

## **ProxyPWD**

Key Name	Default Value	Required
ProxyPWD	None	Yes, if connecting to a proxy server that requires authentication.

### Description

The password that you use to access the proxy server.

## ProxyUID

Key Name	Default Value	Required
ProxyUID	None	Yes, if connecting to a proxy server that requires authentication.

### Description

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

### SSP\_

Key Name	Default Value	Required
SSP_	None	No

### Description

Set a server-side property by using the following syntax, where [SSPKey] is the name of the server-side property and [SSPValue] is the value for that property:

SSP [SSPKey]=[SSPValue]

After the driver applies the server-side property, the SSP\_prefix is removed from the DSN entry, leaving an entry of [SSPKey]=[SSPValue].

#### Note:

- The SSP\_ prefix must be upper case.
- When setting a server-side property in a connection string, it is recommended that you enclose the value in braces ({ }) to make sure that special characters can be properly escaped.

## UseProxy

Key Name	Default Value	Required
UseProxy	Clear (0)	No

This option specifies whether the driver uses a proxy server to connect to the data store.

- Enabled (1): The driver connects to a proxy server based on the information provided in the ProxyHost, ProxyPort, ProxyUID, and ProxyPWD keys.
- Disabled (0): The driver connects directly to the Spark server.

# Third-Party Trademarks

IBM and AIX are trademarks or registered trademarks of IBM Corporation or its subsidiaries in Canada, United States, and/or other countries.

Kerberos is a trademark of the Massachusetts Institute of Technology (MIT).

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 Azure, 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.

Solaris is a registered trademark of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

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

Apache Spark, Apache, and Spark 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 NONINFRINGEMENT 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.

# 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.

# **MIT Kerberos License**

Copyright © 1985-2015 by the Massachusetts Institute of Technology.

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.

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 HOLDER 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.

Downloading of this software may constitute an export of cryptographic software from the United States of America that is subject to the United States Export Administration Regulations (EAR), 15 CFR 730-774. Additional laws or regulations may apply. It is the responsibility of the person or entity contemplating export to comply with all applicable export laws and regulations, including obtaining any required license from the U.S. government.

The U.S. government prohibits export of encryption source code to certain countries and individuals, including, but not limited to, the countries of Cuba, Iran, North Korea, Sudan, Syria, and residents and nationals of those countries.

Documentation components of this software distribution are licensed under a Creative Commons Attribution-ShareAlike 3.0 Unported License. (http://creativecommons.org/licenses/by-sa/3.0/)

Individual source code files are copyright MIT, Cygnus Support, Novell, OpenVision Technologies, Oracle, Red Hat, Sun Microsystems, FundsXpress, and others.

Project Athena, Athena, Athena MUSE, Discuss, Hesiod, Kerberos, Moira, and Zephyr are trademarks of the Massachusetts Institute of Technology (MIT). No commercial use of these trademarks may be made without prior written permission of MIT.

"Commercial use" means use of a name in a product or other for-profit manner. It does NOT prevent a commercial firm from referring to the MIT trademarks in order to convey information (although in doing so, recognition of their trademark status should be given).

The following copyright and permission notice applies to the OpenVision Kerberos Administration system located in kadmin/create, kadmin/dbutil, kadmin/passwd, kadmin/server, lib/kadm5, and portions of lib/rpc:

Copyright, OpenVision Technologies, Inc., 1993-1996, All Rights Reserved

WARNING: Retrieving the OpenVision Kerberos Administration system source code, as described below, indicates your acceptance of the following terms. If you do not agree to the following terms, do not retrieve the OpenVision Kerberos administration system.

You may freely use and distribute the Source Code and Object Code compiled from it, with or without modification, but this Source Code is provided to you "AS IS" EXCLUSIVE OF ANY WARRANTY, INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, OR ANY OTHER WARRANTY, WHETHER EXPRESS OR IMPLIED. IN NO EVENT WILL OPENVISION HAVE ANY LIABILITY FOR ANY LOST PROFITS, LOSS OF DATA OR COSTS OF PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, INCLUDING, WITHOUT LIMITATION, THOSE RESULTING FROM THE USE OF THE SOURCE CODE, OR THE FAILURE OF THE SOURCE CODE TO PERFORM, OR FOR ANY OTHER REASON.

OpenVision retains all copyrights in the donated Source Code. OpenVision also retains copyright to derivative works of the Source Code, whether created by OpenVision or by a third party. The OpenVision copyright notice must be preserved if derivative works are made based on the donated Source Code.

OpenVision Technologies, Inc. has donated this Kerberos Administration system to MIT for inclusion in the standard Kerberos 5 distribution. This donation underscores our commitment to continuing Kerberos technology development and our gratitude for the valuable work which has been performed by MIT and the Kerberos community.

Portions contributed by Matt Crawford crawdad@fnal.gov were work performed at Fermi National Accelerator Laboratory, which is operated by Universities Research Association, Inc., under contract DE-AC02-76CHO3000 with the U.S. Department of Energy.

Portions of src/lib/crypto have the following copyright:

Copyright © 1998 by the FundsXpress, INC.

All rights reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of FundsXpress. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. FundsXpress makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED

WARRANTIES OF MERCHANTIBILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The implementation of the AES encryption algorithm in src/lib/crypto/builtin/aes has the following copyright:

Copyright © 2001, Dr Brian Gladman brg@gladman.uk.net, Worcester, UK.

All rights reserved.

LICENSE TERMS

The free distribution and use of this software in both source and binary form is allowed (with or without changes) provided that:

- 1. distributions of this source code include the above copyright notice, this list of conditions and the following disclaimer;
- 2. distributions in binary form include the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other associated materials;
- 3. the copyright holder's name is not used to endorse products built using this software without specific written permission.

#### DISCLAIMER

This software is provided 'as is' with no explicit or implied warranties in respect of any properties, including, but not limited to, correctness and fitness for purpose.

Portions contributed by Red Hat, including the pre-authentication plug-in framework and the NSS crypto implementation, contain the following copyright:

Copyright © 2006 Red Hat, Inc.

Portions copyright © 2006 Massachusetts Institute of Technology

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 Red Hat, 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.

The bundled verto source code is subject to the following license:

Copyright 2011 Red Hat, Inc.

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.

The MS-KKDCP client implementation has the following copyright:

Copyright 2013,2014 Red Hat, Inc.

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.

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.

The implementations of GSSAPI mechglue in GSSAPI-SPNEGO in src/lib/gssapi, including the following files:

```
lib/gssapi/generic/gssapi err generic.et
lib/gssapi/mechglue/g accept sec context.c
lib/gssapi/mechglue/g acquire cred.c
lib/gssapi/mechglue/g canon name.c
lib/gssapi/mechglue/g compare name.c
lib/gssapi/mechglue/g context time.c
lib/gssapi/mechglue/g_delete sec context.c
lib/gssapi/mechglue/g dsp name.c
lib/gssapi/mechglue/g dsp status.c
lib/gssapi/mechglue/g dup name.c
lib/gssapi/mechglue/g exp sec context.c
lib/gssapi/mechglue/g export name.c
lib/gssapi/mechglue/g glue.c
lib/gssapi/mechglue/g imp name.c
lib/gssapi/mechglue/g imp sec context.c
lib/gssapi/mechglue/g init sec context.c
lib/gssapi/mechglue/g initialize.c
lib/gssapi/mechglue/g inquire context.c
lib/gssapi/mechglue/g inquire cred.c
lib/gssapi/mechglue/g inquire names.c
lib/gssapi/mechglue/g process context.c
lib/gssapi/mechglue/g rel buffer.c
```

```
lib/gssapi/mechglue/g_rel_cred.c
lib/gssapi/mechglue/g_rel_oid_set.c
lib/gssapi/mechglue/g_seal.c
lib/gssapi/mechglue/g_sign.c
lib/gssapi/mechglue/g_store_cred.c
lib/gssapi/mechglue/g_unseal.c
lib/gssapi/mechglue/g_userok.c
lib/gssapi/mechglue/g_utils.c
lib/gssapi/mechglue/g_verify.c
lib/gssapi/mechglue/gssd_pname_to_uid.c
lib/gssapi/mechglue/mglueP.h
lib/gssapi/mechglue/oid_ops.c
lib/gssapi/spnego/gssapiP_spnego.h
lib/gssapi/spnego/spnego_mech.c
```

and the initial implementation of incremental propagation, including the following new or changed files:

```
include/iprop_hdr.h
kadmin/server/ipropd_svc.c
lib/kdb/iprop.x
lib/kdb/kdb_convert.c
lib/kdb/kdb_log.c
lib/kdb/kdb_log.h
lib/krb5/error_tables/kdb5_err.et
slave/kpropd_rpc.c
slave/kproplog.c
```

are subject to the following license:

Copyright © 2004 Sun Microsystems, Inc.

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.

Kerberos V5 includes documentation and software developed at the University of California at Berkeley, which includes this copyright notice:

Copyright © 1983 Regents of the University of California.

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. Neither the name of the University 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 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.

Portions contributed by Novell, Inc., including the LDAP database backend, are subject to the following license:

Copyright © 2004-2005, Novell, 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.
- The copyright holder's name is not 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.

Portions funded by Sandia National Laboratory and developed by the University of Michigan's Center for Information Technology Integration, including the PKINIT implementation, are subject to the following license:

COPYRIGHT © 2006-2007

THE REGENTS OF THE UNIVERSITY OF MICHIGAN

ALL RIGHTS RESERVED

Permission is granted to use, copy, create derivative works and redistribute this software and such derivative works for any purpose, so long as the name of The University of Michigan is not used in any advertising or publicity pertaining to the use of distribution of this software without specific, written prior authorization. If the above copyright notice or any other identification of the University of Michigan is included in any copy of any portion of this software, then the disclaimer below must also be included.

THIS SOFTWARE IS PROVIDED AS IS, WITHOUT REPRESENTATION FROM THE UNIVERSITY OF MICHIGAN AS TO ITS FITNESS FOR ANY PURPOSE, AND WITHOUT WARRANTY BY THE UNIVERSITY OF MICHIGAN OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE REGENTS OF THE UNIVERSITY OF MICHIGAN SHALL NOT BE LIABLE FOR ANY DAMAGES, INCLUDING SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, WITH RESPECT TO ANY CLAIM ARISING OUT OF OR IN CONNECTION WITH THE USE OF THE SOFTWARE, EVEN IF IT HAS BEEN OR IS HEREAFTER ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The pkcs11.h file included in the PKINIT code has the following license:

Copyright 2006 g10 Code GmbH

Copyright 2006 Andreas Jellinghaus

This file is free software; as a special exception the author gives unlimited permission to copy and/or distribute it, with or without modifications, as long as this notice is preserved.

This file is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY, to the extent permitted by law; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

Portions contributed by Apple Inc. are subject to the following license:

Copyright 2004-2008 Apple Inc. All Rights Reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of Apple Inc. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. Apple Inc. makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTIBILITY AND FITNESS FOR A PARTICULAR PURPOSE. The implementations of UTF-8 string handling in src/util/support and src/lib/krb5/unicode are subject to the following copyright and permission notice:

The OpenLDAP Public License

Version 2.8, 17 August 2003

Redistribution and use of this software and associated documentation ("Software"), with or without modification, are permitted provided that the following conditions are met:

- 1. Redistributions in source form must retain copyright statements and notices,
- 2. Redistributions in binary form must reproduce applicable copyright statements and notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution, and
- 3. Redistributions must contain a verbatim copy of this document.

The OpenLDAP Foundation may revise this license from time to time. Each revision is distinguished by a version number. You may use this Software under terms of this license revision or under the terms of any subsequent revision of the license.

THIS SOFTWARE IS PROVIDED BY THE OPENLDAP FOUNDATION AND ITS CONTRIBUTORS "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 OPENLDAP FOUNDATION, ITS CONTRIBUTORS, OR THE AUTHOR(S) OR OWNER(S) OF THE SOFTWARE 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 names of the authors and copyright holders must not be used in advertising or otherwise to promote the sale, use or other dealing in this Software without specific, written prior permission. Title to copyright in this Software shall at all times remain with copyright holders.

OpenLDAP is a registered trademark of the OpenLDAP Foundation.

Copyright 1999-2003 The OpenLDAP Foundation, Redwood City, California, USA. All Rights Reserved. Permission to copy and distribute verbatim copies of this document is granted.

Marked test programs in src/lib/krb5/krb have the following copyright:

Copyright © 2006 Kungliga Tekniska Högskola

(Royal Institute of Technology, Stockholm, Sweden).

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. Neither the name of KTH 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 KTH AND ITS 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 KTH 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.

The KCM Mach RPC definition file used on OS X has the following copyright:

Copyright © 2009 Kungliga Tekniska Högskola

(Royal Institute of Technology, Stockholm, Sweden).

All rights reserved.

Portions Copyright © 2009 Apple 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:

- 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. Neither the name of the Institute 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 INSTITUTE 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 INSTITUTE 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.

Portions of the RPC implementation in src/lib/rpc and src/include/gssrpc have the following copyright and permission notice:

Copyright © 2010, Oracle America, 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:

- 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. Neither the name of the "Oracle America, 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 HOLDER 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.

Copyright © 2006,2007,2009 NTT (Nippon Telegraph and Telephone Corporation). 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 as the first lines of this file unmodified.
- 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.

THIS SOFTWARE IS PROVIDED BY NTT "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 NTT 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.

Copyright 2000 by Carnegie Mellon University

All Rights Reserved

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above

copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of Carnegie Mellon University not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

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.

Copyright © 2002 Naval Research Laboratory (NRL/CCS)

Permission to use, copy, modify and distribute this software and its documentation is hereby granted, provided that both the copyright notice and this permission notice appear in all copies of the software, derivative works or modified versions, and any portions thereof.

NRL ALLOWS FREE USE OF THIS SOFTWARE IN ITS "AS IS" CONDITION AND DISCLAIMS ANY LIABILITY OF ANY KIND FOR ANY DAMAGES WHATSOEVER RESULTING FROM THE USE OF THIS SOFTWARE.

Portions extracted from Internet RFCs have the following copyright notice:

Copyright © The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright © 1991, 1992, 1994 by Cygnus Support.

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Cygnus Support makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

Copyright © 2006 Secure Endpoints Inc.

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.

Portions of the implementation of the Fortuna-like PRNG are subject to the following notice:

Copyright © 2005 Marko Kreen

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.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR 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 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.

Copyright © 1994 by the University of Southern California

EXPORT OF THIS SOFTWARE from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to copy, modify, and distribute this software and its documentation in source and binary forms is hereby granted, provided that any documentation or other materials related to such distribution or use acknowledge that the software was developed by the University of Southern California.

DISCLAIMER OF WARRANTY. THIS SOFTWARE IS PROVIDED "AS IS". The University of Southern California MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED. By way of example, but not limitation, the University of Southern California MAKES NO REPRESENTATIONS OR WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. The University of Southern California shall not be held liable for any liability nor for any direct, indirect, or consequential damages with respect to any claim by the user or distributor of the ksu software.

Copyright © 1995

The President and Fellows of Harvard University

This code is derived from software contributed to Harvard by Jeremy Rassen.

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 acknowledgement:

This product includes software developed by the University of California, Berkeley and its contributors.

4. Neither the name of the University 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 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.

Copyright © 2008 by the Massachusetts Institute of Technology.

Copyright 1995 by Richard P. Basch. All Rights Reserved.

Copyright 1995 by Lehman Brothers, Inc. All Rights Reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of Richard P. Basch, Lehman Brothers and M.I.T. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. Richard P. Basch, Lehman Brothers and M.I.T. make no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

The following notice applies to src/lib/krb5/krb/strptime.c and src/include/k5-queue.h.

Copyright © 1997, 1998 The NetBSD Foundation, Inc.

All rights reserved.

This code was contributed to The NetBSD Foundation by Klaus Klein.

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 acknowledgement:

This product includes software developed by the NetBSD Foundation, Inc. and its contributors.

4. Neither the name of The NetBSD Foundation 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 NETBSD FOUNDATION, INC. 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 FOUNDATION 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 following notice applies to Unicode library files in src/lib/krb5/unicode:

Copyright 1997, 1998, 1999 Computing Research Labs,

New Mexico State University

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 COMPUTING RESEARCH LAB OR NEW MEXICO STATE UNIVERSITY 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.

The following notice applies to src/util/support/strlcpy.c:

Copyright © 1998 Todd C. Miller Todd.Miller@courtesan.com

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" AND THE AUTHOR DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY SPECIAL, DIRECT, 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.

The following notice applies to src/util/profile/argv\_parse.c and src/util/profile/argv parse.h:

Copyright 1999 by Theodore Ts'o.

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" AND THEODORE TS'O (THE AUTHOR) DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY SPECIAL, DIRECT, 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. (Isn't it sick that the U.S. culture of lawsuit-happy lawyers requires this kind of disclaimer?)

The following notice applies to SWIG-generated code in src/util/profile/profile\_tcl.c:

Copyright © 1999-2000, The University of Chicago

This file may be freely redistributed without license or fee provided this copyright message remains intact.

The following notice applies to portiions of src/lib/rpc and src/include/gssrpc:

Copyright © 2000 The Regents of the University of Michigan. All rights reserved.

Copyright © 2000 Dug Song dugsong@UMICH.EDU. All rights reserved, all wrongs reversed.

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. Neither the name of the University 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 "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 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.

Implementations of the MD4 algorithm are subject to the following notice:

Copyright © 1990, RSA Data Security, Inc. All rights reserved.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD4 Message Digest Algorithm" in all material mentioning or referencing this software or this function.

License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD4 Message Digest Algorithm" in all material mentioning or referencing the derived work.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

Implementations of the MD5 algorithm are subject to the following notice:

Copyright © 1990, RSA Data Security, Inc. All rights reserved.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message- Digest Algorithm" in all material mentioning or referencing this software or this function.

License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

The following notice applies to src/lib/crypto/crypto\_tests/t\_
mddriver.c:

Copyright © 1990-2, RSA Data Security, Inc. Created 1990. All rights reserved.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

Portions of src/lib/krb5 are subject to the following notice:

Copyright © 1994 CyberSAFE Corporation.

Copyright 1990,1991,2007,2008 by the Massachusetts Institute of Technology.

All Rights Reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of M.I.T. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. Furthermore if you modify this software you must label your software as modified software and not distribute it in such a fashion that it might be confused with the original M.I.T. software. Neither M.I.T., the Open Computing Security Group, nor CyberSAFE Corporation make any representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

Portions contributed by PADL Software are subject to the following license:

Copyright (c) 2011, PADL Software Pty Ltd. 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. Neither the name of PADL Software 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 PADL SOFTWARE 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 PADL SOFTWARE 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 bundled libev source code is subject to the following license:

All files in libev are Copyright (C)2007,2008,2009 Marc Alexander Lehmann.

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.

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.

Alternatively, the contents of this package may be used under the terms of the GNU General Public License ("GPL") version 2 or any later version, in which case the provisions of the GPL are applicable instead of the above. If you wish to allow the use of your version of this package only under the terms of the GPL and not to allow others to use your version of this file under the BSD license, indicate your decision by deleting the provisions above and replace them with the notice and other provisions required by the GPL in this and the other files of this package. If

you do not delete the provisions above, a recipient may use your version of this file under either the BSD or the GPL.

Files copied from the Intel AESNI Sample Library are subject to the following license:

Copyright © 2010, Intel Corporation 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 Intel Corporation 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.

# **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 aheared to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, Ihash, 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 rouines 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

#### 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.org madler@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, nocharge, 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 Hive

Copyright © 2008-2015 The Apache Software Foundation

#### Apache Spark

Copyright © 2014 The Apache Software Foundation

#### Apache ZooKeeper

Copyright © 2010 The Apache Software Foundation

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.