

Release Notes

Amazon JDBC Connector for Apache Hive 2.6.19

Released 2022-03-04

These release notes provide details of enhancements, features, known issues, and workflow changes in Amazon JDBC Connector for Apache Hive 2.6.19, as well as the version history.

For information about upcoming support deprecations or removals, see the Workflow Changes section. Deprecated features will not receive any updates, but will continue to be usable in their current state until support is removed in a future release.

Enhancements & New Features

Updated log4j third-party libraries

The JDBC 4.2 version of the connector has been updated to version 2.17.1 of the log4j third-party libraries.

The JDBC 4.1 version of the connector has been updated to version 2.12.4 of the log4j third-party libraries. To address security vulnerabilities, do one of the following:

- In PatternLayout in the logging configuration, replace Context Lookups like `{ctx:loginId}` or `$$ {ctx:loginId}` with Thread Context Map patterns (`%X`, `%mdc`, or `%MDC`).
- Otherwise, in the configuration, remove references to Context Lookups like `{ctx:loginId}` or `$$ {ctx:loginId}` where they originate from sources external to the application such as HTTP headers or user input.

Updated Jackson library

The connector now uses version 2.13.1 of the Jackson library. Previously, the connector used Jackson version 2.12.3.

Resolved Issues

The following issue has been resolved in Amazon JDBC Connector for Apache Hive 2.6.19.

- When `close()` is not called on the connection object, the connector prevents the JVM from exiting.

Workflow Changes

The following changes may disrupt established workflows for the driver.

In addition to changes that are already implemented in the current version of the driver, this section describes potentially disruptive changes that will be implemented in a future version of the driver, so that you can plan accordingly.

Upcoming

Removing support for Hive versions earlier than 1.0.0.

As early as December 2021, the connector will no longer support Hive versions earlier than 1.0.0. For a list of supported Hive versions, see the *Installation and Configuration Guide*.

Removing support for Hive Server 1

As early as December 2021, the connector will no longer support Hive Server 1. For a list of supported Hive Server versions, see the *Installation and Configuration Guide*.

Version 2.6.19

Removed slf4j-log4j12 dependency

Beginning with this release, the connector no longer includes the `slf4j-log4j12` library in the `JAR` file.

Version 2.6.17

Removed support for JDBC 4 (Java 6)

Beginning with this release, the connector no longer supports JDBC 4 (Java 6). For a list of supported JDBC versions, see the *Installation and Configuration Guide*.

Version 2.6.1

Removing support for JDBC 3 (Java 5)

Beginning with this release, the driver no longer supports JDBC 3 (Java 5). For a list of supported JDBC versions, see the *Installation and Configuration Guide*.

Version History

Version 2.6.17

Released 2021-12-24

Enhancements & New Features

Updated log4j third-party libraries

The JDBC 4.2 version of the connector has been updated to version 2.17.0 of the log4j third-party libraries.

The JDBC 4.1 version of the connector has been updated to version 2.12.2 of the log4j third-party libraries. To address security vulnerabilities, do one of the following:

- In PatternLayout in the logging configuration, replace Context Lookups like `{ctx:loginId}` or `$$ {ctx:loginId}` with Thread Context Map patterns (`%X`, `%mdc`, or `%MDC`).
- Otherwise, in the configuration, remove references to Context Lookups like `{ctx:loginId}` or `$$ {ctx:loginId}` where they originate from sources external to the application such as HTTP headers or user input.

Third-party library upgrades

The connector has been upgraded with the following third-party libraries:

- Apache Commons Codec 1.15 (previously 1.9)
- Apache HttpClient 4.5.13 (previously 4.5.3)
- Apache HttpCore 4.4.14 (previously 4.4.6)
- Jackson 2.12.3 (previously 2.10.1)

The JDBC 4.1 connector has been updated to use slf4j 1.7.30.

Upgraded ZooKeeper library

The JDBC 4.2 version of the connector now uses version 3.7.0 of the ZooKeeper and ZooKeeper Jute library. Previously, this version of the connector used ZooKeeper version 3.4.14.

The JDBC 4.1 version of the connector continues to use ZooKeeper version 3.4.14 because ZooKeeper versions later than 3.4.14 no longer support Java 7.

SSO authentication

You can now authenticate using single sign-on (SSO) via SAML 2.0 on supported servers. For more information, see the *Installation and Configuration Guide*.

Updated SSL support

The connector now supports BCFKS TrustStores files. To specify the TrustStore type, set the `SSLTrustStoreType` property to the preferred type. For more information, see the *Installation and Configuration Guide*.

Simulate transactions

You can now specify whether the connector simulates or ignores transactions. To do this, set the `IgnoreTransactions` property. For more details, see the *Installation and Configuration Guide*.

Updated documentation

The `FastConnection` and `LoginTimeout` properties are now documented in the *Installation and Configuration Guide*. For more details, see the *Installation and Configuration Guide*.

Updated ZooKeeper support

When connecting with ZooKeeper, the connector now returns available servers in a round-robin rotation.

Support for Java 11

The connector now supports Java 11.

Resolved Issues

The following issues have been resolved in Amazon JDBC Connector for Apache Hive 2.6.17.

- The HiveJDBCDataEngine does not check getFunctions and returns an empty result set.
- When a ZooKeeper connection URL does not contain a namespace for each host the connector does not close the connection properly.
- When connecting to HiveServer1 instances using the Connection isValid method, the connector returns an incorrect value.
- The connector logs incorrect socket timeout values.
- When the connection is not valid, the connector does not notify the ConnectionEventListener object as expected.

Version 2.6.12

Released 2020-09-04

Enhancements & New Features

Binary column length

You can now specify the maximum size of binary columns. To do this, set the `BinaryColumnLength` property to the specified column length. For more details, see the *Installation and Configuration Guide*.

Updated third-party libraries

The driver now uses the following third-party libraries:

- jackson 2.10.1
- log4j 2.13.3
- slf4j 1.7.30
- Thrift 0.13.0

The JDBC 4.1 driver continues to use the previous versions of these libraries.

Resolved Issues

The following issues have been resolved in Amazon JDBC Connector for Apache Hive 2.6.12.

- The driver does not honor the values set via the `Statement.setFetchSize` JDBC API call.

- Queries with a UNION clause are replaced with a UNION ALL clause.
- The JDBC 4.1 version of the driver fails to connect to servers that require encryption using TLS 1.2.
- When you invoke the JDBC 4.2 driver with the Datasource class, an exception occurs.
- When you use the driver with the Denodo application, it returns the following error: "Could not initialize Class".
- In some cases, Hive queries fail with the following error:

```
"com.amazon.hiveserver2.exceptions.ExceptionConverter.toSQLException:  
[JDBC Driver]Index: 4, Size: 4"
```
- SQL statements using the EXISTS predicate return an error.
- The driver returns incorrect results for decimal columns.

Version 2.6.9

Released 2019-10-04

Enhancements & New Features

Updated third-party libraries

The driver now uses the following third-party libraries:

- jackson 2.9.9
- jackson-databind 2.9.9.3
- Log4j 1.2.17
- slf4j-api 1.7.25
- slf4j-log4j12 1.7.25
- Thrift 0.12.0
- ZooKeeper 3.4.14

Support for JDBC 4.2

The driver now supports JDBC 4.2. For more details see the *Installation and Configuration Guide*.

New JDBC class names

The driver now supports the following class names for Driver and DataSource that are independent of the JDBC version used:

- `com.amazon.hive.jdbc.HS1Driver`
- `com.amazon.hive.jdbc.HS2Driver`
- `com.amazon.hive.jdbc.HS1DataSource`
- `com.amazon.hive.jdbc.HS2DataSource`

The previous JDBC-version-specific class names for 4.0 and 4.1 are still supported.

Renaming join columns

By default, the driver does not allow join columns to be renamed. To enable the renaming of join columns, set the `RenameJoinColumn` property to `true`.

Updated Hive Service support

The driver now uses Hive Service 3.1.1 for HiveServer2 connections.

Updated Hive support

The driver now supports Apache Hive versions 0.11 through 3.1.

Resolved Issues

The following issues have been resolved in Amazon JDBC Connector for Apache Hive 2.6.9.

- The driver resolves the host name to an IP address for SSL verification, causing the host name verification step to fail.
- The driver incorrectly treats `SSLTrustStore` and `SSLTrustStorePWD` as server-side properties.
- If a SQL query ends in a semicolon, the driver reports an error.

This issue has been resolved. The driver now removes trailing semicolons from queries before sending them to the server.

- Queries that use REGEX column specification fail.

This issue has been resolved. The driver no longer attempts to parse REGEX column specification.

Version 2.6.2

Released 2018-08-03

Resolved Issues

The following issue was resolved in Amazon JDBC Connector for Apache Hive 2.6.2.

- The driver does not behave according to the `SocketTimeout` setting, and stops responding instead of returning an error when the server takes too long to respond to a request.

This issue has been resolved. The driver now returns an error if the amount of time taken by the server to respond to a request exceeds the number of seconds specified by the `SocketTimeout` setting.

Version 2.6.1

Released 2018-04-24

Enhancements & New Features

Improved installation process

Third-party libraries are now bundled with the driver as a shaded jar file.

Updated hive-service

The driver has been upgraded to use `hive-service-2.0.1` with JDBC 4.1 and `hive-service-1.1.1` for JDBC 4.0.

Updated third-party libraries

The driver now uses the following third-party libraries:

- `http-client 4.5.3`
- `http-core 4.4.6`
- `commons-codec 1.9`
- `commons-logging 1.2`
- `ZooKeeper 3.4.11`

Resolved Issues

The following issues were resolved in Amazon JDBC Connector for Apache Hive 2.6.1.

- A string literal after a back-tick (`) is treated as a comment.
- When the driver attempts to connect to the server using Beeline, an error occurs on the server and an exception is logged by the driver.

This issue has been resolved. Now, when connecting to the server using Beeline, connection properties such as `driver` and `url` are no longer treated as configuration properties for the server, and the driver does not send SET queries to the server with these properties.

- The driver removes query comments and hints when the `UseNativeQuery` property is set to 1.

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://aws.amazon.com/support>.