

Release Notes

Amazon Redshift JDBC Driver 1.2.51

Released November 2020

These release notes provide details of enhancements, features, known issues, and workflow changes in Amazon Redshift JDBC Driver 1.2.51, as well as the version history.

Enhancements & New Features

Return metadata from multiple data stores

The driver can now return metadata from multiple Redshift databases and clusters. To enable this, set the `DatabaseMetadataCurrentDbOnly` property to 0. For more information, see the *Installation and Configuration Guide*.

Query group labels

The driver can now apply a user-defined label to all queries made from a connection. To do this, specify the label with the `QueryGroup` connection property. For more information, see the *Installation and Configuration Guide*.

Improved block comment recognition

Block comment (`/* */`) recognition has been improved, making query processing faster.

Improved driver logging

The driver logging has been improved for IAM connections.

Resolved Issues

The following issues have been resolved in Amazon Redshift JDBC Driver 1.2.51.

- The driver mistakenly recognizes any occurrence of `(tag)` as a dollar-quoted string.

As a workaround, if the same identifier containing a `$` is used in a query, the identifier must be enclosed in double quotes (`"`).

- The driver incorrectly translates `"CALL procedureName(arg1, arg2, ... argi)"` to `"SELECT procedureName(arg1, arg2, ... argi)"` when it fails to retrieve parameter information for the procedure.

- This issue has been resolved. Now, if the CALL statement is not enclosed in curly braces ({ }), there are no modifications done to the query.
- The driver incorrectly treats a quoted identifier as the actual procedure name when trying to retrieve procedure parameter information, and therefore is not able to find the procedure information.

This issue has been resolved. The driver now strips double quotes around the procedure name for quoted identifiers before attempting to retrieve the procedure information.

- In some cases, when using the Browser Azure AD plugin, a race condition occurs and the connection fails intermittently.
- In some cases, using a password in a connection URL causes the connection to fail with the error message "[Amazon](500310) Invalid operation: password authentication failed for user "xxxx".", even when the password is correct.
- In some cases, batches with multiple parameter sets are not initialized correctly.
- When SQLProcedureColumns() is called for a stored procedure with multiple OUT parameters, only one OUT parameter is returned.
- When using getClientInfoProperties, NAME returns APPLICATIONNAME.

This issue has been resolved. Now, ApplicationName is returned.

- getSchemas() does not narrow down results to the specified schemaPattern.

Known Issues

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

- User defined operators containing a question mark (?) are not processed correctly.

A question mark at the beginning and/or end of a user defined operator is recognized as a parameter and not as a part of the operator as there is no way to distinguish if '?' is part of the operator or not in this case.

- Strings containing PWD not parsed correctly.

If the property name following PWD is not a valid property name supported by Redshift, it is treated as a part of the PWD value. This is because PWD can contain '&', ';' and '=' meaning the driver does not know where the PWD value stops until it finds the next valid property name following a ';'.

For example:

- `jdbc:redshift://[HOST]:5439/simba?UID=[UID]&PWD=[PWD]&BlockingRowsMode=0`: the extracted PWD is `[PWD]`.
 - `jdbc:redshift://[HOST]:5439/simba?UID=[UID]&PWD=[PWD]&UnknownSetting=0`: the extracted PWD is `[PWD]&UnknownSetting=0`.
- The driver does not accept multi-statement queries that contain output parameters.
 - The driver does not support daylight saving time when inserting TIMETZ data in PreparedStatement's `setObject()` and `setTime()` functions.
 - The driver cannot return the fractional seconds and timezone components for TIMETZ data in CallableStatement's `getTime()` function when it is used for stored procedure output parameter.
 - The driver cannot return REF_CURSOR 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 REF_CURSOR data as VARCHAR data.

This limitation exists because single-node clusters do not support the FETCH ALL command syntax that is required for returning REF_CURSOR data as OBJECT data. If you attempt to retrieve REF_CURSOR 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 1.2.51

Removed support for DSILogLevel

Beginning with this release, the driver no longer supports the `DSILogLevel` connection property. For more information about logging configuration, see the *Installation and Configuration Guide*.

Version 1.2.50

New third-party dependency

Beginning with this release, the driver uses the third-party library `antlr4-4.5.3.jar`.

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 1.2.50

Released November 2020

Enhancements & New Features

Improved query processing

The query processing engine can now properly identify comments and quoted strings, and therefore process queries correctly.

Filter DbGroups

You can now configure the driver to filter all `DbGroups` that are received from the SAML response in the Azure, Browser Azure, and Browser SAML authentication types. To do this, specify the filter with the `DbGroupsFilter` connection property. For more information, see the *Installation and Configuration Guide*.

Support for registerOutParameter

The driver now supports registerOutParameter with a parameter name for stored procedures.

Rename OMNI datatype to SUPER

Data of type OMNI is now known as SUPER. For more information, see the *Installation and Configuration Guide*.

Resolved Issues

The following issues have been resolved in Amazon Redshift JDBC Driver 1.2.50.

- When using getTypeInfo with PRECISION and MAXIMUM_SCALE columns, the driver returns incorrect values.
- When java.sql.DatabaseMetaData is not used, read-only mode is not set.
- When BlockingRowsMode is enabled, query cancel is unresponsive.
- When executing a parameterized SELECT query with a SUPER (previously named OMNI) column in the WHERE clause, the driver returns an error.
- When the IAM role is set to the China region, the driver fails to parse a SAML assertion.
- When SQLColumns is called for NUMERIC and DECIMAL columns in late binding view, the driver returns an error.
- If more than one TIME data is selected, fractional seconds are not included in the column.

Version 1.2.47

Released July 2020

Enhancements & New Features

Support for getFunctions()

The driver now supports the getFunctions() catalog function.

Support for getFunctionColumns()

The driver now supports the getFunctionColumns() catalog function.

Support for OMNI data

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

Support for TIME data

The driver now supports data of types TIME and TIMETZ. For more information, see the *Installation and Configuration Guide*.

Specify relying party trusts

You can now configure the driver to allow different relying party trusts for AD FS authentication. To do this, when the authentication type is AD FS, specify the relying party with the `loginToRp` connection property. For more information, see the *Installation and Configuration Guide*.

Support for read-only mode

You can now configure the driver to use read-only mode. To do this, set the new `ReadOnly` property to `true`. For more information, see the *Installation and Configuration Guide*.

Additionally, you can change the read-only settings by using the `setReadOnly` API call.

Resolved Issues

The following issues have been resolved in Amazon Redshift JDBC Driver 1.2.47.

- `DatabaseMetaData.supportsBatchUpdates` is incorrectly set to `false`.
This issue has been resolved. `DatabaseMetaData.supportsBatchUpdates` is now set to `true` to match the driver capability.
- When a cluster restarts, the driver becomes unresponsive.
- When using `PreparedStatement`, query cancel is unresponsive after retrieving the first few rows.
- When there are multiple `RAISE` statements in a stored procedure, the driver does not return values.
- When retrieving a `NULL` value using `getString()` from a column of `GEOMETRY` data type, the driver returns a `NullPointerException`.
- When `getProcedureColumns` is called more than once within the same connection, the driver returns an error.

- When using IAM authentication, the SAML response parsing is not secure.
- When using IAM authentication, the profile name in `~/.aws/config` is not detected.
- When using numeric data in late binding views, the driver returns incorrect `DECIMAL_DIGITS` values in `SQLColumns`.

Version 1.2.45

Released May 2020

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.

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 1.2.41.

- 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.

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>.