




# Simba Drill JDBC Driver with SQL Connector

## Installation and Configuration Guide

Simba Technologies Inc.

Version 1.6.0  
October 10, 2018



**Copyright © 2018 Simba Technologies Inc. All Rights Reserved.**

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

**Trademarks**

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

**Contact Us**

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

Tel: +1 (604) 633-0008

Fax: +1 (604) 633-0004

[www.simba.com](http://www.simba.com)

## About This Guide

### Purpose

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

### Audience

The guide is intended for end users of the Simba Drill JDBC Driver.

### Knowledge Prerequisites

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

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

### Document Conventions

*Italics* are used when referring to book and document titles.

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

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

#### **Note:**

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

#### **Important:**

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

# Table of Contents

About the Simba Drill JDBC Driver .....	6
About Apache Drill .....	6
About the Driver .....	6
System Requirements .....	8
Simba Drill JDBC Driver Files .....	9
Installing and Using the Simba Drill JDBC Driver .....	10
Referencing the JDBC Driver Libraries .....	10
Registering the Driver Class .....	11
Building the Connection URL .....	12
Configuring Authentication .....	14
Using Your Drill User Name and Password .....	14
Using Kerberos .....	14
Using MapR-SASL .....	16
Configuring TLS (SSL) Verification .....	18
Configuring Logging .....	19
Features .....	21
Catalog and Schema Support .....	21
Schema Types and File Formats .....	21
Specifying Column Names .....	22
Data Types .....	24
Casting Binary Data .....	24
Security and Authentication .....	25
Driver Configuration Options .....	26
AuthMech or auth .....	26
CastAnyToVarchar .....	27
CatalogSupport .....	27
DelegationUID .....	27
enableTLS .....	28
disableCertificateVerification .....	28
disableHostVerification .....	28
ExcludedSchemas .....	29
GetMetadataWithQueries .....	29

## Installation and Configuration Guide

IncludedSchemas .....	30
KrbHostFQDN or service_host .....	30
KrbKeytab or keytab .....	31
KrbPrincipal or principal .....	31
KrbRealm or realm .....	31
KrbServiceName or service_name .....	32
LogLevel .....	32
LogPath .....	33
NettyMaxDirectMem .....	33
NumNettyDirectArenas .....	34
PWD .....	34
RemoveQryCatalog .....	34
sasl_encrypt .....	35
Schema .....	35
StringColumnLength .....	35
TLSHandshakeTimeout .....	36
TLSProtocol .....	36
TLSProvider .....	36
trustStorePassword .....	37
trustStorePath .....	37
trustStoreType .....	37
UID .....	38
UpdateColMetadataOnSchemaChange .....	38
Third-Party Trademarks .....	39
Third-Party Licenses .....	40

## About the Simba Drill JDBC Driver

### About Apache Drill

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

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

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

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

 **Note:**

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

### About the Driver

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

The Simba Drill JDBC Driver complies with the JDBC 4.0 and 4.1 data standards. JDBC 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 JDBC driver, which connects an application to the database. For more information about JDBC, see *Data*

*Access Standards* on the Simba Technologies

website: <https://www.simba.com/resources/data-access-standards-glossary>.

This guide is suitable for users who want to access data residing within Drill from their desktop environment. Application developers might also find the information helpful. Refer to your application for details on connecting via JDBC.

## System Requirements

Each machine where you use the Simba Drill JDBC Driver must have Java Runtime Environment (JRE) 8.0 installed.



## Simba Drill JDBC Driver Files

The Simba Drill JDBC Driver is delivered in a ZIP archive named `DrillJDBC41-[Version].zip`, where *[Version]* is the version number of the driver.

The archive contains the driver supporting the JDBC API version indicated in the archive name, as well as release notes and third-party license information. In addition, the required third-party libraries and dependencies are packaged and shared in the driver JAR file in the archive.

## Installing and Using the Simba Drill JDBC Driver

To install the Simba Drill JDBC Driver on your machine, extract the files from the ZIP archive to the directory of your choice.

### ! Important:

If you received a license file through email, then you must copy the file into the same directory as the driver JAR file before you can use the Simba Drill JDBC Driver.

To access a Drill data store using the Simba Drill JDBC Driver, you need to configure the following:

- The list of driver library files (see [Referencing the JDBC Driver Libraries](#) on page 10)
- The `Driver` or `DataSource` class (see [Registering the Driver Class](#) on page 11)
- The connection URL for the driver (see [Building the Connection URL](#) on page 12)

### ! Important:

The Simba Drill JDBC Driver provides read-only access to Drill data stores.

## Referencing the JDBC Driver Libraries

Before you use the Simba Drill JDBC Driver, the JDBC application or Java code that you are using to connect to your data must be able to access the driver JAR files. In the application or code, specify all the JAR files that you extracted from the ZIP archive.

### Using the Driver in a JDBC Application

Most JDBC applications provide a set of configuration options for adding a list of driver library files. Use the provided options to include all the JAR files from the ZIP archive as part of the driver configuration in the application. For more information, see the documentation for your JDBC application.

### Using the Driver in Java Code

You must include all the driver library files in the class path. This is the path that the Java Runtime Environment searches for classes and other resource files. For more information, see "Setting the Class Path" in the appropriate Java SE Documentation.

For Java SE 8:

- For Windows:  
<http://docs.oracle.com/javase/8/docs/technotes/tools/windows/classpath.html>
- For Linux and Solaris:  
<http://docs.oracle.com/javase/8/docs/technotes/tools/solaris/classpath.html>

## Registering the Driver Class

Before connecting to the data store, you must register the appropriate class for your application.

The Simba Drill JDBC Driver provides the following fully-qualified class names (FQCNs) to support JDBC 4.1 connections to Drill:

- `com.simba.drill.jdbc41.Driver`
- `com.simba.drill.jdbc41.DataSource`

The `Driver` class extends `java.sql.Driver`, and the `DataSource` class extends `javax.sql.DataSource` and `javax.sql.ConnectionPoolDataSource`.

The following sample code shows how to use the `DriverManager` to establish a connection:

```
private static Connection connectViaDM() throws Exception
{
    Connection connection = null;
    Class.forName(DRIVER_CLASS);
    connection = DriverManager.getConnection(CONNECTION_URL);
    return connection;
}
```

The following sample code shows how to use the `DataSource` class to establish a connection:

```
private static Connection connectViaDS() throws Exception
{
    Connection connection = null;
    Class.forName(DRIVER_CLASS);
    DataSource ds = new com.simba.drill.jdbc41.DataSource();
    ds.setURL(CONNECTION_URL);
    connection = ds.getConnection();
    return connection;
}
```

## Building the Connection URL

Use the connection URL to supply connection information to the data store that you are accessing.

### Connecting to a Drillbit

The following is the format of a connection URL for connecting to a drillbit, where *[Host]* is the DNS or IP address of the server:

```
jdbc:drill:drillbit=[Host]
```

By default, the driver connects to port 31010.

You can specify optional settings such as the number of the TCP port to connect to or any of the connection properties supported by the driver or your Drill implementation. For a list of the properties available in the driver, see [Driver Configuration Options](#) on page 26.

The following is the format of a connection URL that specifies some optional settings:

```
jdbc:drill:drillbit=[Host]:[Port];[Property1]=[Value];  
[Property2]=[Value];...
```

For example, to connect to port 31500 on a drillbit and authenticate the connection using a user name and password, you would use the following connection URL:

```
jdbc:drill:drillbit=192.168.1.1:31500;AuthMech=Plain;  
UID=simba;PWD=simba
```

#### ! Important:

- Properties are case-sensitive.
- Do not duplicate properties in the connection URL.

### Connecting to a ZooKeeper Cluster

The following is the format of a connection URL for connecting to a ZooKeeper cluster, where *[ServerList]* is a comma-separated list of servers in your ZooKeeper cluster and *[ClusterName]* is the unique name of the drillbit cluster that you want to use:

```
jdbc:drill:zk=[ServerList]/drill/[ClusterName]
```

Optionally, you can specify any of the connection properties supported by the driver or your Drill implementation. For a list of the properties available in the driver, see [Driver Configuration Options](#) on page 26.

The following is the format of a connection URL that specifies some optional settings:

```
jdbc:drill:zk=[ServerList]/drill/[ClusterName];[Property1]=[Value];[Property2]=[Value];...
```

For example, to connect to a ZooKeeper cluster and authenticate the connection using a user name and password, you would use the following connection URL:

```
jdbc:drill:zk=192.168.39.43:5181, 192.168.1.1:31010/drill/drillbits1;AuthMech=Plain;UID=simba;PWD=simba
```

**! Important:**

- Properties are case-sensitive.
- Do not duplicate properties in the connection URL.

## Configuring Authentication

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

- [Using Your Drill User Name and Password](#) on page 14
- [Using Kerberos](#) on page 14
- [Using MapR-SASL](#) on page 16

### Note:

For some of the connection properties used for authentication, the driver supports more than one property name. For example, you can set either `AuthMech` or `auth` to specify the authentication mechanism to use. For complete information about each connection property, see [Driver Configuration Options](#) on page 26.

## Using Your Drill User Name and Password

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

You provide this information to the driver in the connection URL. For more information about the syntax of the connection URL, see [Building the Connection URL](#) on page 12.

### To configure user name and password authentication:

1. Set the `AuthMech` property to `Plain`.
2. Set the `UID` property to an appropriate user name for accessing the Drill server.
3. Set the `PWD` property to the password corresponding to the user name you provided.

For example:

```
jdbc:drill:drillbit=192.168.1.1:31500;AuthMech=Plain;
UID=simba;PWD=simba
```

## Using Kerberos

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

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

You can configure Kerberos authentication in the driver using a keytab or a service principal name, depending on the specific requirements of your Kerberos setup.

You provide this information to the driver in the connection URL. For more information about the syntax of the connection URL, see [Building the Connection URL](#) on page 12.

### To configure Kerberos authentication using a keytab:

1. Set the `AuthMech` property to `Kerberos`.
2. Set the `KrbKeytab` property to the full path and file name of your keytab file.
3. Set the `UID` property to your client principal name.
4. Optionally, set the `KrbRealm` property to the realm name of the service principal. If you do not set this property, the driver uses the default realm defined in your Kerberos setup.

For example:

```
jdbc:drill:drillbit=192.168.1.1:31500;AuthMech=Kerberos;  
KrbKeytab=C:\myKeytabs\myUser.keytab;UID=simba;  
KrbRealm=drill.example.com
```

### To configure Kerberos authentication using the service principal name:

1. Run the `kinit` command using the following syntax, where `[Keytab]` is the Kerberos credential and `[Principal]` is the Kerberos user principal to use for authentication:

```
kinit -kt [Keytab] [Principal]
```

2. In your connection URL, set the `AuthMech` property to `Kerberos`.
3. Do one of the following:
  - To specify a complete service principal name including the primary and instance (and optionally, the realm), set the `KrbPrincipal` property to the service principal name.

For example, the following settings are both valid:

```
KrbPrincipal=jsmith/simbadrill  
KrbPrincipal=jsmith/simbadrill@drill.example.com
```

- Or, to specify one or more components of the service principal name individually, configure one or more of the following connection properties as needed:

Component	Configuration Instructions
Primary	<p>Set the <code>KrbServiceName</code> property to the primary name of the service principal.</p> <p>If you do not set this property, the driver uses the default value <code>drill</code> as the primary name.</p>
Instance	<p>Set the <code>KrbHostFQDN</code> property to the instance name of the service principal.</p> <p>If you do not set this property, the driver uses the host name of the drillbit as the instance name.</p>
Realm	<p>Set the <code>KrbRealm</code> property to the realm name of the service principal.</p> <p>If you do not set this property, the driver uses the default realm defined in your Kerberos setup.</p>

For example, to specify a complete service principal name for authentication:

```
jdbc:drill:drillbit=192.168.1.1:31500;AuthMech=Kerberos;
KrbPrincipal=jsmith/simbadrill@drill.example.com
```

As another example, to specify your primary name only so that the driver authenticates the connection using the host name of the drillbit node as the instance name and the default realm defined in your Kerberos setup:

```
jdbc:drill:drillbit=192.168.1.1:31500;AuthMech=Kerberos;
KrbServiceName=jsmith
```

## Using MapR-SASL

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

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



You provide this information to the driver in the connection URL. For more information about the syntax of the connection URL, see [Building the Connection URL](#) on page 12.

**To configure MapR-SASL authentication:**

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

For example:

```
jdbc:drill:drillbit=192.168.1.1:31500;AuthMech=MapRSASL
```

## Configuring TLS (SSL) Verification

If you are connecting to a Drill instance that has Transport Layer Security (TLS) enabled (previously called SSL), you can configure the driver to connect to a TLS-enabled socket. When using TLS to connect to a server, the driver can be configured to verify the identity of the server.

You provide this information to the driver in the connection URL. For more information about the syntax of the connection URL, see [Building the Connection URL](#) on page 12.

### To configure TLS verification:

1. To enable TLS connections, set the `enableTLS` property to `true`.
2. TLS provider defaults to JDK. If you are using OpenSSL as your TLS provider, set the `TLSProvider` property to `OPENSSL`.
3. TLS protocol defaults to TLSV1.2. If your provider is using a different version, set the `TLSProtocol` to the version used by your TLS provider. For a list of supported versions, see [TLSProtocol](#) on page 36.
4. Set the `TLSHandshakeTimeout` to the number of milliseconds you want the driver to wait before alerting the user a connection has not been established (10000 by default).
5. To use the java trust store:
  - Set the `trustStorePath` property to the custom path for the Java TrustStore. (If not set the driver will use a default path. See [trustStorePath](#) on page 37)
  - Set the `trustStoreType` property to the type of TrustStore used.
  - If your trust store uses a password, set the `trustStorePassword` property to the password.
6. If you want to use the Windows trust store, set the `useSystemTrustStore` property to `true`.
7. If you do not want the driver to verify the host against the certificate, set the `disableHostVerification` property to `true`.

#### ! Important:

If `disableHostVerification` is set to `true`, the driver does not verify the host in the certificate is the host being connected to.

8. If you do not want the driver to validate the host certificate against the Trust Store, set the `disableCertificateVerification` property to `true`.

## Configuring Logging

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.

In the connection URL, set the `LogLevel` key to enable logging at the desired level of detail. The following table lists the logging levels provided by the Simba Drill JDBC Driver, in order from least verbose to most verbose.

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

### To enable logging:

1. Set the `LogLevel` property to the desired level of information to include in log files.
2. Set the `LogPath` property to the full path to the folder where you want to save log files. To make sure that the connection URL is compatible with all JDBC applications, escape the backslashes (`\`) in your file path by typing another backslash.

For example, the following connection URL enables logging level 3 and saves the log files in the `C:\temp` folder:

```
jdbc:drill:drillbit=localhost;LogLevel=3;LogPath=C:\\temp
```

3. To make sure that the new settings take effect, restart your JDBC application and reconnect to the server.

The Simba Drill JDBC Driver produces the following log files in the location specified in the `LogPath` property:

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

If the `LogPath` value is invalid, then the driver sends the logged information to the standard output stream (`System.out`).

#### To disable logging:

1. Set the `LogLevel` property to 0.
2. To make sure that the new setting takes effect, restart your JDBC application and reconnect to the server.

## Features

More information is provided on the following features of the Simba Drill JDBC Driver:

- [Catalog and Schema Support](#) on page 21
- [Schema Types and File Formats](#) on page 21
- [Specifying Column Names](#) on page 22
- [Data Types](#) on page 24
- [Casting Binary Data](#) on page 24
- [Security and Authentication](#) on page 25

## Catalog and Schema Support

The Simba Drill JDBC Driver supports both catalogs and schemas to make it easy for the driver to work with various JDBC applications. The Simba Drill JDBC Driver follows convention, adding a catalog named DRILL under which all of the schemas and databases are organized. The driver also maps the JDBC schema to the DRILL catalog. The driver provides access to all of the schemas/databases that are listed under this catalog, ensuring compatibility with standard BI tools.

By default, the driver indicates to the application that it has catalog support through DatabaseMetadata. You can disable this behavior by setting the `CatalogSupport` connection property to `LEGACY`. For more information, see [CatalogSupport](#) on page 27.

## Schema Types and File Formats

The Simba Drill JDBC Driver supports the following schema types:

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

## Specifying Column Names

Use of the SQL asterisk (\*) selector to return all columns in a table is limited. You must modify your queries to specify the columns to return.

Depending on the schema type or file format, refer to the following syntax guidelines and modify your queries as needed. Note that when you write queries for Drill, you must enclose all table and schema names in backticks (`).

### Note:

The Simba ODBC Driver for Apache Drill provides an application called Drill Explorer that helps generate queries for Drill. However, the Simba Drill JDBC Driver currently does not support Drill Explorer. If you would like to simplify the process of formulating valid queries for Drill, consider using the Drill Explorer and the ODBC driver.

## HBase

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

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

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

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

## Parquet

When a query is used as a subquery, such as when you use a SQL view to connect to a data store using SAP Lumira, you must specify column names in the query. For example, SAP Lumira might not connect to your Drill data source successfully when using the following query:

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

When the query is modified to specify column names, SAP Lumira connects to your Drill data store successfully using the SQL view:

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

## JSON

When a query is used as a subquery, such as when you use a SQL view to connect to a data store using SAP Lumira, you must specify column names in the query. For example, SAP Lumira might not connect to your Drill data store successfully when using the following query:

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

When the query is modified to specify column names, SAP Lumira connects to your Drill data store successfully using the SQL view:

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

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

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

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

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

## CSV or TSV

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

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

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

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

## Hive

When using a Hive schema, you can use standard queries without modifying them.

## Data Types

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

The following table lists the supported data type mappings.

Drill Type	SQL Type	Java Type
BIGINT	BIGINT	BIGINT
BIT	BOOLEAN	BOOLEAN
DATE	DATE	DATE
DATETIME	DATE	DATE
FLOAT4	FLOAT	REAL
FLOAT8	DOUBLE	DOUBLE
INT	INT	INTEGER
SMALLINT	SMALLINT	SMALLINT
TIMESTAMP	TIMESTAMP	TIMESTAMP
TINYINT	TINYINT	TINYINT
VARCHAR	VARCHAR	VARCHAR
VARDECIMAL	DECIMAL	BIGDECIMAL

## Casting Binary Data

Given that Drill can work with self-describing data such as HBase and file systems without central metadata definitions, there are scenarios where the file formats do not have defined data types for the data. HBase, for example, always treats data as binary. Drill provides auxiliary functions to cast (or interpret) the data as certain data types. The following are a few examples of how to cast data as a specific data type. For more



information about the SQL queries in Drill, see the Apache Drill documentation: <http://drill.apache.org/docs/>.

**! Important:**

When writing queries for Drill, you must enclose all table and schema names in backticks (`). For examples of this syntax, see the queries below.

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

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

The following query displays the same results in string format:

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

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

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

The following query displays the same results in string format:

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

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

## Security and Authentication

To protect data from unauthorized access, some Drill data stores require connections to be authenticated using user credentials. The Simba Drill JDBC Driver supports the authentication protocols used in Drill by providing a mechanism that allows you to authenticate your connection using your Drill user name and password, the Kerberos protocol, or the MapR-SASL protocol. For detailed configuration instructions, see [Configuring Authentication](#) on page 14.


## Driver Configuration Options

Driver Configuration Options lists and describes the properties that you can use to configure the behavior of the Simba Drill JDBC Driver.

 **Note:**

Your connection string can also include any key and value option pairs supported by your Drill implementation.

You can set configuration properties using the connection URL. For more information, see [Building the Connection URL](#) on page 12.

 **Note:**

Property names and values are case-sensitive.

### AuthMech or auth

Default Value	Data Type	Required
None	String	No

### Description

The authentication mechanism to use. Set the property to one of the following values:

- `Plain`: The driver authenticates the connection using your Drill user name and password.
- `Kerberos`: The driver authenticates the connection using the Kerberos protocol.
- `MapRSASL`: The driver authenticates the connection using the MapR-SASL protocol.

If the server you are connecting to does not require authentication, then do not set this property and do not specify any credentials in the connection URL.

If this property is not set but some credentials are specified in the connection URL, then the Drill client tries to detect and use the appropriate authentication mechanism.

## CastAnyToVarchar

Default Value	Data Type	Required
false	String	No

### Description

This property specifies whether the driver casts columns of type ANY to type VARCHAR.

- `true`: If `SQLColumns` returns columns of type ANY, then the driver casts them to type VARCHAR.
- `false`: The driver does not change the returned columns.

## CatalogSupport

Default Value	Data Type	Required
SUPPORTED	String	No

### Description

This property specifies whether the driver indicates to the application that it has catalog support.

- `SUPPORTED`: The driver indicates to the application that it has catalog support through `DatabaseMetadata`. The driver returns DRILL as the catalog under all catalog functions and indicates that the current catalog is DRILL.
- `LEGACY`: The driver indicates to the application that there is no catalog support through `DatabaseMetadata`. However, the driver stills return DRILL as the catalog name under all catalog functions other than `getSchemas`. `getSchemas` returns the catalog column as Null.
- `NOT_SUPPORTED`: The driver reports it does not support catalogs.

## DelegationUID

Default Value	Data Type	Required
None	String	No

## Description

Use this option to delegate all operations against Drill to a specific user.

## enableTLS

Default Value	Data Type	Required
false	String	No

## Description

This property specifies whether the driver uses SSL.

- `true`: The driver will communicate with the data source using TLS protocol for SSL.
- `false`: The driver will not use SSL when communicating with the data source.

## disableCertificateVerification

Default Value	Data Type	Required
false	String	No

## Description

This property specifies if the driver verifies the certificate being used against the system truststore.

- `false`: The driver verifies the certificate against the system truststore.
- `true`: The driver does not verify the certificate against the system truststore.

## disableHostVerification

Default Value	Data Type	Required
false	String	No

## Description

This property specifies if the driver verifies that the host in the certificate is the host being connected to.

- `false`: The driver verifies the certificate against the host being connected to.
- `true`: The driver does not verify the certificate against the host being connected to.

### ! Important:

If `disableHostVerification` is set to `true`, the driver does not verify the host in the certificate is the host being connected to.

## ExcludedSchemas

Default Value	Data Type	Required
None	String	No

## Description

A comma-separated list of schemas that the driver excludes when reporting schema information to an application.

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

## GetMetadataWithQueries

Default Value	Data Type	Required
None	Integer	No

## Description

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

- 1: The driver uses queries to retrieve metadata.
- 0: The driver uses native API calls when it is connected to a server that supports native API calls.

## IncludedSchemas

Default Value	Data Type	Required
None	String	No

### Description

A comma-separated list of schemas that the driver includes when reporting schema information to an application, if the specified schemas exist in the Drill data store. When this option is set, any schemas that are not specified in the list do not appear in the application.

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

## KrbHostFQDN or service\_host

Default Value	Data Type	Required
The host name of the drillbit.	String	No

### Description

The instance name of the service principal to use for Kerberos authentication.

## KrbKeytab or keytab

Default Value	Data Type	Required
None	String	Yes, if authenticating the connection through Kerberos without providing the service principal name.

### Description

The full path and file name of the keytab file to use for Kerberos authentication.

## KrbPrincipal or principal

Default Value	Data Type	Required
None	String	Yes, if using a complete service principal name to authenticate the connection through Kerberos.

### Description

The complete service principal name to use for Kerberos authentication, including the primary and instance (and optionally, the realm).

For example, `jsmith/simbadrill` or `jsmith/simbadrill@drill.example.com`.

## KrbRealm or realm

Default Value	Data Type	Required
Depends on your Kerberos configuration	String	No

## Description

The realm of the service principal.

If your Kerberos configuration already defines the realm of the service principal as the default realm, then you do not need to configure this property.

## KrbServiceName or service\_name

Default Value	Data Type	Required
drill	String	No

## Description

The primary name of the service principal to use for Kerberos authentication.

## LogLevel

Default Value	Data Type	Required
0	Integer	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.

Set the property to one of the following numbers:

- 0: Disable all logging.
- 1: Enable logging on the FATAL level, which logs very severe error events that will lead the driver to abort.
- 2: Enable logging on the ERROR level, which logs error events that might still allow the driver to continue running.



- 3: Enable logging on the WARNING level, which logs events that might result in an error if action is not taken.
- 4: Enable logging on the INFO level, which logs general information that describes the progress of the driver.
- 5: Enable logging on the DEBUG level, which logs detailed information that is useful for debugging the driver.
- 6: Enable logging on the TRACE level, which logs all driver activity.

When logging is enabled, the driver produces the following log files in the location specified in the `LogPath` property:

- A `DrillJDBC_driver.log` file that logs driver activity that is not specific to a connection.
- A `DrillJDBC_connection_[Number].log` file for each connection made to the database, where `[Number]` is a number that distinguishes each log file from the others. This file logs driver activity that is specific to the connection.

If the `LogPath` value is invalid, then the driver sends the logged information to the standard output stream (`System.out`).

## LogPath

Default Value	Data Type	Required
The current working directory.	String	No

### Description

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

## NettyMaxDirectMem

Default Value	Data Type	Required
None	Long	No

### Description

This property sets the `io.netty.maxDirectMemory` system property. This is only used if the `java.vendor` is IBM Corporation.

## NumNettyDirectArenas

Default Value	Data Type	Required
2	Integer	No

### Description

This property specifies the `io.netty allocator.numDirectArenas` system properties. This is only used if the `java.vendor` is IBM Corporation.

## PWD

Default Value	Data Type	Required
None	String	Yes, if you are authenticating the connection using your Drill credentials.

### Description

The password corresponding to the user name that you provided using the property [UID](#) on page 38.

## RemoveQryCatalog

Default Value	Data Type	Required
1	Integer	No

### Description

This option controls whether the driver removes the catalog from the query string.

- 1: The driver removes the catalog from the query string.
- 0: The driver does not remove the catalog from the query string.

## sasl\_encrypt

Default Value	Data Type	Required
false	String	No

### Description

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

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

### Schema

Default Value	Data Type	Required
None	String	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.

## StringColumnLength

Default Value	Data Type	Required
65535	Integer	No

### Description

The maximum length for character and binary type columns.

## TLShandshakeTimeout

Default Value	Data Type	Required
10000	Integer	No

### Description

The time in milliseconds the driver will wait to timeout if the TLS handshake fails. If the driver gets no response from the data source in this time, it will alert the user.

## TLSProtocol

Default Value	Data Type	Required
TLSv1.2	String	No

### Description

This property specifies the TLS protocol version used for SSL. Valid options include:

- TLS
- TLSV1
- TLSv1.1
- TLSv1.2

## TLSProvider

Default Value	Data Type	Required
JDK	String	No

### Description

This property specifies the TLS provider the driver uses for SSL. Has two valid values:

- JDK
- OPENSSL

## trustStorePassword

Default Value	Data Type	Required
None	String	No

### Description

The password for the host system TrustStore.

## trustStorePath

Default Value	Data Type	Required
None	String	No

### Description

This property specifies the location of the system TrustStore.

If not specified, the driver will default to the following paths:

1. `Java-home/lib/security/jssecacerts`
2. `Java-home/lib/security/cacerts`

## trustStoreType

Default Value	Data Type	Required
JKS	String	No

### Description

This property specifies the TrustStore protocol used by the host system.

- **JKS:** If the `SystemTrustStore` option is set to `true`, this value changes to `Windows-MY`.
- **PKCS12:** If the `SystemTrustStore` option is set to `true`, this value changes to `Windows-ROOT`.

## UID

Default Value	Data Type	Required
None	String	Yes, if you are authenticating the connection using your Drill credentials or a Kerberos keytab.

### Description

If the `AuthMech` (or `auth`) property is set to `Plain`, then set this property to the user name that you use to access the Drill server.

Or, if the `AuthMech` (or `auth`) property is set to `Kerberos`, then set this property to the client principal name.

## UpdateColMetadataOnSchemaChange

Default Value	Data Type	Required
None	Integer	No

### Description

This property determines whether the driver dynamically updates the column metadata of the `ResultSet` based on the data type of each row that is fetched.

- 1: The driver dynamically updates the column metadata of the `ResultSet`.
- 0: The driver does not dynamically update the column metadata of the `ResultSet`.

## Third-Party Trademarks

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

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

All other trademarks are trademarks of their respective owners.

## Third-Party Licenses

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

### **ANTLR License**

Copyright (c) 2012 Terence Parr and Sam Harwell  
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 author 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.

### **Google Protocol Buffers License**

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

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



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

Copyright 2014, Google Inc. All rights reserved.

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

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

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

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

### **JSON in Java License**

Copyright © 2015 Douglas Crockford

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

### **Reflections License**

Licensed under the WTFPL license.

## Simple Logging Façade for Java (SLF4J) License

Copyright © 2004-2015 [QOS.ch](http://QOS.ch)

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, 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.

## 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 Commons**

Copyright © 2001-2015 The Apache Software Foundation

**Apache Curator**

Copyright © 2011-2015 The Apache Software Foundation

**Apache Drill**

Copyright © 2012-2014 The Apache Software Foundation

**Apache Hadoop Common**

Copyright © 2014 The Apache Software Foundation

**Apache Logging Services**

Copyright © 1999-2012 The Apache Software Foundation

**Apache ZooKeeper**

Copyright © 2010 The Apache Software Foundation

**Google Guava**

Copyright © 2006-2015 Google, Inc.

**Jackson**

Copyright © 2009-2011 FasterXML, LLC

**Java Native Access**

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

**Javassist**

Copyright © 1999-2015 by Shigeru Chiba, All Rights Reserved.

**Joda-Time**

Copyright © 2005–2015 Joda.org. All rights reserved.

**Metrics**

Copyright © 2010-2014 Coda Hale, [Yammer.com](http://Yammer.com)

**Netty**

Copyright © 2015 The Netty Project

**Typesafe**

Copyright © 2015 Typesafe

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.