

Release Notes

Cloudera JDBC Driver for Apache Hive 2.5.20

Released 2017-12-22

These release notes provide details of enhancements, features, and known issues in Cloudera JDBC Driver for Apache Hive 2.5.20, as well as the version history.

Enhancements & New Features

Hive 2.1 now supported

The driver now supports Apache Hive versions 0.11 through 2.1.

Hadoop Delegation Tokens now supported

The driver now supports authentication using Hadoop Delegation Tokens. For more information, see the Installation and Configuration Guide.

[00097507] SubjectAlternativeNames host name matching now supported

The driver now supports SubjectAlternativeNames host name matching for SSL certificates..

AllowAllHostNames option deprecated

The driver configuration option AllowAllHostNames, which could be used to disable TLS/SSL host name verification, has been deprecated. Instead, use the option CAIssuedCertsMismatch. For more information about configuring TLS/SSL options in the driver, see the Install and Configuration Guide.

Support for executeAsync and getYarnATSGuid

Users can now use the IHadoopStatement interface to execute a SQL statement asynchronously, and retrieve the Yarn ATS GUID for an executing query.

Updated Kerberos support

The driver now supports using Kerberos authentication with http transportMode.

Resolved Issues

The following issues have been resolved in Cloudera JDBC Driver for Apache Hive 2.5.20.

- [00097540] When next() is called for a second time on a ResultSet that does not have any more rows, an incorrect exception is thrown.

This issue has been resolved. The driver now returns False.

- If a query contains both a DISTINCT and an ORDER BY clause, the ORDER BY clause is ignored.
- [00095678] If there are over 400 concurrent connections, an Out Of Memory error may occur.
- [00095678] If there are over 400 concurrent connections, the log files may be deleted.
- [00095678] Exceptions are not logged or thrown during catalog function calls.
- [00098137] AllowSelfSignedCerts does not disable SSL.

This issue has been resolved. Now, when the AllowSelfSignedCerts property is set to 1, SSL verification is disabled. The driver does not verify the server certificate against the trust store, and does not verify if the server's host name matches the CN or Subject Alternative Names in the server certificate.

This behavior is consistent with Cloudera JDBC Driver for Apache Hive 2.5.18 and earlier.

- Translations are incorrect when queries have both JOIN and GROUP BY clauses.
- Joins involving more than three tables fail.
- Scalar function column names without aliases are returned with extra information appended.

Version History

Version 2.5.19

Released 2017-04-06

Enhancements & New Features

AsyncExecPollInterval configuration option

You can now configure how often the driver polls the server for query execution status.

Resolved Issues

The following issues have been resolved in Cloudera JDBC Driver for Apache Hive 2.5.19.

- In some configurations users would receive a NullPointerException system message when attempting to connect to the server.
- The driver sometimes returns fewer rows than expected if the JVM is running low on memory, even though no errors are reported.

Version 2.5.18

Released 2016-11-30

Enhancements & New Features

[Salesforce: 00090277] Support added for specifying how the driver obtains the Kerberos Subject

When using Kerberos authentication, you can now configure the `KrbAuthType` connection property to specify how the driver obtains the Kerberos Subject. For more information, see the *Installation and Configuration Guide*.

Resolved Issues

The following issues were resolved in Cloudera JDBC Driver for Apache Hive 2.5.18.

- In some cases, when the heap size is restricted, the driver retrieves fewer rows than expected but does not return any exceptions or errors.

This issue has been resolved. The driver now retrieves the correct number of rows when the heap size is restricted.

Version 2.5.17

Released 2016-09-29

Enhancements & New Features

Support added for connecting via the HTTP transport protocol

You can now use the new `transportMode` and `httpPath` connection properties to configure the driver to connect through the HTTP transport mode. The default `UID` and `PWD` values in the driver have been updated for improved compatibility with these new connection properties, and the default `AuthMech` value now depends on the `transportMode` setting. For more information, see the *Installation and Configuration Guide*.

Support added for configuring host name verification for SSL connections

You can now use the `AllowAllHostNames` connection property to specify whether host name verification is enabled for SSL connections.

Updated handling of socket timeouts

The driver now uses 30 as the default value for the `SocketTimeout` property. If a socket read operation takes longer than 30 seconds to complete, then the driver closes the connection. Also, error messages for socket timeouts have been improved for clarity.

Resolved Issues

The following issues were resolved in Cloudera JDBC Driver for Apache Hive 2.5.17.

- The driver throws a null pointer exception upon calling `Driver.getPropertyInfo()`.
- The buffer length and consequently char octet length defaults to 0 for CHAR and VARCHAR types.
- [13952] Query translation drops brackets around OR clauses.
- Driver returns error when casting to `TIMESTAMP` and `DATE` in SQL statement.
- Inserting a false Boolean value as 0 inserts `TRUE`.

- When the `setQueryTimeout()` method is called and the query processing time exceeds the query timeout value, the driver returns a socket timeout error.

The driver now returns the correct error for query timeouts (`SqlTimeoutException`).

Version 2.5.16

Released 2015-11-12

Resolved Issues

The following issues were resolved in Cloudera JDBC Driver for Apache Hive 2.5.16.

- `DatabaseMetaData.getColumns()` does not return complex type fields through the JDBC driver.

The driver now partially supports Array, Struct, Map, and Union SQL types. These types are mapped to Varchar.

- Hive Server 1 and Hive Server 2 drivers cannot run at the same time in the same JVM.

This issue has been resolved. Before, the `HS1Driver` and `HS2Driver` could not coexist in the same JVM; loading both drivers at the same time caused a Hive Server 2 connection failure. You can now run both drivers at the same time in the same JVM.

Version 2.5.15

Released 2015-07-20

Resolved Issues

The following issues were resolved in Cloudera JDBC Driver for Apache Hive 2.5.15.

- In some cases, the `getTables` catalog function returns table types incorrectly.
- The `REMARKS` table metadata column does not contain the correct data.

This issue has been resolved. Previously, the `REMARKS` table metadata column was hard-coded, so the data in the column could not be changed. You can now work with the column as expected.

- In some cases, driver performance is slow when using catalog functions to retrieve metadata.

This issue has been resolved. Previously, the way that the driver retrieved metadata involved retrieving and processing more data than was necessary. Now, the behavior in the driver is optimized and performance has improved significantly.

- In CDH 4.x or earlier, in some cases nested calls cause the driver to return a "Read a negative frame size" error.

Version 2.5.14

Released 2015-07-02

Enhancements & New Features

Support added for case-insensitive treatment of catalog names, schema names, table names, and column names in catalog function calls

The driver is now able to work with catalogs, schemas, tables, and columns as expected regardless of whether the names are spelled with upper- or lower-case characters.

Optimized metadata retrieval

The driver is now designed to push catalog function restrictions down to the server for processing when possible, which improves driver performance during metadata retrieval.

Resolved Issues

The following issues were resolved in Cloudera JDBC Driver for Apache Hive 2.5.14.

- Driver returns duplicate data for the `getColumns` JDBC API when the table restriction contains a wildcard.
- In some cases, the driver fails to retrieve columns metadata from tables that contain partitioned columns.
- Driver sends a separate `DESC <table>` query for each `DECIMAL`, `VARCHAR`, and `CHAR` column in the queried table.

This issue has been resolved. Now, the driver sends only one `DESC <table>` query for the queried table.

- The `RowsFetchedPerBlock` setting limits the amount of schemas, tables, and columns metadata retrieved from Hive Server 2 instances.

This issue has been resolved. The `RowsFetchedPerBlock` setting now works as expected, and only limits the maximum number of rows retrieved per fetch call.

Version 2.5.13

Released 2015-06-16

Resolved Issues

The following issues were resolved in Cloudera JDBC Driver for Apache Hive 2.5.13.

- Driver does not support the `LENGTH` scalar function for non-character columns.

This issue has been resolved. You can now call the `LENGTH` function on any data type that can be converted to `STRING`.

- Driver cannot execute "SET <key>=<value>" statements using `executeUpdate()`.
- Driver cannot execute `SET` statements with leading or trailing spaces.

This problem has been resolved. Before, the driver returned an error if a `SET` statement started or ended with spaces. The driver is now able to parse the statement and execute it successfully.

- TCP connections do not close when `connection.close()` is called.

This issue has been resolved. Before, TCP connections did not close when `connection.close()` was called, and only closed when the client application closed. Now, TCP connections close immediately after `connection.close()` is called.

- When handling multiple connections at the same time, in some cases the driver will retrieve the wrong data.

This issue has been resolved. Driver support for multithreading has been fixed and the driver will now return the correct data when there are multiple connections.

Version 2.5.12

Released 2015-05-06

Resolved Issues

The following issues were resolved in Cloudera JDBC Driver for Apache Hive 2.5.12.

- The `getColumns()` function returns an empty result set if it uses both schema and table restrictions, and the table restriction contains an escaped wildcard character.

This issue has been resolved. The `getColumns()` function now returns the correct result set.

- Query translation does not use the database context specified by a `USE <database> query`.

This issue has been resolved. Now, when you change the database context by executing a `USE <database> query`, any subsequent query translation will use that database context.

Version 2.5.11

Released 2015-04-16

Enhancements & New Features

SSL now configured separately from authentication

Before, you would enable SSL in the driver by setting the authentication mechanism to "User Name and Password with Secure Sockets Layer" (`AuthMech=4`) or "No Authentication with Secure Sockets Layer" (`AuthMech=5`). You can now use the new `SSL` property to enable or disable SSL connections, and use the `AuthMech` property solely to configure authentication.

Resolved Issues

The following issues were resolved in Cloudera JDBC Driver for Apache Hive 2.5.11.

- Driver does not report an error when a MapReduce job fails on the server side during query execution.

This issue has been resolved. The driver now reports an error if a MapReduce job fails.

Version 2.5.10

Released 2015-02-25

Enhancements & New Features

Support added for write-back when connected to Hive 0.14 or later

You can now execute INSERT, UPDATE, and DELETE statements when connected to Apache Hive 0.14 or later.

Support added for dynamic service discovery with Apache ZooKeeper

You can now connect to Hive servers that are registered against a ZooKeeper service by connecting to the ZooKeeper service.

Version 2.5.9

Released 2015-02-18

Resolved Issues

The following issues were resolved in Cloudera JDBC Driver for Apache Hive 2.5.9.

- Using Java services to load the driver classes may cause a Hive Server 2 connection failure.

When working in Java 6 or later, you have the option of loading the required driver classes using Java services instead of using `Class.forName()`. Java services sometimes loads `HS1Driver` before loading `HS2Driver`, which prevents `HS2Driver` from operating correctly. As a short-term workaround for this issue, this release of the driver prevents Java services from performing the loading.

Version 2.5.8

Released 2015-01-27

Resolved Issues

The following issues were resolved in Cloudera JDBC Driver for Apache Hive 2.5.8.

- Driver misinterprets newer version of Hive.

This issue has been resolved. Before, the Hive version was being misinterpreted, which caused the driver to establish connections using V3 of the protocol and to apply Hive 0.10 translations to queries. The driver now correctly exhibits updated Hive behavior.

- When using Kerberos authentication, the driver returns an error message.

This issue has been resolved. Before, the driver did not retrieve the JAAS configuration file from the **KRB5_CONFIG** system property as expected, and attempting to authenticate through Kerberos would return the following error message: "unable to find principal name for authentication". The driver now correctly handles Kerberos authentication.

Version 2.5.6

Released 2014-12-04

Enhancements & New Features

JDBC 4.1 now supported

The driver now supports JDBC 4.1. To use the driver with JDBC 4.1, use the HiveJDBC41_2.5.14 package.

Support added for Kerberos authentication through AccessControlContext

The driver now provides an alternative method for obtaining Kerberos tickets. Instead of obtaining a ticket-granting ticket (TGT) from the ticket cache, the driver can now check whether there are any Subjects associated with the AccessControlContext and whether those Subjects have a TGT. If so, the driver can use the TGT from the Subject.

Support added for Kerberos authentication on IBM Java 1.6

You can now use Kerberos authentication when running IBM Java 1.6.

Support added for the Connection.isValid() and Connection.getClientInfo() methods for JDBC 4 and JDBC 4.1

The driver now supports the Connection.isValid() and Connection.getClientInfo() methods.

Support added for connection pooling

The driver now supports the JDBC ConnectionPoolDataSource interface via the following classes:

- `com.cloudera.hive.jdbc3.DataSource`
- `com.cloudera.hive.jdbc4.DataSource`
- `com.cloudera.hive.jdbc41.DataSource`

SSLTrustStore and SSLTrustStorePwd configuration parameters implemented

The `SSLTrustStore` and `SSLTrustStorePwd` parameters are optional parameters that you can use in the connection string. Use these parameters to configure the driver to use a specific TrustStore when connecting through SSL. If these parameters are not set, then the driver uses the default TrustStore located in `jre\lib\security\cacerts`.

Resolved Issues

The following issues were resolved in Cloudera JDBC Driver for Apache Hive 2.5.6.

- Query cancellation requests fail when schema is not defined in the connection string.

Fixed an issue where the query cancellation from version 2.5.5 does not work properly if the schema is not defined in the connection string.

- Queries that include comments fail to execute.

This issue has been resolved. Now, queries that contain comments can be processed.

Version 2.5.5

Released 2014-10-21

Enhancements & New Features

Support for Hive 0.13.0 data types added to the Hive Server 1 client

Support for the `CHAR(n)` and `DECIMAL(precision, scale)` data types have been added to the Hive Server 1 client. The client now fully supports all the data types in Hive 0.13.0.

Support added for PreparedStatements with parameters

Previously, the driver did not support the use of `PreparedStatements` with parameters because the Thrift API did not support this functionality. You can now use parameters in `PreparedStatements`. However, the workaround for supporting this feature is not a typical method for doing `PreparedStatements`. The driver

replaces the question mark (?) in the query with the actual parameter value and executes the query during the execution stage. It is not recommended that you use `prepareStatement.getParameterMetadata` or `prepareStatement.getResultSetMetadata` before `prepareStatement.executeQuery`, because the `PreparedStatement` does not return the expected values.

Direct integration with Kerberos Key Distribution Center added

You can now configure the driver to get a ticket from your Key Distribution Center directly. To do this, in the JVM environment, configure a JAAS configuration file that directs the driver to use a generated keytab file as the credentials.

Resolved Issues

The following issues were resolved in Cloudera JDBC Driver for Apache Hive 2.5.5.

- Column metadata fails to be retrieved from Hive 0.9 servers.

Fixed an issue where column metadata failed to be retrieved from servers running Hive 0.9.

- Query cancellation request fails to stop queries.

Fixed an issue where the driver ignores query cancellation requests from the application and continues to execute the query. Now, when the application sends a query cancellation request, the query stops executing and the server stops processing the query.

- REPLACE commands that include a meta-character as a parameter are not translated correctly.

Fixed an issue where queries that use the REPLACE command "replace(expression, pattern, replacement)" with a meta-character as a parameter are not translated correctly. The solution for the issue enables escaping of all meta-characters, including the following:

. ^ \$ * + ? { } [] | ()

Version 2.5.4

Released 2014-09-08

Resolved Issues

The following issues were resolved in Cloudera JDBC Driver for Apache Hive 2.5.4.

- Certain known connection string properties could be incorrectly "SET" on the server.

In order to support server property configuration in Hive at startup, any unrecognized properties specified in the connection string would cause a "SET" command to be executed automatically on the server side. When this happened, certain combinations of known properties stopped working correctly. This issue has been resolved.

Version 2.5.3

Released 2014-09-04

Enhancements & New Features

statement.setMaxRows function implemented

The statement.setMaxRows function has been implemented. This function allows you to limit the number of rows fetched from the result set.

Resolved Issues

The following issues were resolved in Cloudera JDBC Driver for Apache Hive 2.5.3.

- Wrong class path in MANIFEST file

Fixed an issue where the MANIFEST file contained an incorrect class path caused by branding.

- Multiple TCP connections created by ZooKeeper from the server side

Fixed an issue where generating ROWCOUNT set results caused ZooKeeper to open new TCP connections until ZooKeeper reached the maximum number of connections that it can handle. This issue was fixed by properly closing the operation that is opened on the client when a ROWCOUNT set is generated.

Version 2.5.2

Released 2014-08-28

Enhancements & New Features

Hive variables supported in connection string

The driver now supports the use of Hive variables in the connection string.

DatabaseMetadata added

DatabaseMetadata information has been added for DBMS name and DBMS version.

Resolved Issues

The following issues were resolved in Cloudera JDBC Driver for Apache Hive 2.5.2.

- USE statement causing memory leaks

Fixed an issue where using the USE statement caused server-side memory leaks to occur.

- SET statement causing errors

Fixed an issue where using the SET statement caused errors to occur.

Version 2.5.1

Released 2014-07-15

Enhancements & New Features

LIMIT ZERO query feature added for prepareStatement

This feature supports queries with LIMIT 0 attached during the query prepare stage, improving the performance of the prepareStatement process.

Version 2.5.0

Released 2014-06-15

Version 2.5.0 was the initial release of the Cloudera JDBC Driver for Apache Hive.

Contact Us

If you are having difficulties using the driver, our [Community Forum](#) may have your solution. In addition to providing user to user support, our forums are a great place to share your questions, comments, and feature requests with us.

If you are a Subscription customer you may also use the [Cloudera Support Portal](#) to search the Knowledge Base or file a Case.

Important: To help us assist you, prior to contacting Cloudera Support please prepare a detailed summary of the client and server environment including operating system version, patch level, and configuration.