Release Notes

Cloudera ODBC Driver for Apache Hive 2.5.25

The release notes provide details of enhancements and features in Cloudera ODBC Driver for Apache Hive 2.5.25, as well as the version history.

Enhancements & New Features

Updated Kerberos configuration

The Host FQDN and Service Name configuration options (the **KrbHostFQDN** and **KrbServiceName** connection properties) are now optional. Also, the driver now automatically canonicalizes the server SPN when using Active Directory Kerberos.

Resolved Issues

The following issues have been resolved in Cloudera ODBC Driver for Apache Hive 2.5.25.

- AD Kerberos authentication may use the wrong user when enabling constrained delegation in Tableau server on Windows.
- SQLGetInfo(SQL_KEYWORDS) returns an empty string.
  The driver now returns a comma-separated list of keywords natively defined by the data source except for the ones that are also ODBC reserved words.

Version History

Version 2.5.24

Enhancements & New Features

Updated configuration option

The name of the Delegation User ID case configuration option has changed from **DelegationUIDCase** to **DelegationUserIDCase**.

Resolved Issues

The following issues were resolved in Cloudera ODBC Driver for Apache Hive 2.5.24.

- The statement "set hive.support.quoted.identifiers=column" is executed incorrectly.
- Driver reports only read operations are supported in the driver log.
- The INIFINITY value for FLOAT columns are not correctly retrieved.

- Error/status code returned for query timeout error is incorrect.
- NaN value retrieved for FLOAT columns is incorrect.
- The operation handle for metadata API calls is not closed correctly.

Version 2.5.23

Enhancements & New Features

The following are highlights of the new features and functionalities that have been added to Cloudera ODBC Driver for Apache Hive 2.5.25.

[00093344] Convert letter case of delegated user names

You can now configure the driver to change the Delegation UID (or `DelegationUID`) value to all upper-case or all lower-case. To do this, set the `DelegationUIDCase` connection property.

Version 2.5.22

Enhancements & New Features

The following are highlights of the new features and functionalities that have been added to Cloudera ODBC Driver for Apache Hive 2.5.22.

Updated third-party library dependencies

The driver has been updated to incorporate newer versions of the third-party library dependencies.

Resolved Issues

The following issues have been resolved in Cloudera ODBC Driver for Apache Hive 2.5.22.

When attempting to use the Windows trust store on Windows Server 2016, an access violation exception occurs

This issue has been resolved.

Error message contains typo, misspelling "URI" as "UIR"

This issue has been resolved.

Driver converts COALESCE function to less efficient CASE statement

This issue has been resolved.

In the Cloudera Hive ODBC Driver Configuration tool for DSN-less connections, the Enable Auto Reconnect option is disabled by default

This option is now enabled by default, which is the expected default setting.

In some cases, when upgrading the macOS driver from an earlier version to a later one, the installation process fails to update the odbcinst.ini configuration file

Before, upgrading the driver sometimes required you to make changes to the `~/Library/ODBC/odbcinst.ini` file manually. The driver installation process has been updated so that manual changes are no longer necessary.

Segmentation fault in Driver Manager detection on Linux

This issue has been resolved.

If you use an alias to refer to the Hive server host name while authenticating through Kerberos, the driver fails to connect

This issue has been resolved.

Version 2.5.21

Enhancements & New Features

The following are highlights of the new features and functionalities that have been added to Cloudera ODBC Driver for Apache Hive 2.5.21.

Support for the Windows trust store

You now have the option to use the CA certificates in the Windows trust store for server verification when using SSL.

Auto Reconnect

You can now configure the driver to automatically attempt reconnection to the Hive server if communications are lost.

Resolved Issues

The following issues have been resolved in Cloudera ODBC Driver for Apache Hive 2.5.21.

Driver fails to connect to the server when using TLS 1.2.

This issue has been resolved.

[109843 (SalesForce.com KIR 00089551)] When executing a query that contains a CASE statement, the driver replaces the greater than sign (>) with an equal sign (=).

This issue has been resolved.

Driver does not correctly translate queries that contain the IS NOT NULL operator.

This issue has been resolved.

Version 2.5.20

Enhancements & New Features

The following are highlights of the new features and functionalities that have been added to Cloudera ODBC Driver for Apache Hive 2.5.20.

Delegate Kerberos credentials

You can now have the driver forward your Kerberos user credentials to the server to simplify the authentication process.

Optimized Fast SQLPrepare behavior

The Fast SQLPrepare driver configuration option (the FastSQLPrepare key) is now disabled for non-SELECT queries. This ensures that the driver retrieves the necessary result set metadata at prepare time.

Resolved Issues

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

Unicode characters in parameter values causing errors

This issue has been resolved.

Returning errors for some queries with dates in them

This issue has been resolved.

Unable to create new tables when Unicode character types option set

This issue has been resolved.

Version 2.5.19

Enhancements & New Features

The following are highlights of the new features and functionalities that have been added to Cloudera ODBC Driver for Apache Hive 2.5.19.

Support added for Red Hat Enterprise Linux (RHEL) 7 and CentOS 7

You can now install and run the Linux version of the driver on machines that run RHEL 7 or CentOS 7.

Updated authentication support

For consistency with Microsoft Azure HDInsight updates, the Windows Azure HDInsight Emulator authentication mechanism (`AuthMech=5`) is now deprecated. The driver no longer supports connections to Hive server instances that run on Windows Azure HDInsight Emulator.

Resolved Issues

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

When attempting to connect to the server using a non-Windows version of the driver with Service Discovery Mode enabled, the client stops working

This issue has been resolved.

When executing a parameterized INSERT statement on a DATE, DECIMAL, or TIMESTAMP column, an error occurs

This issue has been resolved.

When executing a query that contains a large number of filters using OR operators, an error occurs

This issue has been resolved.

When executing a query that uses regular expressions in the fields, an error occurs

This issue has been resolved. Before, this problem occurred when querying Hive servers that are configured to not allow quoted identifiers (the `hive.support.quoted.identifiers` parameter is set to none), because these servers cannot parse the backquotes (`) used for regular expressions.

Driver truncates passwords that exceed the maximum character length of the Password field in the Cloudera Hive ODBC Driver DSN Setup dialog box

The maximum character length of the Password field has been increased to 5000 characters.

In some cases, the driver stops working intermittently

This issue has been resolved. Before, this problem occurred due to improper timing and sequence of operations between the execution context and the fetch thread.

In some cases, after being used in Tableau, the driver caches sessions and then tries to reuse them the next time Tableau is started

This issue has been resolved.

When executing a parameterized INSERT statement on a BINARY column, an error occurs

This issue has been resolved.

Version 2.5.18

Enhancements & New Features

The following are highlights of the new features and functionalities that have been added to Cloudera ODBC Driver for Apache Hive 2.5.18.

Updated handling of result set metadata

When connected to Hive 0.14 or later, the driver now reports result set columns as being updatable, improving compatibility with the DotNet Odbc library and OdbcDataAdapter.

Updated default socket timeout threshold

The default value of the Socket Timeout option (the SocketTimeout key) has been changed from 30 to 60.

Resolved Issues

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

When attempting to connect to the server using a non-Windows version of the driver with Service Discovery Mode enabled, the client stops working

This issue has been resolved.

[12726] When executing a parameterized INSERT statement on a DATE, DECIMAL, or TIMESTAMP column, an error occurs

This issue has been resolved.

When executing a query that contains a large number of filters using OR operators, an error occurs

This issue has been resolved.

Version 2.5.17

Enhancements & New Features

The following are highlights of the new features and functionalities that have been added to Cloudera ODBC Driver for Apache Hive 2.5.17.

Debian 6 now supported

You can now install and use the driver on computers that run Debian 6.

Data encryption and integrity support implemented for the SASL SSPI plugin

The driver now supports data encryption and integrity support in the SASL SSPI plugin, providing improved support for connections that are authenticated through Kerberos.

Optimized retrieval of result set schemas

When connected to Hive 0.14 or later, the driver uses the LIMIT 0 query to improve performance for retrieving result set schemas.

Support added for ODBC 3.80

The driver now supports ODBC 3.80. Previously, the driver supported ODBC 3.52.

Improved implementation of SQLCancel

The driver now implements industry-standard behavior for when SQLCancel is called while query results are being fetched.

Upgraded OpenSSL library

The driver now uses OpenSSL 1.0.1l. Previously, the driver used OpenSSL 1.0.0q.

Upgraded ICU library

The driver now uses ICU 53.1. Previously, the driver used ICU 3.8.1.

Implemented driver configuration option for automatically opening a new session against Hive Server 2 when the existing session becomes invalid

You can now use the Invalid Session Auto Recover option (the InvalidSessionAutoRecover key) to configure the driver to automatically open a new session when it detects that the current session is no longer valid. This feature is available only for connections to Hive Server 2 instances.

Automatically configure settings in cloudera.hiveodbc.ini file

On non-Windows platforms, the DriverManagerEncoding and ODBCInst settings are now configured automatically based on the driver manager (iODBC or unixODBC) that is used.

Updated driver logging options

You can now specify the maximum size of each log file by using the "Max File Size" option (the LogFileSize key), and specify the maximum number of log files to keep by using the "Max Number Files" option (the LogFileCount key). Also, in the Logging Options dialog box for the Windows driver, you now have the option of browsing for the folder where you want to save log files instead of typing the path.

Resolved Issues

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

Driver returns incorrect ODBC error code for authentication-related errors

Before, when an authentication-related error occurred, the driver returned error code HY000, which indicates a general error. The driver now correctly returns error code 28000 instead.

When using the Windows version of the driver, a SideBySide error occurs due to a conflict in the OpenSSL and libcurl libraries being loaded

This issue has been resolved.

When attempting to execute a non-parameterized write-back statement multiple times, the driver only executes the statement against Hive once

This issue has been resolved.

Driver returns incorrect results for SET <property name> queries

Before, the driver returned row counts after executing the SET <property name> query. The driver now correctly returns a result set containing the value of the specified property.

When connected to Hive 0.13 or earlier, the driver cannot insert result sets from queries into tables

Before, when preparing or executing an "INSERT INTO <table> SELECT ..." statement, the driver would return an error. The driver is now able to handle these statements correctly and insert result sets from queries into tables.

Driver does not correctly translate "INSERT OVERWRITE TABLE <table> PARTITION(...) SELECT ..." statements correctly

This issue has been resolved.

DSN configurations cannot be saved when you are working as a non-administrator user

This issue has been resolved.

When query execution results in a MapReduce error, the driver closes the current connection and opens a new one

This issue has been resolved.

Driver terminates unexpectedly when preparing a parameterized INSERT statement

This issue has been resolved.

Driver terminates unexpectedly when the application unloads the driver without closing all connections

This issue has been resolved.

Version 2.5.16

Enhancements & New Features

The following are highlights of the new features and functionalities that have been added to Cloudera ODBC Driver for Apache Hive 2.5.16.

Driver now supports the use of "_HOST" as a special value for the Host FQDN option

When configuring Kerberos authentication, you can now set the value of the Host FQDN option (the KrbHostFQDN key) to **_HOST** in order to use the Hive server host name as the fully qualified domain name for Kerberos authentication.

Exposed a custom connection attribute for returning the Simba driver version number

You can now get the Simba driver version number using the custom connection attribute 30000.

Support added for using the USE <database> command to set the current schema context in the driver

You can now use the USE <database> command to set the current schema context in a connection. Previously, the current schema context was set using the Database option (the Schema connection attribute) and could not be changed otherwise.

Interval data types now supported

The driver now supports the Hive data types for time intervals in query expressions and predicates, which are available as of Hive 1.2.0.

Dialog box added for configuring logging in the Windows driver

On Windows, the driver now provides a dialog box (accessible from the DSN Setup dialog box and the Driver Configuration Tool) for configuring driver logging options. For detailed information about logging, see the Installation and Configuration Guide.

Support added for configuring socket timeouts

You can now use the Socket Timeout option (the SocketTimeout key) to specify the number of seconds that an operation can remain idle before it is closed. This setting applies only when asynchronous query execution is being used against HS2.

Implemented driver configuration option for using the SSPI plugin to handle Kerberos authentication

By default, the Windows driver uses MIT Kerberos to handle Kerberos authentication, and only uses the SSPI plugin if the gssapi library is not available. You can now configure the Windows driver to use the SSPI plugin by default by setting the Use Only SSPI option (the UseOnlySSPI key).

Updated the identifier quote character reported by the driver

The driver now reports back quotes (`) instead of double quotes (") as the character being used to quote identifiers.

Support added for translating INSERT INTO/OVERWRITE SELECT statement syntax

The driver now supports query translation for INSERT INTO/OVERWRITE SELECT statements.

Implemented improvements to return better connection-related error messages

Previously, the driver returned generic error messages for some of the connection-related errors. The error messages have been improved to provide more meaningful information about the errors.

Support added for generating the proper DATE and TIMESTAMP literals

When connecting to Hive 1.2 or later, the driver will generate DATE and TIMESTAMP literals that are supported by Hive 1.2 and later.

Disabled top-level UNION query translation workaround when connecting to Hive 0.13 or later

Hive 0.13 added support for top-level UNION. The driver now disables the top-level UNION query translation workaround when connecting to Hive 0.13 or later.

Disable set UNION query translation workaround when connecting to Hive 1.2 or later

Hive 1.2 added support for set UNION. The driver now disables the set UNION query translation workaround when connecting to Hive 1.2 or later.

Resolved Issues

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

Driver cannot connect to Hive Server 1 instances that use versions of Hive before 0.8

This issue has been resolved.

On Windows, the ZooKeeper library stops working if the host and port of the servers are not specified correctly

This issue has been resolved.

When using Kerberos over HTTP(S), the driver returns a Kerberos error if the Host FQDN part of the service principal has mixed upper and lower case

This issue has been resolved.

When the result set includes a column that has a hash character (#) in its name, the driver does not return that column or any of the subsequent columns in the result set

This issue has been resolved.

Version 2.5.15

Enhancements & New Features

The following are highlights of the new features and functionalities that were added to Cloudera ODBC Driver for Apache Hive 2.5.15.

Debian 7 now supported

You can now install and use the driver on computers that run Debian 7 (Ubuntu 12.04 LTS and Ubuntu 14.04 LTS).

Two-way SSL verification added

Before, the driver only supported one-way SSL verification, where the client verifies the server. You can now configure connections where the client and the server both verify each other using SSL.

Custom HTTP headers now supported

You can now specify custom HTTP headers in connections that use HTTP as the transport protocol in the Thrift layer.

SSL and Thrift transport protocol now configured separately from authentication

You can now configure SSL and the Thrift transport protocol by using the new "Enable SSL" option (the SSL key) and the new "Thrift Transport" option (the ThriftTransport key), respectively. Use a combination of these two new options and the "Mechanism" option (the AuthMech key) as needed to configure authentication to the server.

As a result of this update, the following authentication mechanisms have been deprecated:

- User Name and Password (SSL) (AuthMech=4)
- HTTP (AuthMech=7)
- HTTPS (AuthMech=8)
- Kerberos over HTTP (AuthMech=9)

- Kerberos over HTTPS (AuthMech=10)

Support added for saving passwords as encrypted strings in the Windows Registry

When using the user interface of the Windows driver to configure a connection, you can now save the password for authentication (the PWD value) or the password for the client's SSL private key (the ClientPrivateKeyPassword value) by selecting the "Save Password" check box. The saved password will be obscured (not saved in plain text). However, it is still possible for the encrypted password to be copied and used.

Resolved Issues

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

Driver reports "STRING" as the data type name for VARCHAR columns

This issue has been resolved.

The fallback mechanism for executing queries is triggered even after a query cancellation exception occurs

This issue has been resolved.

When using SQLEngine to derive the result set metadata, the driver reports the result set metadata incorrectly

This issue has been resolved.

In some cases, the driver cannot find temporary tables that were created using the CREATE TEMPORARY TABLE syntax in non-native query mode

This issue has been resolved.

Driver retrieves tables metadata for partitioned tables incorrectly

This issue has been resolved.

Driver does not return results for queries that contain common table expressions

This issue has been resolved.

Version 2.5.14

Enhancements & New Features

The following are highlights of the new features and functionalities that were added to Cloudera ODBC Driver for Apache Hive 2.5.14.

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

The driver now supports query translation and parameterized INSERT, UPDATE, and DELETE statements when connected to Hive 0.14 or later.

Implemented dynamic service discovery feature

You can now configure the driver to discover Hive Server 2 services via the ZooKeeper service by setting the Service Discovery Mode option (the ServiceDiscoveryMode connection attribute).

Temporary Table feature in the driver integrated with the natively supported feature from Hive 0.14

When connecting to Hive 0.14 or later, the driver now uses the Temporary Table feature that is natively supported by Hive instead of the Temporary Table feature that is included in the driver. When you connect to Hive 0.14 or later, the Temporary Table feature is always enabled and you do not need to configure the feature in the driver.

Optimized query translations

API calls that are not needed for query translation have been removed. Before, some APIs for data retrieval were called during query translation.

Updated SQL translation logic for handling multi-row INSERT operations

The SQL translation logic in the driver has been updated to account for recent changes in the AE tree structure.

Windows 8 now supported

You can now install and use the driver on computers that run Windows 8.

Resolved Issues

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

Connections that authenticate through Kerberos over HTTPS do not work correctly

This issue has been resolved.

Driver cannot translate INSERT statements that contain a row with a NULL value in a numeric column and another row with a numeric literal in the same numeric column

This issue has been resolved.

On non-Windows platforms, driver-wide configurations set in the cloudera.hiveodbc.ini file are not applied

This issue has been resolved.

Driver opens a new session whenever a query fails to execute

This issue has been resolved.

During query translation, the driver retrieves an excessive amount of metadata

This issue has been resolved.

Version 2.5.13

Enhancements & New Features

The following are highlights of the new features and functionalities that were added to Cloudera ODBC Driver for Apache Hive 2.5.13.

Optimized metadata retrieval processes

The performance of the driver has improved for ODBC catalog function calls.

Resolved Issues

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

Driver cannot connect to a particular build of Hive 0.10

This issue has been resolved.

When an authentication error occurs, the driver returns an incorrect error message

This issue has been resolved. The driver now returns the correct SqlState and an improved error message if an authentication error occurs when using one of the following authentication mechanisms:

- Kerberos
- User Name
- User Name and Password
- User Name and Password (SSL)

Version 2.5.12

Enhancements & New Features

The following are highlights of the new features and functionalities that were added to Cloudera ODBC Driver for Apache Hive 2.5.12.

Updated the default location of the cloudera.hiveodbc.ini configuration file

By default, the cloudera.hiveodbc.ini configuration file is now located in the lib directory corresponding to the bitness of the driver. For example, the configuration file for the 64-bit driver on Linux is located in /opt/cloudera/hiveodbc/lib/64

Driver environment variable renamed to CLOUDERAHIVEINI

The environment variable for locating the cloudera.hiveodbc.ini configuration file has been renamed to CLOUDERAHIVEINI. The default value of CLOUDERAHIVEINI is the default installation location of the cloudera.hiveodbc.ini file.

Support added for translating ATAN2, RIGHT, LEFT, and CHAR scalar functions

The driver is now able to translate ATAN2, RIGHT, LEFT, and CHAR scalar functions.

Improved the process for detecting the Hive version and the HS2 protocol version

The driver is now better able to determine the Hive and HS2 protocol versions that are being used.

OpenSSL 1.0.0m supported

The driver has been upgraded to use OpenSSL 1.0.0m

Resolved Issues

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

The SEARCHED CASE syntax is not translated correctly

This issue has been resolved.

Version 2.5.11

Enhancements & New Features

The following are highlights of the new features and functionalities that were added to Cloudera ODBC Driver for Apache Hive 2.5.11.

Implemented driver configuration option for the temporary table feature

You can now choose to enable or disable the temporary table feature.

Resolved Issues

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

When query translation fails, the driver does not remove the semicolon from the end of the query string

This issue has been resolved.

Driver activity is logged in standard output (stdout) by the Thrift library

This issue has been resolved. Driver activity is no longer logged in stdout.

Driver cannot create log files if the file name contains special characters

This issue has been resolved.

Parameter values are placed incorrectly in parameterized queries

This issue has been resolved.

Version 2.5.10

Enhancements & New Features

The following are highlights of the new features and functionalities that were added to Cloudera ODBC Driver for Apache Hive 2.5.10.

IBM AIX now supported

You can now install and use the driver on computers that run IBM AIX.

Improved HIVE_SYSTEM table feature

The HIVE_SYSTEM table is now hidden from catalog functions by default. You can use driver configuration options to show or hide the HIVE_SYSTEM table.

"Unicode SQL character types" configuration option implemented

You can now use the "Unicode SQL character types" option (the UseUnicodeSqlCharacterTypes key) to control whether the driver returns SQL_WCHAR/SQL_WVARCHAR or SQL_CHAR/SQL_VARCHAR for CHAR, STRING, and VARCHAR columns.

Resolved Issues

The following issues were resolved in Cloudera ODBC Driver for Apache Hive 2.5.10.

The "Rows Fetched Per Block" setting affects the number of columns that the driver reports

This issue has been resolved. The setting for the "Rows Fetched Per Block" option (the RowsFetchedPerBlock key) no longer affects the number of columns that are returned.

Some HS2 operations are not closed as expected

This issue has been resolved.

Driver sometimes returns an ambiguous column reference error when there are too many columns in the SELECT list

This issue has been resolved.

Version 2.5.9

Enhancements & New Features

The following are highlights of the new features and functionalities that were added to Cloudera ODBC Driver for Apache Hive 2.5.9.

Improved usage of the HS2 protocol

The driver now uses one HS2 session per ODBC connection, so that any changes that are applied to the session will take effect and persist for all ODBC statements that are from the same ODBC connection.

Support added for the USE <database> query

You can now use the USE <database> query to run queries in a different context than the database that was set for the connection.

Optimized connection processes

The driver is now better able to reuse physical connections to Hive server.

OpenSSL and cURL libaries statically linked in the Linux and Mac OS X drivers

In the Linux and Mac OS X versions of the driver, the OpenSSL and cURL libraries are now statically linked. You no longer need to add the lib directory of the driver to the LD_LIBRARY_PATH and DYLD_LIBRARY_PATH environment variables.

Authentication using Kerberos over HTTP and HTTPS implemented

You can now configure the driver to authenticate a connection using Kerberos over HTTP or Kerberos over HTTPS.

Query cancellation implemented

You can now cancel queries that have been issued.

Support added for user credentials in the HTTP authentication mechanism

The HTTP authentication mechanism now accepts login through user name and password.

Resolved Issues

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

Subqueries are not translated correctly

This issue has been resolved.

Subqueries inside WHERE clauses are not handled correctly

This issue has been resolved.

Version 2.5.8

Resolved Issues

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

Driver sometimes terminates unexpectedly during connection establishment if Hive returns an error while executing the "set -v" query

This issue has been resolved.

Version 2.5.7

Enhancements & New Features

The following are highlights of the new features and functionalities that were added to Cloudera ODBC Driver for Apache Hive 2.5.7.

Support added for temporary tables

You can now work with temporary tables while using the driver.

CHAR(n) and DECIMAL(p,s) data types supported

The driver now supports the CHAR(n) and DECIMAL(p,s) data types.

Hive 0.13 Thrift result set serialization supported

The driver now supports result set serialization.

Support added for transforming CTAS queries

Query transformation now removes extra parentheses from the SELECT sub-tree.

"Get Tables With Query" feature implemented

You can now configure the driver to retrieve the names of tables in the Hive database by using the SHOW TABLES query instead of the GetTables Thrift API call.

Implemented configuration option for converting server-side property key names into lower case characters

You can now use the "Convert Key Name to Lower Case" option (the LCaseSspKeyName key) to configure the driver to convert all server-side property key names into lower case characters.

Windows driver UI updated to support new configuration options

The DSN Setup dialog box and the Driver Configuration tool now support the "Convert Key Name to Lower Case" and "Use Async Exec" options.

Support added for DATE and VARCHAR data types for HS1 and HS2

The driver now supports the DATE and VARCHAR data types for both the HS1 protocol and the HS2 protocol.

Resolved Issues

The following issues were resolved in Cloudera ODBC Driver for Apache Hive 2.5.7.

Query translation causes "Ambiguous column reference" error message in Hive

This issue has been resolved.

Queries that contain non-equal joins are not translated correctly

This issue has been resolved.

Incorrect results returned for queries containing the SIGN function

This issue has been resolved.

Queries that concatenate two strings are not translated correctly

This issue has been resolved.

Simple case and searched case clauses are not translated correctly

This issue has been resolved.

STRING column type appears as VARCHAR in the SQLColumns result set

This issue has been resolved.

HS1 client API calls for metadata retrieval do not work in Hive versions 0.9.0 to 0.11.0

This issue has been resolved. As a workaround, the driver now retrieves metadata for schemas, tables, and columns by executing a query instead of using API calls when connecting to HS1 in the affected versions of Hive.

Driver Configuration tool uses a system call that is not compatible with Windows XP

This issue has been resolved.

Version 2.5.6

Enhancements & New Features

The following are highlights of the new features and functionalities that were added to Cloudera ODBC Driver for Apache Hive 2.5.6.

Active Directory supported

The Windows version of the driver now supports Active Directory, enabling you to use Active Directory Kerberos as an authentication method.

Version 2.5.5

Resolved Issues

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

In some builds of Hive, the driver cannot open a session against Hive Server 2 if the "Delegation UID" option (the DelegationUID key) is not set

This issue has been resolved.

Version 2.5.4

Enhancements & New Features

The following are highlights of the new features and functionalities that were added to Cloudera ODBC Driver for Apache Hive 2.5.4.

Improved handling of driver configuration settings

Driver configuration settings in connection strings now take precedence over settings in DSNs.

Support added for DSNs and driver configurations from previous versions of the driver

The driver is now backwards compatible with DSNs and driver configurations that were created using previous versions of the driver.

"Delegation UID" configuration option implemented

You can now delegate all operations against Hive to a user that is different than the authenticated user for the connection by using the "Delegation UID" option (the DelegationUID key).

Authentication using HTTP and HTTPS implemented

You can now configure the driver to authenticate a connection using HTTP or HTTPS.

Support added for vendor-specific default settings for DSNs

The driver now displays vendor-specific default settings when you create a new DSN.

Improved handling of right brace characters in the Test Connection dialog box

The Test Connection dialog box can now escape right brace characters (}) in the connection attribute value and then enclose the value in braces ({}).

Support added for braces ({}) in connection attribute values

The driver now accepts attribute values that are enclosed in braces ({}).

"Binary column length" configuration option implemented

You can now specify a maximum data length for BINARY columns by using the "Binary column length" option (the BinaryColumnLength key).

Updated default values for connection properties

The values of the following connection properties are now set to 128:

- SQL_MAX_CATALOG_NAME_LEN
- SQL_MAX_COLUMN_NAME_LEN
- SQL_MAX_IDENTIFIER_LEN
- SQL_MAX_SCHEMA_NAME_LEN
- SQL_MAX_TABLE_NAME_LEN

Support added for asynchronous ExecuteStatement call

When connected to HS2 on HDInsight, you can now use the asynchronous version of the ExecuteStatement Thrift Hive client call.

Support added for using alias configuration keys for UID and PWD

You can now use "User ID" and "Password" as alias configuration keys for the UID and PWD configuration keys, respectively.

Updated Driver Configuration tool

You can now use the Driver Configuration tool to configure advanced driver options and all of the authentication options that the driver supports.

Updated default port numbers for DSNs

When creating a DSN, the default port number now varies depending on your connection:

- When connecting to HDInsight on Azure, the default port number is 563
- When connecting to HDInsight, the default port number is 10001
- For all other connections, the default port number is 10000

Authentication using User Name and Password with SSL implemented

You can now configure the driver to authenticate a connection using Secure Sockets Layer (SSL) and user login.

Support added for returning list of Hive functions for SQLProcedures

The driver is now able to return a list of Hive functions for SQLProcedures.

Improved performance for fetching metadata

The performance of the driver has improved for fetching metadata.

Authentication using User Name and Password implemented

You can now configure the driver to authenticate a connection using user login.

Resolved Issues

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

Application stops working if SQLCancel is called

This issue has been resolved. SQLCancel is not supported. Now, when SQLCancel is called, SQL_ERROR is reported.

In the USE <database> query, the database name is not quoted

This issue has been resolved.

Identifiers in the generated HiveQL are not quoted

This issue has been resolved.

In some cases, Hive returns an empty error message

This issue has been resolved. The driver now returns more meaningful error messages.

When a query fails to execute, the driver does not report the failure until the result is fetched

This issue has been resolved.

Special characters in string literals and parameters are not escaped correctly

This issue has been resolved.

Braces ({}) syntax for attribute values in connection strings is not handled correctly during parsing

This issue has been resolved.

Transform/Map-Reduce query result is treated as a ROWCOUNT result

This issue has been resolved.

SQL_VARCHAR data is returned for BINARY columns

This issue has been resolved. The driver now returns SQL_VARBINARY data for BINARY columns.

Queries that contain the NOT predicate are not transformed correctly

This issue has been resolved.

Driver does not correctly execute Base64 encoding of User Name and Password

This issue has been resolved.

Server-side properties specified in the DSN are not applied

This issue has been resolved.

Version 2.5.3

Enhancements & New Features

The following are highlights of the new features and functionalities that were added to Cloudera ODBC Driver for Apache Hive 2.5.3.

Driver Configuration Tool updated to support "Default string column length" option

The Driver Configuration tool in the Windows driver now supports the "Default string column length" option.

Version 2.5.2

Enhancements & New Features

The following are highlights of the new features and functionalities that were added to Cloudera ODBC Driver for Apache Hive 2.5.2.

Updated driver documentation

The information in the Installation and Configuration Guide has been revised and improved.

Version 2.5.1

Enhancements & New Features

The following are highlights of the new features and functionalities that were added to Cloudera ODBC Driver for Apache Hive 2.5.1.

Updated driver documentation

The information in the Installation and Configuration Guide has been revised and improved.

Version 2.5.0.1001

Enhancements & New Features

The following are highlights of the new features and functionalities that were added to Cloudera ODBC Driver for Apache Hive 2.5.0.1001.

Implemented option for controlling how the driver applies server-side properties when connecting to a Hive Server 2 instance

Normally, the driver applies server-side properties by executing SET queries. Some configurations of Hive Server 2 support a more efficient method for applying server-side properties that does not involve additional network round-tripping. Use the "Apply Server Side Properties with Queries" option (the ApplySSPWithQueries key) to specify how server-side properties are applied.

Version 2.5.0.1000

Version 2.5.0.1000 was the initial release of the Cloudera ODBC 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.