

Release Notes

Amazon Redshift JDBC Driver 1.2.27

Released April 2019

These release notes provide details of enhancements, features, and known issues in Amazon Redshift JDBC Driver 1.2.27, as well as the version history.

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 REF_CURSOR data, and can only return it as data of type VARCHAR. For more information, see the "Known Issues" section.

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 cannot return REF_CURSOR data as OBJECT data.

When using a stored procedure to retrieve a result set, the driver cannot register REF_CURSOR parameters as data of type OBJECT in the stored procedure.

In this release of the driver, REF_CURSOR data can only be returned as data of type VARCHAR.

- Connection behavior is different compared to the open-source PostgreSQL driver.

When using the Amazon Redshift JDBC Driver, calling `Connection.prepareStatement()` sends all queries to the server for compilation prior to returning. The open-source PostgreSQL driver does not send queries to the server at this step.

- 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.

Version History

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.

Version 1.2.15

Released May 2018

Enhancements & New Features

Improved large query parsing performance

The driver now more efficiently parses large queries, such as queries that have multiple insert values or many columns.

Specify SelectorProvider class

You can now specify a `SelectorProvider` class to open a specific `SocketChannel`. To do this, set the `SelectorProvider` property to the fully qualified class path, and set the `SelectorProviderArg` property to the argument to be passed to the class constructor. For more information, see the Installation and Configuration Guide.

Resolved Issues

The following issues have been resolved in Amazon Redshift JDBC Driver 1.2.15.

- `SELECT {fn DAYNAME({fn CURDATE()}) }` causes an error.
- Invalid `TRUNCATE` syntax does not throw an exception.
- `TRUNCATE` does not report permission errors to the user.

This issue has been resolved. Previously, if the user attempted a `TRUNCATE` operation on a relation that they did not own, the operation failed silently and the driver logged an error message. Now the driver returns an error.

- Cluster identifiers ending in `.cn` cause an error.

Version 1.2.14

Released April 2018

Enhancements & New Features

Internal driver improvements

The driver has been updated with improvements for the security features.

Version 1.2.12

Released February 2018

Enhancements & New Features

Partner SPID support

When configuring authentication using the PingFederate service, you can now specify a partner SPID (service provider ID) value. To do this, set the new `partner_spid` property to the ID value. For more information, see the Installation and Configuration Guide.

Okta application name support

When configuring authentication using the Okta service, you can now specify the Okta application name. To do this, set the new `App_Name` property to the name of the application. For more information, see the Installation and Configuration Guide.

Resolved Issues

The following issue has been resolved in Amazon Redshift JDBC Driver 1.2.12.

- Autocommit is changed from TRUE to FALSE after a ROLLBACK is issued.

Version 1.2.11

Released December 2017

Enhancements & New Features

Support for instance profile credentials

The driver now supports IAM authentication using instance profile credentials when connecting with an EC2 instance.

Resolved Issues

The following issues have been resolved in Amazon Redshift JDBC Driver 1.2.11.

- Autocommit is changed from TRUE to FALSE after a BEGIN statement.
- Incorrect metadata is returned for datatypes of some parameters in prepared statements.
- Single sign-on fails for some versions of AD FS.

Version 1.2.10

Resolved Issues

The following issues have been resolved in Amazon Redshift JDBC Driver 1.2.10.

- TRUNCATE and COPY commands throw an exception when autocommit mode is enabled.
- Driver class does not load when there is a space in the file path to the Redshift JDBC driver jar.

Version 1.2.8

Released September 2017

Enhancements & New Features

Improved jar packaging for jar with no AWS SDK

Resolved Issues

The following issue was resolved in Amazon Redshift JDBC Driver 1.2.8.

- Driver incorrectly reports ssl and sslmode properties are in conflict.

Version 1.2.7

Released August 2017

Enhancements & New Features

Support for Federated Authentication with Single Sign-On

The driver now supports IAM authentication, and federated authentication through AD FS, PingFederate, and Okta Identity Providers.

Enhanced Support for Viewing External Redshift Spectrum Tables

The driver now supports external tables in Amazon Redshift Spectrum.

Changed default SSL mode

The default SSL mode is now `verify-ca`.

Bundled SSL root certificates

The driver now includes the required CA root certificate for SSL verification.

Resolved Issues

The following issues were resolved in Amazon Redshift JDBC Driver 1.2.7.

- `setQueryTimeout` does not timeout a query.
- The driver does not send the appropriate `BEGIN` to keep the connection in `autocommit off` for `TRUNCATE` or `END`.

- Wildcard characters that are a part of a table name are not escaped properly.
- TIMESTAMP with more than four digits in the year is parsed incorrectly.

When using the Amazon Redshift JDBC Driver to parse a TIMESTAMP field that contains a year with more than four digits, the TIMESTAMP is parsed incorrectly.

This has been resolved.

- The TIMESTAMPTZ data type is not supported in the driver.

The Amazon Redshift JDBC Driver does not support the TIMESTAMPTZ data type. As a workaround, use the TIMESTAMP data type.

This has been resolved.

Version 1.2.1

Released November 2016

Enhancements & New Features

Enhanced Connection.isValid()

The Connection.isValid() method can now detect whether the driver is using a valid database connection, even if the database has been shut down unexpectedly. To do this, when isValid() is called, the driver submits an empty database query, and reports the result.

To disable this empty database query, set the DisableIsValidQuery property to True. When this property is set, isValid() may not be able to detect if a connection is invalid because the database has shut down.

Use BEGIN instead of SAVEPOINT

The driver now uses BEGIN statements to start a transaction, instead of SAVEPOINT.

New driver class name

The driver now uses the class names com.amazon.redshift.jdbc.Driver and com.amazon.redshift.jdbc.Datasource to connect to the data store. Previously, different class names were used depending on the version of JDBC that is used.

Resolved Issues

The following issues have been resolved in Amazon Redshift JDBC Driver 1.2.1.

- Multi-byte connection URL or username causes connection failure.
- Closed socket is not properly identified.
- Parser failure when "\" is used.

Using `javax.sql.DataSource` causes an error.

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>.