

Release Notes

Amazon Redshift JDBC Driver 1.2.45

Released May 2020

These release notes provide details of enhancements, features, known issues, and workflow changes in Amazon Redshift JDBC Driver 1.2.45, as well as the version history.

Enhancements & New Features

Lowercase DbGroups

You can now configure the driver to lowercase all DbGroups that are received from the identity provider. To do this, set the new `ForceLowercase` property to `true`. For more information, see the *Installation and Configuration Guide*.

Support for GEOMETRY data

The driver now supports data of type GEOMETRY. For more information, see the *Installation and Configuration Guide*.

Updated Jackson libraries

The JDBC 4.2 version of the driver now uses the following versions of the libraries for the Jackson JSON parser:

- jackson-annotations-2.10.1
- jackson-core-2.10.1
- jackson-databind-2.10.1
- jackson-dataformat-cbor-2.10.1

The JDBC 4 and 4.1 versions of the driver use version 2.6.x of these libraries.

Resolved Issues

The following issues have been resolved in Amazon Redshift JDBC Driver 1.2.45.

- When a query in SQLWorkbench is cancelled, the driver becomes unresponsive.
- When the Java application uses `com.amazon.redshift.api.PGTimestamp` directly, an exception occurs.

- The function `getProcedureColumns()` returns incorrect `COLUMN_TYPE`, `ORDINAL_POSITION`, and `NULLABLE` columns.
- The function `getColumns()` does not return the `IS_GENERATEDCOLUMN` column.
- When `getProcedures()` is called without a schema name filter and with a procedure name filter, it returns incorrect results.
- When the schema name includes unicode characters and the JVM is set to use the windows-1252 character set, the driver is unable to fetch table names using catalog functions.
- The `getColumns()` metadata for an external table returns "unknown type".
- When the driver uses SSO authentication, it sends duplicate parameters and the authentication request fails.
- When the driver uses PingFederate authentication, the driver incorrectly sends password to a "passwordReset" field and the authentication fails.

Known Issues

The following are known issues that you may encounter due to limitations in the data source, the driver, or an application.

- The driver does not support the `getFunctions()` function.
- The driver does not support the use of parameter names in `CallableStatement`'s `registerOutParameter()` function.

As a workaround, use an indexed argument call instead of a named argument call.

- The driver does not correctly parse strings that contain both escaped backslash (`\`) and escaped single quote (`'`) characters.
- In some cases, the driver incorrectly parses strings using C-style escapes.
- The driver cannot return `REFCURSOR` data as `OBJECT` data for Redshift instances running on a single-node cluster.

When working with a single-node cluster, the driver can only return `REFCURSOR` data as `VARCHAR` data.

This limitation exists because single-node clusters do not support the `FETCH ALL` command syntax that is required for returning `REFCURSOR` data as

OBJECT data. If you attempt to retrieve REFCURSOR data this way, the driver returns the following error message: "FETCH ALL not supported in single-node clusters".

- Temporal functions are returned as UTC dates or timestamps.

In the latest version of the open-source PostgreSQL driver, temporal functions are returned in the user's local time zone. Version 8.4 of the open-source driver returns temporal functions as UTC dates or timestamps instead. The Amazon Redshift JDBC Driver is consistent with version 8.4 of the open-source driver.

Workflow Changes

The following changes may disrupt established workflows for the driver.

Version 1.2.41

Removed support for JDBC 4.0 and 4.1 (Java 6 and 7)

Beginning with this release, the driver no longer supports JDBC 4.0 (Java 6) and JDBC 4.1 (Java 7). For a list of supported JDBC versions, see the *Installation and Configuration Guide*.

Version History

Version 1.2.43

Released April 2020

Resolved Issues

The following issues have been resolved in Amazon Redshift JDBC Driver 1.2.43.

- The driver cannot parse dollar-quoted string constants with a tag of one or more characters.
- The driver does not allow the usage of DCL, DDL or DML keywords as table names or column names in a statement when `setReadOnly` is set.
- Improper handling when user interrupts driver execution.

This issue has been resolved. The driver now handles `InterruptedException` and sends a cancel request to the server.

Version 1.2.41

Released February 2020

Enhancements & New Features

IAM authentication with browser plugin

You can now use a browser plugin to authenticate your connection through your identity provider's website. For more information, see the *Installation and Configuration Guide*.

Begin transactions with START TRANSACTION

You can now use START TRANSACTION as a command to start a transaction.

Improved Azure AD error messages

The driver now provides more comprehensive error messages for the Azure AD plugin.

Removal of SLF4J dependency

The driver no longer includes SLF4J as a dependency.

Resolved Issues

The following issues have been resolved in Amazon Redshift JDBC Driver 1.2.45.

- In some cases, the driver incorrectly reports a column's nullability as "Not Nullable" instead of "Unknown".
- If you cancel a query before the query is in execution, the driver returns an error.

This issue has been resolved. Now, if you cancel a query before the query is in execution, the cancellation is ignored. However, the cancel request is still recorded in the driver logs.

- If the data source does not provide a column length, the driver returns an error.
- The driver does not return time zone information when returning a string representation of a timestamp.

Version 1.2.37

Released November 2019

Enhancements & New Features

Azure AD authentication

The driver now supports authentication through Azure AD. For more information, see the *Installation and Configuration Guide*.

SAML assertion included in driver log

If single sign-on authentication is used and the `DSILogLevel` connection property is set to `6`, the driver now always logs the SAML assertion.

Use DbUser from SAML assertion

The driver can now override the specified `DbUser` value with a value from the SAML assertion. To enable this, set the new `AllowDBUserOverride` connection property to `1`. For more information, see the *Installation and Configuration Guide*.

Resolved Issues

The following issue has been resolved in Amazon Redshift JDBC Driver 1.2.37.

- In some cases, the driver is unable to parse complex queries.

Version 1.2.36

Released October 2019

Enhancements & New Features

Optimization of catalog functions

Prepared statement reuse now includes the external schema check query. This improves the performance of catalog functions.

Resolved Issues

The following issue has been resolved in Amazon Redshift JDBC Driver 1.2.36.

- If the driver receives a Notice response while reading data, it terminates unexpectedly.

This issue has been resolved. Now the driver continues reading data until it receives a Ready For Query or Error response.

Version 1.2.35

Released October 2019

Resolved Issues

The following issues were resolved in Amazon Redshift JDBC Driver 1.2.35.

- The driver sends TCP keepalive packets on a different connection than the one being used to access Redshift data.
- The driver ignores empty fields in data arrays.
- When you query data from an S3 folder that contains the "like" keyword in its name, the query fails due to a parsing issue.
- In some cases, when `Connection.setReadOnly()` is set to `true`, queries that involve large amounts of data run more slowly than expected.

Version 1.2.34

Released July 2019

Enhancements & New Features

Enhanced support for query timeout

The driver now supports `Statement.setQueryTimeout`, to indicate how many seconds the driver waits for the `Statement` object to execute before sending a cancel message to the server.

Resolved Issues

The following issues were resolved in Amazon Redshift JDBC Driver 1.2.34.

- Error reporting does not catch all exceptions that occur during data retrieval.
- Driver does not correctly parse strings with complex combinations of escaped and non-escaped single quote (`'`) and double quote (`"`) characters.

- Java file descriptors are retained after the connection is closed.
- Driver fails to properly escape question (?) marks in a PreparedStatement.

Version 1.2.32

Released June 2019

Enhancements & New Features

Full support for REFCURSOR parameters in stored procedures

When connected to Redshift instances that are running on a multi-node cluster, the driver is now able to retrieve REFCURSOR parameters as OBJECT data.

Support for read-only connections

The driver now supports the `Connection.setReadOnly()` method. This is used to indicate if the connection should have read-only privileges.

Resolved Issues

The following issues were resolved in Amazon Redshift JDBC Driver 1.2.32.

- PreparedStatement parameters in some multi-threaded environments are not correctly populated.
- In some cases, the driver fails to connect when using Okta and the Oracle SQL Developer client application.
- Stored procedure calls with the syntax `? = call [sp_name]` are not parsed correctly.
- In some cases, parameters are not bound correctly for stored procedures that only have IN parameters.
- In some cases, quotation marks (") are not escaped correctly.
- In cases where there are multiple stored procedures with the same name, and the user calls one of them, the driver picks the first one in the list.
- This issue has been resolved. In this situation, the driver now displays an error stating that there are multiple stored procedures with that name, and the query fails.

Version 1.2.30

Released May 2019

Enhancements & New Features

Improved support for REFCURSOR parameters in stored procedures

When connected to Redshift instances that are running on a multi-node cluster, the driver is now able to retrieve REFCURSOR parameters as OBJECT data.

For information about the REFCURSOR support that is available for a single-node cluster, see the "Known Issues" section.

Version 1.2.27

Released April 2019

Enhancements & New Features

Stored procedures

You can now use the driver to execute stored procedures. The driver provides full support for all Amazon Redshift stored procedures.

The driver provides limited support for REFCURSOR data, and can only return it as data of type VARCHAR.

Version 1.2.20

Released February 2019

Enhancements & New Features

Improved getColumn support

The catalog function `getColumn` now also returns the column list for use with late binding views.

Improved connection string parsing

The driver can now parse connection strings more efficiently.

Improved security

The driver now provides stronger security for catalog functions, improving protection against SQL injection attacks.

Resolved Issues

The following issues have been resolved in Amazon Redshift JDBC Driver 1.2.20.

- When using `Statement.addBatch()`, the driver may not properly split multiple queries.
- The driver incorrectly reports that the data source is read-only.
- The `CONVERT` function is not always parsed correctly.
- In some cases, driver may have issues parsing queries containing repetitive word patterns.
- In some cases, the driver may have issues parsing queries containing large numbers of blanks in the column list.
- The `ResultSet.isLast()` function always returns `false`.
- In some cases, the driver fails to correctly parse some escape characters.
- The `getSchemas()` function reports an error when a null argument is passed to the catalog function.
- In some cases, the question mark (?) character is not escaped properly in queries.

Version 1.2.16

Released June 2018

Enhancements & New Features

Improved query execution performance

The driver now more efficiently executes queries that do not use `BlockingRowMode`.

Resolved Issues

The following issue has been resolved in Amazon Redshift JDBC Driver 1.2.16.

- The driver is unable to connect to the data store using Okta authentication.

Contact Us

For support, check the Amazon Redshift Forum at <https://forums.aws.amazon.com/forum.jspa?forumID=155> or open a support case using the AWS Support Center at <https://aws.amazon.com/support>.