



Amazon JDBC Driver
for
Apache Hive

Installation and Configuration
Guide

September 8, 2014

Amazon Web Services, Inc.



Copyright © 2014 Amazon Web Services, 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 Amazon Web Services, Inc.

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.

Simba, the Simba logo, SimbaEngine, SimbaEngine C/S, SimbaExpress and SimbaLib are registered trademarks of Simba Technologies Inc. All other trademarks and/or servicemarks are the property of their respective owners.

The Apache Software Foundation

Licensed to the Apache Software Foundation (ASF) under one or more contributor license agreements. See the NOTICE file distributed with this work for additional information regarding copyright ownership. The ASF licenses this file to you 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.

Simple Logging Façade for Java (SLF4J)

Copyright (c) 2004-2013 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.

Contact Us

For support, check the EMR Forum at <https://forums.aws.amazon.com/forum.jspa?forumID=52> or open a support case using the AWS Support Center at <https://forums.aws.amazon.com/forum.jspa?forumID=52>

Printed in Canada

Rev.	Date	Author	Description
1.0	2014-Sep-11	OS	Initial draft

Table of Contents

Introduction	5
System Requirements	5
Amazon JDBC Driver for Apache Hive Files	5
Using the Amazon JDBC Driver for Apache Hive	6
Setting the Class Path.....	7
Amazon JDBC Driver for Apache Hive Classes	7
Hive and Hive Server Protocol Versions	7
Building the Connection URL	8
Java Sample Code.....	9
Configuring Authentication.....	11
Using No Authentication.....	11
Using Kerberos.....	11
Using User Name.....	12
Using User Name and Password	12
Using User Name and Password with Secure Sockets Layer	12
Features	13
SQL Query versus HiveQL Query	13
Data Types	13
Catalog and Schema Support.....	13
Contact Us.....	14
Appendix A: Authentication Options.....	15
Using No Authentication.....	16
Using Kerberos.....	16
Using User Name.....	16
Using User Name and Password	16
Using User Name and Password with Secure Sockets Layer	16
Appendix B: Kerberos Authentication for Windows Configuration.....	17
Appendix C: Driver Configuration Options	18

Introduction

Welcome to the Amazon JDBC Driver for Apache Hive. 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.

Amazon JDBC Driver for Apache Hive is used for direct SQL and HiveQL access to Apache Hadoop / Hive distributions, enabling Business Intelligence (BI), analytics and reporting on Hadoop / Hive-based data. The driver efficiently transforms an application's SQL query into the equivalent form in HiveQL. Hive Query Language is a subset of SQL-92. If an application is Hive-aware, then the driver is configurable to pass the query through. The driver interrogates Hive to obtain schema information to present to a SQL-based application. Queries, including joins, are translated from SQL to HiveQL. For more information about the differences between HiveQL and SQL, refer to the section *Features* on page 13.

Amazon JDBC Driver for Apache Hive complies with the JDBC 3.0 and JDBC 4.0 data standards.

The *Amazon JDBC Driver for Apache Hive Installation and Configuration Guide* is suitable for users who want to access data residing within Hive from their desktop environment. Application developers may also find the information helpful. Please refer to your application for details on connecting via JDBC.

Note: Adobe Reader is recommended for viewing PDFs. Adobe Reader is available for Windows, Mac OS X and Linux platforms from <http://get.adobe.com/reader/>.

System Requirements

To use the Amazon JDBC Driver for Apache Hive with the JDBC 3.0 API, each computer where you use the driver must meet the following system requirements:

- Java Runtime Environment (JRE) version 1.4 or 5.0

To use the Amazon JDBC Driver for Apache Hive with the JDBC 4.0 API, each computer where you use the driver must meet the following minimum system requirements:

- Java Runtime Environment (JRE) version 6.0

The Amazon JDBC Driver for Apache Hive is tested using Hive 0.10, 0.11, 0.12 and 0.13

Amazon JDBC Driver for Apache Hive Files

The Amazon JDBC Driver for Apache Hive is delivered in two ZIP archives having the name Amazon_HiveJDBC3_<version>.zip and Amazon_HiveJDBC4_<version>.zip. Each archive contains the driver supporting the JDBC API version in the archive name.

The Amazon_HiveJDBC3_<version>.zip archive contains the following file and folder structure:

- HiveJDBC3
 - hive_metastore.jar
 - hive_service.jar
 - HiveJDBC3.jar
 - libfb303-0.9.0.jar
 - libthrift-0.9.0.jar
 - log4j-1.2.14.jar
 - ql.jar
 - slf4j-api-1.5.8.jar
 - slf4j-log4j12-1.5.8.jar
 - TCLIServiceClient.jar

The Amazon_HiveJDBC4_<version>.zip archive contains the following file and folder structure:

- HiveJDBC4
 - hive_metastore.jar
 - hive_service.jar
 - HiveJDBC4.jar
 - libfb303-0.9.0.jar
 - libthrift-0.9.0.jar
 - log4j-1.2.14.jar
 - ql.jar
 - slf4j-api-1.5.8.jar
 - slf4j-log4j12-1.5.8.jar
 - TCLIServiceClient.jar

Using the Amazon JDBC Driver for Apache Hive

To access a Hive data warehouse using the Amazon JDBC Driver for Apache Hive, you need to set:

- Class path
- Driver class
- Connection URL

The Amazon JDBC Driver for Apache Hive provides read-only access to Hive data.

Setting the Class Path

The class path is the path that the Java Runtime Environment searches for classes and other resource files. For details on setting the class path, refer to <http://docs.oracle.com/javase/7/docs/technotes/tools/windows/classpath.html>.

To use the Amazon JDBC Driver for Apache Hive, you must include all JAR files included in the HIVE_JDBC.zip archive in the class path.

Amazon JDBC Driver for Apache Hive Classes

The Amazon JDBC Driver for Apache Hive makes classes having the following fully qualified class names (FQCNs) available, supporting JDBC 3.0 and JDBC 4.0 in connections to Hive Server 1 and Hive Server 2 instances:

- `com.amazon.hive.jdbc3.HS1Driver`
- `com.amazon.hive.jdbc3.HS2Driver`
- `com.amazon.hive.jdbc4.HS1Driver`
- `com.amazon.hive.jdbc4.HS2Driver`

Initialize the class appropriate for the Hive Server instance and your application as needed prior to connecting to the Hive Server.

Hive and Hive Server Protocol Versions

The following table illustrates versions of Hive Server supported by the Driver for each version of Hive.

Hive Version	Hive Server 1	Hive Server 2
0.13	Not Currently Supported	Yes
0.12	Yes	Yes
0.11	Yes	Yes
0.10	Yes	Yes

Building the Connection URL

Using the connection URL, you supply connection information to the data source that you are accessing. The connection URL for the Amazon JDBC Driver for Apache Hive takes the following form:

```
jdbc:Subprotocol://Host:Port[/Schema];Property1=Value;Property2=Value;...
```

Placeholders in the connection URL are defined as follows:

- *Subprotocol* is the value **hive** if you are connecting to a Hive Server 1 system. If you are connecting to a Hive Server 2 system, use the value **hive2**
- *Host* is the DNS or IP address of the server hosting the Hive data warehouse.
- *Port* is the port to connect to on *Host*
- *Schema* is the name of the schema/database you want to access. Specifying a schema is optional. If you do not specify a schema, then the schema named default is used.

Note: Queries can still be issued on other schemas by explicitly specifying the schema in the query. To inspect your databases and determine the appropriate database schema to use, type **show databases** at the Hive command prompt.

- You can specify one or more connection properties. For details on all available properties, see *Appendix C: Driver Configuration Options* on page 18.

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

- For any connection property with a property key that does not match any of the connection properties specified in "Appendix C: Driver Configuration Options" on page 18, the driver will attempt to apply the property as a Hive server-side property for the client session.

For example, to connect to a Hive Server 2 instance installed on the local computer by using a user name and password:

```
jdbc:hive2://localhost:10000;AuthMech=3;UID=UserName;PWD=Password
```

UserName and *Password* specify credentials for an existing user on the host running Hive Server 2.

Note: The UID and PWD properties apply when employing User Name and Password or User Name and Password with Secure Sockets Layer (SSL) authentication mechanisms. The UID property also applies when User Name authentication is enabled. For further details on properties available in the connection URL, see *Appendix C: Driver Configuration Options* on page 18.

Note: If using Hiveserver2 (Hive2) and no parameters are specified, the UID will default to "hive" and AuthMech will default to "2".

Java Sample Code

The following Java code provides an example demonstrating how to use the JDBC API to:

- Register the Amazon JDBC Driver for Apache Hive
- Establish a connection to a Hive database
- Query the database
- Parse a result set
- Handle exceptions
- Clean up to avoid memory leakage

Important: To use the Amazon JDBC Driver for Apache Hive in an application, you must include all JAR files included in the ZIP archive in the class path for your Java project.

```
// java.sql packages are required
import java.sql.*;

class AmazonJDBCHiveExample {

    // Define a string as the fully qualified class name (FQCN) of the
    // desired JDBC driver
    static String JDBCdriver = "com.amazon.hive.jdbc3.HS1Driver";
    // Define a string as the connection URL
    static String ConnectionURL = "jdbc:hive://192.168.1.1:10000";

    public static void main(String[] args) {

        Connection con = null;
        Statement stmt = null;
        ResultSet rs = null;

        // Define the SQL statement for a query
        String query = "SELECT first_name, last_name, emp_id FROM
            default.emp";

        try {

            // Register the driver using the class name
            Class.forName(JDBCdriver);

            // Establish a connection using the connection URL
            con = DriverManager.getConnection(ConnectionURL);

            // Create a Statement object for sending SQL statements
            // to the database
            stmt = con.createStatement();

            // Execute the SQL statement
            rs = stmt.executeQuery(query);
```

```
// Display a header line for output appearing in the
// Console View
System.out.printf("%20s%20s%20s\r\n", "FIRST NAME",
    "LAST NAME" , "EMPLOYEE ID");

// Step through each row in the result set returned from
// the database
while(rs.next()) {
    // Retrieve values from the row where the cursor is
    // currently positioned using column names
    String FirstName = rs.getString("first_name");
    String LastName = rs.getString("last_name");
    String EmployeeID = rs.getString("emp_id");

    // Display values in columns 20 characters wide in
    // the Console View using the Formatter
    System.out.printf("%20s%20s%20s\r\n", FirstName,
        LastName, EmployeeID);
}

} catch (SQLException se) {
    // Handle errors encountered during interaction with the
    // data source
    se.printStackTrace();
} catch (Exception e) {
    // Handle other errors
    e.printStackTrace();
} finally {
    // Perform clean up
    try {
        if (rs != null) {
            rs.close();
        }
    } catch (SQLException se1) {
        // Log this
    }

    try {
        if (stmt != null) {
            stmt.close();
        }
    } catch (SQLException se2) {
        // Log this
    }

    try {
        if (con != null) {
            con.close();
        }
    } catch (SQLException se3) {
        // Log this
    }
}
```

```
        se3.printStackTrace();
    } // End try

    } // End try

} // End main

} // End AmazonJDBCHiveExample
```

Configuring Authentication

When using the Amazon JDBC Driver for Apache Hive, you configure authentication via properties specified in the connection URL.

For details on selecting an appropriate authentication mechanism when using the Amazon JDBC Driver for Apache Hive, see *Appendix A: Authentication Options* on page 15.

For details on properties you can use in the connection URL, see *Appendix C: Driver Configuration Options* on page 18.

Using No Authentication

To use no authentication:

- Set the AuthMech property to 0

For example:

```
jdbc:hive2://localhost:10000;AuthMech=0
```

Using Kerberos

For information on operating Kerberos, refer to the documentation for your operating system.

To configure the Amazon JDBC Driver for Apache Hive to use Kerberos authentication:

1. Set the AuthMech property to 1
2. If your Kerberos setup does not define a default realm or if the realm of your Hive server is not the default, then set the appropriate realm using the KrbRealm property.
3. Set the KrbHostFQDN property to the fully qualified domain name of the Hive Server 2 host.
4. Set the KrbServiceName property to the service name of the Hive Server 2.

For example:

```
jdbc:hive2://localhost:10000;AuthMech=1;KrbRealm=EXAMPLE.COM;KrbHostFQDN=hs2.example.com;KrbServiceName=hive
```

Using User Name

To configure User Name authentication:

1. Set the AuthMech property to 2
2. Set the UID property to the appropriate user name recognized by the Hive server.

For example:

```
jdbc:hive2://localhost:10000;AuthMech=2;UID=hs2
```

Using User Name and Password

To configure User Name and Password authentication:

1. Set the AuthMech property to 3
2. Set the UID property to the appropriate user name recognized by the Hive server.
3. Set the PWD property to the password corresponding to the user name you provided in step 2.

For example:

```
jdbc:hive2://localhost:10000;AuthMech=3;UID=hs2;PWD=*****
```

Using User Name and Password with Secure Sockets Layer

To configure User Name and Password authentication using SSL:

1. Create a KeyStore containing your signed, trusted SSL certificate.
2. Set the AuthMech property to 4
3. Set the SSLKeyStore property to the full path of the KeyStore you created in step 1, including the file name.
4. Set the SSLKeyStorePwd property to the password for the KeyStore you created in step 1.
5. Set the UID property to the appropriate user name recognized by the Hive server.
6. Set the PWD property to the password corresponding to the user name you provided in step 5.

For example:

```
jdbc:hive2://localhost:10000;AuthMech=4;SSLKeyStore=C:\\Users\\bsmith\\  
\\Desktop\\keystore.jks;SSLKeyStorePwd=*****;UID=hs2;PWD=*****
```

Note: By default, the SSL certificate used by the server cannot be self-signed. Use the AllowSelfSignedCerts property to allow the SSL certificate used by the server to be self-signed. By default, the Common Name (CN) contained in a certificate signed by a Certification Authority (CA) must match the host name of the Hive server. You can allow mismatches using the CAIssuedCertNamesMismatch property. The driver always allows the Common Name contained in

self-signed certificates to mismatch the host name. See *Appendix C: Driver Configuration Options* on page 18 for more information.

Features

SQL Query versus HiveQL Query

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

Data Types

The following data types are supported:

- TINYINT
- SMALLINT
- INT
- BIGINT
- FLOAT
- DOUBLE
- DECIMAL

Note: In Hive 0.13 and later, users can specify scale and precision when creating tables using the DECIMAL data type.

- BOOLEAN
- STRING
- CHAR

Note: The CHAR data type is supported in Hive 0.13 and later.

- BINARY
- TIMESTAMP
- DATE
- VARCHAR

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

Catalog and Schema Support

The Amazon JDBC Driver for Apache Hive supports both catalogs and schemas in order to make it easy for the driver to work with various JDBC applications. Since Hive only organizes tables into schema/database, a synthetic catalog named Hive is added, under which all of the

schemas/databases are organized. The driver also maps the JDBC schema to the Hive schema/database.

Note: "CatalogSchemaSwitch" being set to 1 on the connection string will cause Hive Catalogs to be treated as Schemas in the Driver as a restriction for filtering.

Contact Us

For support, check the EMR Forum at <https://forums.aws.amazon.com/forum.jspa?forumID=52> or open a support case using the AWS Support Center at <https://forums.aws.amazon.com/forum.jspa?forumID=52>

Appendix A: Authentication Options

Hive Server 1 supports the following authentication mechanisms:

- No Authentication

Hive Server 2 supports the following authentication mechanisms:

- No Authentication
- Kerberos
- User Name
- User Name and Password
- User Name and Password with Secure Sockets Layer

To determine the authentication mechanism configured for your Hive Server 2, examine the following properties in your `hive-site.xml` file:

- **hive.server2.authentication**—The property sets the authentication mode for Hive Server 2. The following values are available:
 - **NOSASL** disables the Simple Authentication and Security Layer (SASL).
 - **KERBEROS** enables Kerberos authentication.
 - **NONE** enables plain SASL transport. **NONE** is the default value.
 - **PLAINSASL** enables user name and password based authentication using a cleartext password mechanism.
- **hive.server2.enable.doAs**—If set to the default value of **TRUE**, then Hive processes queries as the user submitting the query. If set to **FALSE**, then queries are run as the user that runs the `hiveserver2` process.

For more details on Hive Server authentication mechanisms, see the documentation for your Hadoop / Hive distribution.

Table 1 lists authentication mechanisms to configure for the Amazon JDBC Driver for Apache Hive based settings of `hive.server2.authentication` and `hive.server2.enable.doAs` properties in the `hive-site.xml` file.

<code>hive.server2.authentication</code>	<code>hive.server2.enable.doAs</code>	Driver Authentication Mechanism
NOSASL	FALSE	No Authentication
KERBEROS	TRUE or FALSE	Kerberos
NONE	TRUE or FALSE	User Name
PLAINSASL	TRUE or FALSE	User Name and Password

Table 1 Amazon JDBC Driver for Apache Hive Authentication Mechanism Configurations

More details on selecting the appropriate authentication mechanism for the Amazon JDBC Driver for Apache Hive are provided below.

For examples showing how to configure each authentication mechanism, see *Configuring Authentication* on page 11.

Using No Authentication

When `hive.server2.authentication` is set to `NOSASL`, you must configure your connection to use no authentication.

Note: Setting `hive.server2.authentication` to `NOSASL` and `hive.server2.enable.doAs` to `TRUE` creates an error. While the service starts, the configuration results in an unusable service.

Using Kerberos

When connecting to a Hive server of type Hive Server 2 and `hive.server2.authentication` is set to `KERBEROS`, then you must configure your connection to use Kerberos.

Using User Name

When connecting to a Hive server of type Hive Server 2 and `hive.server2.authentication` is set to `NONE`, then you must configure your connection to use User Name. Validation of the credentials that you include depends on `hive.server2.enable.doAs`:

- If `hive.server2.enable.doAs` is set to `TRUE`, then the User Name in the driver configuration must be an existing operating system user on the host running Hive Server 2.
- If `hive.server2.enable.doAs` is set to `FALSE`, then the User Name in the driver configuration is ignored.

If no User Name is specified in the driver configuration, then the driver defaults to using the user name `anonymous`.

Note: If you deploy Hadoop using Apache Ambari, then by default the authentication method is User Name.

Using User Name and Password

When connecting to a Hive server of type Hive Server 2 configured to use plain SASL authentication, then you must configure your connection to include a User Name and Password.

Using User Name and Password with Secure Sockets Layer

When connecting to a Hive server of type Hive Server 2 configured to use plain SASL authentication and configured to use SSL, then you must configure your connection to include a User Name and Password.

Appendix B: Kerberos Authentication for Windows Configuration

Note: If you do not use Kerberos authentication, then you do not need to configure Kerberos authentication.

On the Windows platform, using Kerberos authentication as part of an Active Directory implementation is recommended. To implement Active Directory, contact your system administrator.

When using Kerberos for enhanced authentication in an Active Directory implementation, users are granted a Ticket to Get Tickets (TGT) automatically when logging on to the network. The Kerberos Key Distribution Center (KDC) is installed as part of the domain controller and performs two service functions: the Authentication Service (AS) and the Ticket-Granting Service (TGS). The KDC has access to Active Directory user account information. After the Authentication Service verifies the user, the Ticket-Granting Service provides the Ticket to Get Tickets.

In an Active Directory implementation using Kerberos, Kerberos uses the Active Directory cache to store credentials.

The Amazon JDBC Driver for Apache Hive requests a service ticket in the process of establishing a connection to the Hive database.

For more details on Kerberos authentication and support for Kerberos authentication in the Amazon JDBC Driver for Apache Hive, see the *SimbaEngine SDK SimbaClientServer User's Guide*, available at <http://www.simba.com/wp-content/uploads/2013/05/SimbaClientServer-Users-Guide1.pdf>.

Appendix C: Driver Configuration Options

Table 2 lists and describes the properties available to configure the behaviour of the Amazon JDBC Driver for Apache Hive.

Note: You can set configuration properties using the connection URL. For details on the connection URL, see *Building the Connection URL* on page 8.

Property	Default Value	Description
AllowSelfSignedCerts	0	When set to the default value of 0, the SSL certificate used by the server cannot be self-signed. When set to 1, the SSL certificate used by the server can be self-signed. Note: The setting is only applicable to the User Name and Password with SSL authentication mechanism, and is ignored by other authentication mechanisms. (Optional)
AuthMech	0	The authentication mechanism to use. Set the value to 0 for no authentication, 1 for Kerberos, 2 for User Name, 3 for User Name and Password, 4 for User Name and Password with SSL. (Optional)
CAIssuedCertNamesMismatch	0	The property controls whether to allow the CA issued SSL certificate name to not match the host name of the Hive server. Mismatches are not allowed when the property is set to 0. To allow mismatches, set the property to 1. Note: The setting is only applicable to the User Name and Password with SSL authentication mechanism, and is ignored by other authentication mechanisms. (Optional)
DecimalColumnScale	10	The maximum number of digits to the right of the decimal point for numeric data types (Optional)

Property	Default Value	Description
DefaultStringColumnLength	255	The maximum data length for string columns Note: Hive does not provide the maximum data length for String columns in the columns metadata. This option allows you to tune the maximum data length for String columns. The range of DefaultStringColumnLength is 0 to 32,767. (Optional)
DelegationUID	N/A	Used to delegate all operations against Hive to a user that is different than the authenticated user for the connection. Note: The setting is only applicable when connecting to a Hive Server 2 that supports the delegation feature. Otherwise, the setting has no effect. (Optional)
KrbHostFQDN		The fully qualified domain name of the Hive Server 2 host used. (Required if AuthMech is Kerberos)
KrbRealm	Depends on Kerberos configuration	If there is no default realm configured, or if the realm of the Hive Server 2 host is different from the default realm for your Kerberos setup, then define the realm of the Hive Server 2 host using this option. (Optional)
KrbServiceName		The Kerberos service principal name of the Hive Server 2. (Required if AuthMech is Kerberos)
PWD		The password set up for User Name and Password authentication mechanisms. (Required if AuthMech is User Name and Password, or User Name and Password with SSL)
RowsFetchedPerBlock	10000	The maximum number of rows that a query returns at a time. Any positive 32-bit integer is a valid value but testing has shown that performance gains are marginal beyond the default value of 10000 rows.

Property	Default Value	Description
SocketTimeout	0	The number of seconds after which Hive closes the connection with the client application if the connection is idle. The default value of 0 indicates that an idle connection is not closed.
SSLKeyStore	N/A	The full path, including the file name, of the Java KeyStore containing an SSL certificate to use during authentication. See also the SSLKeyStorePwd property. (Required if AuthMech is User Name and Password with SSL)
SSLKeyStorePwd	N/A	The password required to access the Java KeyStore specified using the SSLKeyStore property (Required if AuthMech is User Name and Password with SSL)
UID		The user name of an existing user on the host running Hive Server 2 Important: You must set the hive.server2.authentication property in the hive-site.xml file for the Hive Server 2 to NONE OR You must set up the Hive Server 2 credential when using the User Name and Password authentication mechanism. (Required if AuthMech is User Name, User Name and Password, or User Name and Password with SSL)
UseNativeQuery	0	Enabling the UseNativeQuery option using a value of 1 disables the SQL Connector feature. The SQL Connector feature has been added to the driver to apply transformations to the queries emitted by an application to convert them into an equivalent form in HiveQL. If the application is Hive aware and already emits HiveQL, then turning off the SQL Connector feature avoids the extra overhead of query transformation. (Optional)

Property	Default Value	Description
CatalogSchemaSwitch	0	Enabling the CatalogSchemaSwitch option using a value of 1 on the connection string will cause Impala Catalogs to be treated as Schemas in the Driver as a restriction for filtering.
PreparedMetaLimitZero	0	Enabling PreparedMetaLimitZero will cause the PreparedStatement.getMetadata() call to request metadata from the server with "LIMIT 0".

Table 2 Amazon JDBC Driver for Apache Hive Configuration Options