



Architecting the Future of Big Data

Hortonworks ODBC Driver with SQL Connector for Apache Spark

User Guide

Revised: March 17, 2017



© 2012-2015 Hortonworks Inc. All Rights Reserved.

Parts of this Program and Documentation include proprietary software and content that is copyrighted and licensed by Simba Technologies Incorporated. This proprietary software and content may include one or more feature, functionality or methodology within the ODBC, JDBC, ADO.NET, OLE DB, ODBO, XMLA, SQL and/or MDX component(s).

For information about Simba's products and services, visit: www.simba.com

Table of Contents

About the Hortonworks ODBC Driver with SQL Connector for Apache Spark	4
Contact Us	5
Windows Driver	6
Windows System Requirements	6
Installing the Driver on Windows	6
Creating a Data Source Name on Windows	7
Configuring a DSN-less Connection on Windows	9
Configuring Authentication on Windows	11
Configuring Advanced Options on Windows	15
Configuring HTTP Options on Windows	17
Configuring SSL Verification on Windows	17
Configuring Server-Side Properties on Windows	19
Configuring Logging Options on Windows	20
Configuring Kerberos Authentication for Windows	21
Verifying the Driver Version Number on Windows	26
macOS Driver	27
macOS System Requirements	27
Installing the Driver on macOS	27
Verifying the Driver Version Number on macOS	27
Linux Driver	29
Linux System Requirements	29
Installing the Driver Using the RPM File	29
Verifying the Driver Version Number on Linux	31
Configuring the ODBC Driver Manager on Non-Windows Machines	32
Specifying ODBC Driver Managers on Non-Windows Machines	32
Specifying the Locations of the Driver Configuration Files	33
Configuring ODBC Connections	35
Creating a Data Source Name on a Non-Windows Machine	35

Configuring a DSN-less Connection on a Non-Windows Machine	37
Configuring Authentication on a Non-Windows Machine	39
Configuring SSL Verification on a Non-Windows Machine	43
Configuring Server-Side Properties on a Non-Windows Machine	44
Testing the Connection	44
Configuring Logging Options	46
Authentication Mechanisms	48
Shark Server	48
Using a Connection String	50
DSN Connection String Example	50
DSN-less Connection String Examples	50
Features	53
SQL Connector for HiveQL	53
Data Types	53
Catalog and Schema Support	54
spark_system Table	55
Server-Side Properties	55
Get Tables With Query	55
Active Directory	56
Write-back	56
Security and Authentication	56
Driver Configuration Options	58
Configuration Options Appearing in the User Interface	58
Configuration Options Having Only Key Names	77
Contact Us	81
Third-Party Trademarks	82
Third-Party Licenses	83

About the Hortonworks ODBC Driver with SQL Connector for Apache Spark

The Hortonworks ODBC Driver with SQL Connector for Apache Spark 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 53.

The Hortonworks ODBC Driver with SQL Connector for Apache Spark 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 the *Data Access Standards Glossary*: <http://www.simba.com/resources/data-access-standards-library>. For complete information about the ODBC specification, see the *ODBC API Reference*: [http://msdn.microsoft.com/en-us/library/windows/desktop/ms714562\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/ms714562(v=vs.85).aspx).

The *User 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.



Architecting the Future of Big Data

Contact Us

If you have difficulty using the Hortonworks ODBC Driver with SQL Connector for Apache Spark, please contact our support staff. We welcome your questions, comments, and feature requests.

Please have a detailed summary of the client and server environment (OS version, patch level, Hadoop distribution version, Spark version, configuration, etc.) ready, before you call or write us. Supplying this information accelerates support.

By telephone:

USA: (855) 8-HORTON

International: (408) 916-4121

On the Internet:

Visit us at www.hortonworks.com

Windows Driver

Windows System Requirements

The Hortonworks ODBC Driver with SQL Connector for Apache Spark supports Apache Spark versions 0.8 through 2.1.

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

- One of the following operating systems:
 - Windows Vista, 7, 8, or 10
 - Windows Server 2008 or later
- 100 MB of available disk space
- Visual C++ Redistributable for Visual Studio 2013 installed (both 32- and 64-bit). You can download the installation packages at <https://www.microsoft.com/en-ca/download/details.aspx?id=40784>.

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

Installing the Driver on Windows

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

- `HortonworksSparkODBC32.msi` for 32-bit applications
- `HortonworksSparkODBC64.msi` for 64-bit applications

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

To install the Hortonworks ODBC Driver with SQL Connector for Apache Spark on Windows:

1. Depending on the bitness of your client application, double-click to run **HortonworksSparkODBC32.msi** or **HortonworksSparkODBC64.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**.

Creating a Data Source Name on Windows

Typically, after installing the Hortonworks ODBC Driver with SQL Connector for Apache Spark, you need to create a Data Source Name (DSN).

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

To create a Data Source Name on Windows:

1. Open the ODBC Administrator:
 - If you are using Windows 7 or earlier, click **Start**  > **All Programs** > **Hortonworks ODBC Driver with SQL Connector for Apache Spark 1.2** > **ODBC Administrator**.
 - Or, if you are using Windows 8 or later, on the Start screen, type **ODBC administrator**, and then click the **ODBC Administrator** search result.

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 Hortonworks ODBC Driver with SQL Connector for Apache Spark appears in the alphabetical list of ODBC drivers that are installed on your system.
3. Choose one:
 - To create a DSN that only the user currently logged into Windows can use, click the **User DSN** tab.
 - Or, to create a DSN that all users who log into Windows can use, click the **System DSN** tab.

Note:

It is recommended that you create a System DSN instead of a User DSN. Some applications, such as Sisense, 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 **Hortonworks ODBC Driver with SQL Connector for Apache Spark** and then click **Finish**. The Hortonworks ODBC Driver with SQL Connector for Apache Spark 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. In the **Host** field, type the IP address or host name of the Spark server.
10. In the **Port** field, type the number of the TCP port that the Spark server uses to listen for client connections.
11. In the **Database** field, type the name of the database schema to use when a schema is not explicitly specified in a query.

Note:

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

12. In the Authentication area, configure authentication as needed. For more information, see "Configuring Authentication on Windows" on page 11.

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 48.

13. 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.

14. 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 48.

15. 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 17.
16. To configure client-server verification over SSL, click **SSL Options**. For more information, see "Configuring SSL Verification on Windows" on page 17.

17. To configure advanced driver options, click **Advanced Options**. For more information, see "Configuring Advanced Options on Windows" on page 15.
18. 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 19.
19. To configure logging behavior for the driver, click **Logging Options**. For more information, see "Configuring Logging Options on Windows" on page 20.
20. To test the connection, click **Test**. Review the results as needed, and then click **OK**.

Note:

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

21. To save your settings and close the Hortonworks Spark ODBC Driver DSN Setup dialog box, click **OK**.
22. 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 Hortonworks Spark ODBC Driver Configuration tool that is installed with the Hortonworks ODBC Driver with SQL Connector for Apache Spark. The following section explains how to use the driver configuration tool. For information about using connection strings, see "Using a Connection String" on page 50.

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

1. Choose one:
 - If you are using Windows 7 or earlier, click **Start**  > **All Programs** > **Hortonworks ODBC Driver with SQL Connector for Apache Spark 1.2** > **Driver Configuration**.
 - Or, if you are using Windows 8 or later, click the arrow button at the bottom of the Start screen, and then click **Hortonworks ODBC Driver with SQL Connector for Apache Spark 1.2** > **Driver Configuration**.

Note:

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. In the Authentication area, configure authentication as needed. For more information, see "Configuring Authentication on Windows" on page 11.

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 48.

5. 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.

6. 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 48.

7. 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 17.
8. To configure client-server verification over SSL, click **SSL Options**. For more information, see "Configuring SSL Verification on Windows" on page 17.
9. To configure advanced options, click **Advanced Options**. For more information, see "Configuring Advanced Options on Windows" on page 15.
10. 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 19.
11. To save your settings and close the Hortonworks Spark ODBC Driver Configuration tool, click **OK**.

Configuring Authentication on Windows

Some Spark servers are configured to require authentication for access. To connect to a Spark server, you must configure the Hortonworks ODBC Driver with SQL Connector for Apache Spark 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 48.

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

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

Important:

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

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 Hortonworks 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 17.
4. To save your settings and close the dialog box, click **OK**.

Example connection string for Shark Server:

```
Driver=Hortonworks Spark ODBC Driver;Host=server_host;  
Port=10000;SparkServerType=1;AuthMech=0;Schema=Spark_database
```

Example connection string for Shark Server 2:

```
Driver=Hortonworks Spark ODBC Driver;Host=server_host;  
Port=10000;SparkServerType=2;AuthMech=0;Schema=Spark_database
```

Example connection string for Spark Thrift Server:

```
Driver=Hortonworks Spark ODBC Driver;Host=server_host;  
Port=10000;SparkServerType=3;AuthMech=0;Schema=Spark_database
```

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 21. For information about setting the Use Only SSPI advanced option, see "Configuring Advanced Options on Windows" on page 15.

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 Hortonworks 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. To allow the driver to pass your credentials directly to the server for use in authentication, select **Delegate Kerberos Credentials**.
6. In the **Service Name** field, type the service name of the Spark server.
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 17.
9. To save your settings and close the dialog box, click **OK**.

Example connection string:

```
Driver=Hortonworks Spark ODBC Driver;Host=server_host;  
Port=10000;SparkServerType=2;AuthMech=1;Schema=Spark_database;  
KrbRealm=Kerberos_realm;KrbHostFQDN=domain_name;  
KrbServiceName=service_name
```

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 Hortonworks 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**.

Example connection string:

```
Driver=Hortonworks Spark ODBC Driver;Host=server_host;  
Port=10000;SparkServerType=2;AuthMech=2;Schema=Spark_database;  
UID=user_name
```

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 Hortonworks 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.
7. 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 17.
8. To save your settings and close the dialog box, click **OK**.

Example connection string:

```
Driver=Hortonworks Spark ODBC Driver;Host=server_host;  
Port=10000;SparkServerType=2;AuthMech=3;Schema=Spark_database;  
UID=user_name;PWD=password
```

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 Hortonworks 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 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.

Note:

If necessary, you can create custom HTTP headers. For more information, see "Configuring HTTP Options on Windows" on page 17.

6. Click **SSL Options** and configure SSL settings as needed. For more information, see "Configuring SSL Verification on Windows" on page 17.
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.

Example connection string:

```
Driver=Hortonworks Spark ODBC Driver;Host=server_host;  
Port=10000;SparkServerType=2;AuthMech=6;  
Schema=Spark_database;UID=user_name;PWD=password;  
HTTPPath=Spark_HTTP_path
```

Configuring Advanced Options on Windows

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

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 Hortonworks Spark ODBC Driver Configuration tool, and then click **Advanced Options**.
2. To disable the SQL Connector feature, select the **Use Native Query** check box.

3. To defer query execution to SQLExecute, select the **Fast SQLPrepare** check box.
4. To allow driver-wide configurations to take precedence over connection and DSN settings, select the **Driver Config Take Precedence** check box.
5. To use the asynchronous version of the API call against 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.

7. To enable the driver to return SQL_WVARCHAR instead of SQL_VARCHAR for STRING and VARCHAR columns, and SQL_WCHAR instead of SQL_CHAR for CHAR columns, select the **Unicode SQL Character Types** check box.
8. To enable the driver to return the 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.

Note:

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.

Note:

For information about how to determine if your Spark server supports the HTTP transport protocol, see "Authentication Mechanisms" on page 48.

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 Hortonworks 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.

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 Hortonworks Spark ODBC Driver Configuration tool, and then click **SSL Options**.
2. Select the **Enable SSL** check box.
3. To allow self-signed certificates from the server, select the **Allow Self-signed Server Certificate** check box.
4. To allow the common name of a CA-issued SSL certificate to not match the host name of the 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 configure the driver to load certificates from a specific PEM file, specify the full path to the file in the **Trusted Certificates** field.
 - Or, to use the certificates in the PEM file that is installed with the driver, leave the **Trusted Certificates** field empty.
 - 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. 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.

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 Hortonworks 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.

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 check box is clear, then select the check box.

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

Configuring Logging Options on Windows

To help troubleshoot issues, you can enable logging. In addition to functionality provided in the Hortonworks ODBC Driver with SQL Connector for Apache Spark, 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.

To enable driver logging on Windows:

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

Logging Level	Description
OFF	Disables all logging.
FATAL	Logs severe error events that lead the driver to abort.
ERROR	Logs error events that might allow the driver to continue running.
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.
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. 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 Hortonworks ODBC Driver with SQL Connector for Apache Spark produces a log file named `SparkODBC_driver.log` at the location that you specify in the Log Path field.

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 Hortonworks ODBC Driver with SQL Connector for Apache Spark 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. For more information, see "Setting up One-Way Trust with Active Directory" in the Hortonworks documentation:
http://docs.hortonworks.com/HDPDocuments/HDP2/HDP-2.1.7/bk_installing_manually_book/content/ch23s05.html.

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-4.0.1-amd64.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-4.0.1-i386.msi>.

Note:

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

2. To run the installer, double-click the `.msi` file that you downloaded 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** , 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** , 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`.

Note:

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

Kerberos, 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 Hortonworks ODBC Driver with SQL Connector for Apache Spark 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. Open the ODBC Administrator:
 - If you are using Windows 7 or earlier, click **Start** > **All Programs** > **Hortonworks ODBC Driver with SQL Connector for Apache Spark 1.2** > **ODBC Administrator**.
 - Or, if you are using Windows 8 or later, on the Start screen, type **ODBC administrator**, and then click the **ODBC Administrator** search result.

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 Hortonworks ODBC Driver with SQL Connector for Apache Spark 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 Hortonworks ODBC Driver with SQL Connector for Apache Spark supports Apache Spark versions 0.8 through 2.1.

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:

- macOS version 10.9, 10.10, or 10.11
- 100 MB of available disk space
- iODBC 3.52.7 or later

Installing the Driver on macOS

The Hortonworks ODBC Driver with SQL Connector for Apache Spark is available for macOS as a .dmg file named `spark-odbc-native.dmg`. The driver supports both 32- and 64-bit client applications.

To install the Hortonworks ODBC Driver with SQL Connector for Apache Spark on macOS:

1. Double-click **spark-odbc-native.dmg** to mount the disk image.
2. Double-click **spark-odbc-native.dmg** 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/hortonworks/spark` directory.

6. To accept the installation location and begin the installation, click **Install**.
7. When the installation completes, click **Close**.

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 32.

Verifying the Driver Version Number on macOS

If you need to verify the version of the Hortonworks ODBC Driver with SQL Connector for Apache Spark 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.hortonworks.sparkodbc
```

The command returns information about the Hortonworks ODBC Driver with SQL Connector for Apache Spark that is installed on your machine, including the version number.

Linux Driver

Linux System Requirements

The Hortonworks ODBC Driver with SQL Connector for Apache Spark supports Apache Spark versions 0.8 through 2.1.

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

- One of the following distributions:
 - Red Hat® Enterprise Linux® (RHEL) 5, 6, or 7
 - CentOS 5, 6, or 7
 - SUSE Linux Enterprise Server (SLES) 11 or 12
- 150 MB of available disk space
- One of the following ODBC driver managers installed:
 - iODBC 3.52.7 or later
 - unixODBC 2.2.14 or later
- 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 to install and use the version of the driver that matches the bitness of the client application:

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

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

The placeholders in the file names are defined as follows:

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

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

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

To install the Hortonworks ODBC Driver with SQL Connector for Apache Spark using the RPM File:

1. Log in as the root user, and then navigate to the folder containing the RPM package for the driver.
2. 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 Hortonworks ODBC Driver with SQL Connector for Apache Spark files are installed in the `/usr/lib/spark/lib/native/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.

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 32.

Verifying the Driver Version Number on Linux

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

To verify the version number:

- Depending on your version of Linux, at the command prompt, run one of the following commands:

- `yum list | grep spark-odbc-native`
- `rpm -qa | grep spark-odbc-native`
- `dpkg -l | grep spark-odbc-native`

The command returns information about the Hortonworks ODBC Driver with SQL Connector for Apache Spark 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 Hortonworks ODBC Driver with SQL Connector for Apache Spark, do the following:

- Make sure that your machine uses the correct ODBC driver manager by setting the library path environment variable. For more information, see "Specifying ODBC Driver Managers on Non-Windows Machines" on page 32.
- If the driver configuration files are not stored in the default locations, then make sure that the ODBC driver manager locates and uses those files by setting environment variables. For more information, see "Specifying the Locations of the Driver Configuration Files" on page 33.

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 page 35.

Specifying ODBC Driver Managers on Non-Windows Machines

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

macOS

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

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

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

Linux

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

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

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

Specifying the Locations of the Driver Configuration Files

By default, ODBC driver managers are configured to use hidden versions of the `odbc.ini` and `odbcinst.ini` configuration files (named `.odbc.ini` and `.odbcinst.ini`) located in the home directory, as well as the `hortonworks.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 `HORTONWORKSSPARKINI` to the full path and file name of the `hortonworks.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 `HORTONWORKSSPARKINI` to the full path and file name of the `hortonworks.sparkodbc.ini` file.

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

For unixODBC:

```
export ODBCINI=/usr/local/odbc/odbc.ini
export ODBCSYSINI=/usr/local/odbc
export HORTONWORKSSPARKINI=/etc/hortonworks.sparkodbc.ini
```

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

1. If the `HORTONWORKSSPARKINI` 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 `hortonworks.sparkodbc.ini`.
3. The driver searches the current working directory of the application for a file named `hortonworks.sparkodbc.ini`.

4. The driver searches the home directory for a hidden file named `.hortonworks.sparkodbc.ini` (prefixed with a period).
5. The driver searches the `/etc` directory for a file named `hortonworks.sparkodbc.ini`.

Configuring ODBC Connections

The following sections describe how to configure ODBC connections when using the Hortonworks ODBC Driver with SQL Connector for Apache Spark on non-Windows platforms:

- "Creating a Data Source Name on a Non-Windows Machine" on page 35
- "Configuring a DSN-less Connection on a Non-Windows Machine" on page 37
- "Configuring Authentication on a Non-Windows Machine" on page 39
- "Configuring SSL Verification on a Non-Windows Machine" on page 43
- "Configuring Server-Side Properties on a Non-Windows Machine" on page 44
- "Configuring Logging Options" on page 46
- "Testing the Connection" on page 44

Creating a Data Source Name on a Non-Windows Machine

When connecting to your data store using a DSN, you only need to configure the `odbc.ini` file. Set the properties in the `odbc.ini` file to create a DSN that specifies the connection information for your data store. For information about configuring a DSN-less connection instead, see "Configuring a DSN-less Connection on a Non-Windows Machine" on page 37.

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

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

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

Note:

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

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

For example, on a macOS machine:

```
[ODBC Data Sources]
Sample DSN=Hortonworks Spark ODBC Driver
```

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

```
[ODBC Data Sources]
Sample DSN=Hortonworks 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  
/hortonworks/spark/lib/universal/libhortonworkssparkodbc  
.dylib
```

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

```
Driver=/usr/lib/spark/lib/native/Linux-i386-  
32/libhortonworkssparkodbc32.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, Spark 1.1 or later, then 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 39.
 - 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 43.
 - 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 44.
 - 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 Hortonworks ODBC Driver with SQL Connector for Apache Spark, see "Driver Configuration Options" on page 58.
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 33.

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=Hortonworks Spark ODBC Driver
[Sample DSN]
Driver=/
Library
/hortonworks/spark/lib/universal/libhortonworkssparkodbc.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 machine, containing a DSN that connects to a Thrift Server instance:

```
[ODBC Data Sources]
Sample DSN=Hortonworks Spark ODBC Driver 32-bit
[Sample DSN]
Driver=/usr/lib/spark/lib/native/Linux-i386-
32/libhortonworkssparkodbc32.so
SparkServerType=2
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]
Hortonworks ODBC Driver with SQL Connector for Apache
Spark=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
/hortonworks/spark/lib/universal/libhortonworkssparkodbc
.dylib
```

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

```
Driver=/usr/lib/spark/lib/native/Linux-i386-
32/libhortonworkssparkodbc32.so
```

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

For example:

```
Description=Hortonworks ODBC Driver with SQL Connector
for Apache Spark
```

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 33.

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

```
[ODBC Drivers]
Hortonworks Spark ODBC Driver=Installed
[Hortonworks Spark ODBC Driver]
```

```
Description=Hortonworks ODBC Driver with SQL Connector for  
Apache Spark  
Driver=  
Library  
/hortonworks/spark/lib/universal/libhortonworkssparkodbc.dylib
```

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

```
[ODBC Drivers]  
Hortonworks Spark ODBC Driver 32-bit=Installed  
Hortonworks Spark ODBC Driver 64-bit=Installed  
[Hortonworks Spark ODBC Driver 32-bit]  
Description=Hortonworks ODBC Driver with SQL Connector for  
Apache Spark (32-bit)  
Driver=/usr/lib/spark/lib/native/Linux-i386-  
32/libhortonworkssparkodbc32.so  
[Hortonworks Spark ODBC Driver 64-bit]  
Description=Hortonworks ODBC Driver with SQL Connector for  
Apache Spark (64-bit)  
Driver=/usr/lib/spark/lib/native/Linux-amd64-  
64/libhortonworkssparkodbc64.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 50.

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

- "Configuring Authentication on a Non-Windows Machine" on page 39
- "Configuring SSL Verification on a Non-Windows Machine" on page 43
- "Configuring Server-Side Properties on a Non-Windows Machine" on page 44

For detailed information about all the connection properties that the driver supports, see "Driver Configuration Options" on page 58.

Configuring Authentication on a Non-Windows Machine

Some Spark servers are configured to require authentication for access. To connect to a Spark server, you must configure the Hortonworks ODBC Driver with SQL Connector for Apache Spark 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 48.

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 `hortonworks.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 58.

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 43.

Example connection string for Shark Server:

```
Driver=Hortonworks Spark ODBC Driver;Host=server_host;  
Port=10000;SparkServerType=1;AuthMech=0;Schema=Spark_database
```

Example connection string for Shark Server 2:

```
Driver=Hortonworks Spark ODBC Driver;Host=server_host;  
Port=10000;SparkServerType=2;AuthMech=0;Schema=Spark_database
```

Example connection string for Spark Thrift Server:

```
Driver=Hortonworks Spark ODBC Driver;Host=server_host;  
Port=10000;SparkServerType=3;AuthMech=0;Schema=Spark_database
```

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 43.

Example connection string:

```
Driver=Hortonworks Spark ODBC Driver;Host=server_host;  
Port=10000;SparkServerType=2;AuthMech=1;Schema=Spark_database;  
KrbRealm=Kerberos_realm;KrbHostFQDN=domain_name;  
KrbServiceName=service_name
```

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.

Example connection string:

```
Driver=Hortonworks Spark ODBC Driver;Host=server_host;  
Port=10000;SparkServerType=2;AuthMech=2;Schema=Spark_database;  
UID=user_name
```

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 43.

Example connection string:

```
Driver=Hortonworks Spark ODBC Driver;Host=server_host;  
Port=10000;SparkServerType=2;AuthMech=3;Schema=Spark_database;  
UID=user_name;PWD=password
```

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 79.
6. Configure SSL settings as needed. For more information, see "Configuring SSL Verification on a Non-Windows Machine" on page 43.
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.

Example connection string:

```
Driver=Hortonworks Spark ODBC Driver;Host=server_host;  
Port=10000;SparkServerType=2;AuthMech=6;Schema=Spark_database;  
UID=user_name;PWD=password;HTTPPath=Spark_HTTP_path
```

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 connecting to a server over SSL, 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 `hortonworks.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 self-signed certificates from the server, 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.

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 `hortonworks.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`.

Testing the Connection

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.

Configuring Logging Options

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

Important:

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

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

To enable logging:

1. Open the `hortonworks.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 `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 `hortonworks.sparkodbc.ini` configuration file.
7. Restart your ODBC application to make sure that the new settings take effect.

The Hortonworks ODBC Driver with SQL Connector for Apache Spark produces a log file named `SparkODBC_driver.log` at the location you specify using the `LogPath` key.

To disable logging:

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

Authentication Mechanisms

To connect to a Spark server, you must configure the Hortonworks ODBC Driver with SQL Connector for Apache Spark 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:

<code>hive.server2.authentication</code>	Authentication Mechanism
NOSASL	No Authentication

<code>hive.server2.authentication</code>	Authentication Mechanism
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:

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

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

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

For detailed instructions on how to configure authentication when using a non-Windows driver, see "Configuring Authentication on a Non-Windows Machine" on page 39.

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 58.

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:

- *[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=Hortonworks Spark ODBC Driver;SparkServerType=1;  
Host=[Server];Port=[PortNumber];
```

For example:

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

For example:

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

For example:

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

For example:

```
Driver=Hortonworks Spark ODBC Driver;Host=192.168.222.160;  
Port=10000;AuthMech=1;KrbRealm=HORTONWORKS;
```

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

For example:

```
Driver=Hortonworks Spark ODBC Driver;Host=192.168.222.160;  
Port=10000;AuthMech=3;UID=hortonworks;PWD=hortonworks;
```

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

For example:

```
Driver=Hortonworks Spark ODBC Driver;Host=192.168.222.160;  
Port=10000;AuthMech=5;UID=hortonworks;PWD=hortonworks;  
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=Hortonworks Spark ODBC Driver;Host=[Server];  
Port=[PortNumber];AuthMech=6;UID=[YourUserName];  
PWD=[YourPassword];HTTPPath=[ServerURL];
```

For example:

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

Features

For more information on the features of the Hortonworks ODBC Driver with SQL Connector for Apache Spark, see the following:

- "SQL Connector for HiveQL" on page 53
- "Data Types" on page 53
- "Catalog and Schema Support" on page 54
- "spark_system Table" on page 55
- "Server-Side Properties" on page 55
- "Get Tables With Query" on page 55
- "Active Directory" on page 56
- "Write-back" on page 56
- "Security and Authentication" on page 56

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 Hortonworks ODBC Driver with SQL Connector for Apache Spark 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.

Catalog and Schema Support

The Hortonworks ODBC Driver with SQL Connector for Apache Spark 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 Hortonworks ODBC Driver with SQL Connector for Apache Spark 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 Hortonworks 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 `hortonworks.sparkodbc.ini` file. Properties specified in the driver configuration tool or the `hortonworks.sparkodbc.ini` file apply to all connections that use the Hortonworks ODBC Driver with SQL Connector for Apache Spark.

For more information about setting server-side properties when using the Windows driver, see "Configuring Server-Side Properties on Windows" on page 19. 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 44.

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 Hortonworks ODBC Driver with SQL Connector for Apache Spark 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. For more information, see "Setting up One-Way Trust with Active Directory" in the Hortonworks documentation:
http://docs.hortonworks.com/HDPDocuments/HDP2/HDP-2.1.7/bk_installing_manually_book/content/ch23s05.html.

Write-back

The Hortonworks ODBC Driver with SQL Connector for Apache Spark supports translation for INSERT syntax when connecting to a Spark Thrift Server instance that is running Spark 1.3 or later.

Spark does not support UPDATE or DELETE syntax.

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 Hortonworks ODBC Driver with SQL Connector for Apache Spark 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 48. For detailed driver configuration instructions, see "Configuring Authentication on Windows" on page 11 or "Configuring Authentication on a Non-Windows Machine" on page 39.

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 17 or "Configuring SSL Verification on a Non-Windows Machine" on page 43.

Driver Configuration Options

Driver Configuration Options lists the configuration options available in the Hortonworks ODBC Driver with SQL Connector for Apache Spark 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 Hortonworks Spark ODBC Driver Configuration tool and the following dialog boxes:

- Hortonworks 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 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 `hortonworks.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 Hortonworks ODBC Driver with SQL Connector for Apache Spark, or via the key name when using a connection string or configuring a connection from a Linux/macOS machine:

- "Allow Common Name Host Name Mismatch" on page 59
- "Allow Self-Signed Server Certificate" on page 60
- "Apply Properties with Queries" on page 60
- "Async Exec Poll Interval" on page 60
- "Binary Column Length" on page 61
- "Client Certificate File" on page 61
- "HTTP Path" on page 66
- "Invalid Session Auto Recover" on page 67
- "Log Level" on page 67
- "Log Path" on page 68
- "Max File Size" on page 68
- "Max Number Files" on page 69
- "Mechanism" on page 69
- "Password" on page 70

- "Client Private Key File" on page 61
- "Client Private Key Password" on page 62
- "Convert Key Name to Lower Case" on page 62
- "Database" on page 62
- "Decimal Column Scale" on page 63
- "Default String Column Length" on page 63
- "Delegate Kerberos Credentials" on page 63
- "Delegation UID" on page 64
- "Driver Config Take Precedence" on page 64
- "Enable Auto Reconnect" on page 64
- "Enable SSL" on page 65
- "Fast SQLPrepare" on page 65
- "Get Tables With Query" on page 65
- "Host(s)" on page 66
- "Host FQDN" on page 66
- "Port" on page 70
- "Realm" on page 70
- "Rows Fetched Per Block" on page 71
- "Save Password (Encrypted)" on page 71
- "Service Name" on page 71
- "Show System Table" on page 72
- "Socket Timeout" on page 72
- "Spark Server Type" on page 72
- "Thrift Transport" on page 73
- "Trusted Certificates" on page 74
- "Two Way SSL" on page 74
- "Unicode SQL Character Types" on page 75
- "Use Async Exec" on page 75
- "Use Native Query" on page 75
- "Use Only SSPI" on page 76
- "Use System Trust Store" on page 76
- "User Name" on page 77

Allow Common Name Host Name Mismatch

Key Name	Default Value	Required
CAIssuedCertNamesMismatch	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 Self-Signed Server Certificate

Key Name	Default Value	Required
AllowSelfSignedServerCert	Clear (0)	No

Description

This option specifies whether the driver allows self-signed certificates from the server.

- Enabled (1): The driver authenticates the Spark server even if the server is using a self-signed certificate.
- Disabled (0): The driver does not allow self-signed certificates from the server.

Note:

This setting is applicable only when SSL is enabled.

Apply Properties with Queries

Key Name	Default Value	Required
ApplySSPWithQueries	Selected (1)	No

Description

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.

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 62.

Note:

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

Client Private Key Password

Key Name	Default Value	Required
<code>ClientPrivateKeyPassword</code>	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
<code>LCaseSspKeyName</code>	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 lower-case characters.
- Disabled (0): The driver does not modify the server-side property key names.

Database

Key Name	Default Value	Required
<code>Schema</code>	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
DefaultStringColumnLength	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

Description

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.

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

The IP address or host name of the Spark server.

Host FQDN

Key Name	Default Value	Required
KrbHostFQDN	<code>_HOST</code>	Yes, if the authentication mechanism is Kerberos.

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	<p><code>/spark</code> if using Windows Azure HDInsight Service (6).</p> <p><code>/</code> if using non-Windows Azure HDInsight Service with Thrift Transport set to HTTP (2).</p>	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
<code>InvalidSessionAutoRecover</code>	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
<code>LogLevel</code>	OFF (0)	No

Description

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

Important:

- Only enable logging long enough to capture an issue. Logging decreases performance and can consume a large quantity of disk space.
- 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 `hortonworks.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 a log file named `SparkODBC_driver.log` in the location specified in the Log Path (`LogPath`) property.

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 `hortonworks.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 `hortonworks.sparkodbc.ini` file.

Max Number Files

Key Name	Default Value	Required
LogFileCount	50	No

Description

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 `hortonworks.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)

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).

Port

Key Name	Default Value	Required
Port	<ul style="list-style-type: none"> • non-HDInsight clusters: 10000 • Windows Azure HDInsight Emulator: 10001 • Windows Azure HDInsight Service: 443 	Yes

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 Hortonworks 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 Name

Key Name	Default Value	Required
KrbServiceName	spark	Yes, if the authentication mechanism is Kerberos.

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.

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 48.

Trusted Certificates

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

Description

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.

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 61, "Client Private Key File" on page 61, and "Client Private Key Password" on page 62.
- Disabled (0): The server does not verify the client. Depending on whether one-way SSL is enabled, the client might verify the server. For more information, see "Enable SSL" on page 65.

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 65.

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 them into an equivalent form in HiveQL.

- 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.

Note:

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

Use Only SSPI

Key Name	Default Value	Required
UseOnlySSPI	Clear (0)	No

Description

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

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

Important:

This option is available only in the Windows driver.

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.

Note:

This option is only available on Windows.

User Name

Key Name	Default Value	Required
UID	For User Name (2) authentication only, the default value is <code>anonymous</code>	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.

Configuration Options Having Only Key Names

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

- "ADUserNameCase" on page 77
- "Driver" on page 78
- "ForceSynchronousExec" on page 78
- "http.header." on page 79
- "SSP_" on page 79

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 possible:

- `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	Hortonworks 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 (Hortonworks 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 73.

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.



Architecting the Future of Big Data

Contact Us

If you have difficulty using the Hortonworks ODBC Driver with SQL Connector for Apache Spark, please contact our support staff. We welcome your questions, comments, and feature requests.

Please have a detailed summary of the client and server environment (OS version, patch level, Hadoop distribution version, Spark version, configuration, etc.) ready, before you call or write us. Supplying this information accelerates support.

By telephone:

USA: (855) 8-HORTON

International: (408) 916-4121

On the Internet:

Visit us at www.hortonworks.com

Third-Party Trademarks

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.

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.

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 NON-INFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

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

Cyrus SASL License

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

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

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

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

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

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

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

dtoa License

The author of this software is David M. Gay.

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

Permission to use, copy, modify, and distribute this software for any purpose without fee is hereby granted, provided that this entire notice is included in all copies of any software

which is or includes a copy or modification of this software and in all copies of the supporting documentation for such software.

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

Expat License

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

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

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

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

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 NON-INFRINGEMENT 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.

OpenSSL License

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

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

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

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

4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment:

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

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR

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

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

Original SSLeay License

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

All rights reserved.

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

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

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

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

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

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

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

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

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

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

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

RapidJSON License

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

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

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

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

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

Other dependencies and licenses:

Open Source Software Licensed Under the BSD License:

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

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

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

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

Open Source Software Licensed Under the JSON License:

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

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

Terms of the JSON License:

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

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

The Software shall be used for Good, not Evil.

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

Terms of the MIT License:

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

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

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

Stringencoders License

Copyright 2005, 2006, 2007

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

All rights reserved.

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

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

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

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

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

This is the standard "new" BSD license:

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

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

jloup@gzip.org

Mark Adler

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

Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

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

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

5. **Submission of Contributions.** Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. **Trademarks.** This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. **Disclaimer of Warranty.** Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
8. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or

consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

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

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

Copyright [yyyy] [name of copyright owner]

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

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

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

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

Apache Hive

Copyright © 2008-2015 The Apache Software Foundation

Apache Spark

Copyright © 2014 The Apache Software Foundation



Architecting the Future of Big Data

Apache Thrift

Copyright © 2006-2010 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.