

# Release Notes

---

## Amazon Redshift JDBC Driver 1.2.1

**Released November 2016**

These release notes provide details of enhancements, features, and known issues in Amazon Redshift JDBC Driver 1.2.1, as well as the version history.

### 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.

## Known Issues

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

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

When using the Amazon Redshift JDBC Driver, calling `Connection.prepare()` 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.

- `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.

- 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.

- 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 1.1.17

**Released June 2016**

These release notes provide details of enhancements, features, and known issues in Amazon Redshift JDBC Driver 1.1.17, as well as the version history.

## Enhancements & New Features

### SSL with client certificate verification

You can use the new connection properties `SSLRootCert`, `SSLCert`, `SSLKey`, and `SSLPassword` to establish a TLS/SSL connection to a proxy server that uses client certificate verification.

### JDBC 4.2 now supported

The driver can now use JDBC 4.2 to connect to a data source.

### Improved support for large rowcount results in the JDBC 4.2 driver

When using the `executeLargeBatch()` or `executeLargeUpdate()` methods, the JDBC 4.2 driver can now report numbers of rows modified that are larger than `INTEGER.MAX_VALUE`.

### Position field added to error messages

Error messages now include a Position field to indicate the location in the query where the problem is occurring.

## Resolved Issues

The following issues have been resolved in Amazon Redshift JDBC Driver 1.1.17.

- Driver returns incorrect metadata values for `getExportedKeys`.
- Driver returns incorrect metadata values for `getCrossReference`.
- Driver removes comments from all queries.
- Driver incorrectly escapes query content that is enclosed in quotation marks (" or ').
- Driver fails to escape braces ({} ) from queries.
- Driver returns an error when attempting to connect to the database with the connection properties set to `AuthMech=REQUIRE` and `SSL=TRUE`.
- In some cases, when the driver attempts to execute many queries in a single statement but stops before all the queries are completed, the driver stops working.

This issue has been resolved. The problem occurred when the statement contained hundreds of queries, and query execution was interrupted at a specific point in the process (for example, if a server error occurred or if the user cancelled the query). The driver is now able to handle these cases successfully.

# Version 1.1.13

Released February 2016

## Enhancements & New Features

### SSLMode connection property

You can now set the SSLMode property to specify whether the driver verifies host names when validating TLS/SSL certificates.

## Resolved Issues

The following issues were resolved in Amazon Redshift JDBC Driver 1.1.13.

- ResultsetMetadata not setting IsNullable correctly.
- Statement reuse causes server-side named statements to not close.
- Batch parameter sets with multiple statements missing results.
- Intermittent return of 0 rows when more exist.
- High CPU utilization potential infinite loop after switching drivers.
- SSL Error writing large queries on slow connections.
- When calling a scalar function and a comment contains another function, the driver returns an error.
- When the file.encoding property is set to UTF-16, the driver fails to connect.
- When calling a REGEXP\_REPLACE function that contains a question mark (?), the driver returns an error.
- When calling a PreparedStatement.executeBatch() method that contains multiple parameter sets and a SQL statement containing multiple commands, the driver omits all parameters except for the last set of parameters.
- Statements and prepared statements containing multiple semicolon-delimited commands are not executed atomically.

The driver now executes these statements atomically so that either all or none of the commands are successfully executed.

- Driver reports incorrect precision for DECIMAL data reported from DatabaseMetaData.getColumns().

## Version 1.1.10

Released October 2015

### Enhancements & New Features

#### Statement.setMaxRows() JDBC method implemented

The driver now supports the Statement.setMaxRows() JDBC method, which is used to limit the number of rows returned over the network to the application.

#### UnknownLength connection property implemented

You can now specify the length and precision to be reported for fields of unknown length by setting the UnknownLength connection property.

#### Support added for UDF functions

The driver is now able to handle the syntax for the new UDF functionality in Redshift.

## Version 1.1.9

Released September 2015

### Resolved Issues

The following issues were resolved in Amazon Redshift JDBC Driver 1.1.9.

- Autocommit behaviour after an explicit "COMMIT" or "ROLLBACK" call is incorrect.

Before, if autocommit was disabled, it would become enabled after explicit "COMMIT" or "ROLLBACK" calls were made. Now, when autocommit is disabled, the driver behaves correctly after the "COMMIT" or "ROLLBACK" keyword is executed.

- Driver returns an assertion error when calling getPrimaryKeys on database metadata.

The driver is now able to get primary keys metadata successfully.

- When executing a prepared statement that uses the `setTimestamp` method on a `Calendar` object, the driver returns an incorrect timestamp in the result set.

The driver now correctly stores and retrieves the timestamp in UTC format.

- When autocommit is disabled and a savepoint has been created, the driver uses the `BEGIN` statement redundantly.

Before, when autocommit is disabled and a savepoint has been created, the driver would send `BEGIN` at the start of the transaction and again when the savepoint is created. Now, the driver only sends `BEGIN` at the start of the transaction.

## Version 1.1.7

Released August 2015

### Enhancements & New Features

#### Updated `DatabaseMetadata` to indicate that the driver can only support one active statement per connection

`DatabaseMetadata` now indicates to applications that the driver can only support one active statement per connection. This update prevents applications from potentially executing multiple statements in a connection and failing.

#### Improved error messaging for region issues

The driver now displays more informative error messages for cases where a Redshift cluster and an S3 bucket are in different regions.

#### Support added for loading the open-source PostgreSQL driver and the Redshift driver into the same class path

You can now enable both drivers to be loaded into the same class path by setting the `OpenSourceSubProtocolOverride` connection attribute to `TRUE`.

### Resolved Issues

The following issues were resolved in Amazon Redshift JDBC Driver 1.1.7.

- In some cases, when using SSL, the driver fails to process very large queries.

This issue has been resolved. The logging of SSL diagnostics has also been improved, and the logs no longer display `FATAL` for non-fatal messages.

- In some cases, when executing a prepared statement, numeric values are not rounded correctly.
- Driver returns an error message stating "(500150) Error setting/closing connection: {0}".

Before, the error occurred whenever the driver encountered an empty array. The driver is now able to handle empty arrays successfully.

- Driver returns unrelated error messages when an incorrect Redshift password is provided.

The driver now displays a SQLException when an incorrect Redshift password is provided.

## Version 1.1.6

Released June 2015

### Resolved Issues

The following issues were resolved in Amazon Redshift JDBC Driver 1.1.6.

- Driver does not support the use of multiple SQL commands in a single prepared statement.

You can now use multiple SQL commands (separated by semicolons) within a single prepared statement.

## Version 1.1.5

Released June 2015

### Resolved Issues

The following issues were resolved in Amazon Redshift JDBC Driver 1.1.5.

- In some cases, the driver returns incorrect result set metadata after executing a prepared statement.

The driver now updates the result set metadata after executing the prepared statement, which ensures that the correct metadata is returned.

- Driver cannot retrieve database metadata from non-selectable tables.

Before, the driver could not retrieve metadata from system tables that cannot be selected, such as shadow tables and TOAST tables. Now, the driver is able to retrieve metadata from these tables.

- When calling `isNullable` on parameter metadata, a null pointer exception occurs.
- In some cases, query cancellation causes a connection exception to occur.

The issue occurred because the way that the query cancellation functionality was implemented in the driver interfered with the thread used for cancelling query statements. The driver is now able to cancel queries properly.

- When calling `getDataType` on a column that uses an unsupported data type, the driver returns "NULL".

Now, the driver now returns "OTHER" for unsupported column data types.

## Version 1.1.2

Released March 2015

### Enhancements & New Features

#### Changed the prefix of the driver log file name

The prefix for the log file that the driver generates has been changed from "PostgreSQL" to "Redshift". Before, the name of the log file was "PostgreSQL\_connection\_<connection number>.log". Now, the name of the log file is "Redshift\_connection\_<connection number>.log".

### Resolved Issues

The following issues were resolved in Amazon Redshift JDBC Driver 1.1.2.

- When working in a connection pool, thread safety issues sometimes occur.

This issue occurred because the driver contained static references that caused thread safety issues to occur whenever there were many concurrent connections. These static references have been removed from the driver, and the thread safety issues have been resolved.

- Warning messages returned during statement execution are not posted on the statement-level warning listener.



In the driver, there are "connection-level" listeners for warnings about the status of a connection, and "statement-level" listeners for warnings about the execution of a query statement. Before, all warnings were posted to the connection-level listener. Now, warnings regarding statement execution are posted to the statement-level listener instead.

- Result set metadata does not return the correct display size for CHAR columns.

This issue occurred because the driver did not recognize that columns of type BPCHAR are of variable length, so BPCHAR columns returned from the server were handled incorrectly. The driver now returns the correct display size for these columns.

## Version 1.1.1

Released March 2015

### Resolved Issues

The following issues were resolved in Amazon Redshift JDBC Driver 1.1.1.

- When executing a query that has nested strings containing escaped characters, the driver sometimes returns a syntax error.

This issue occurred because the driver did not escape single quotation marks (') correctly when parsing queries. Now, the driver correctly parses and executes queries that have nested strings containing escaped characters.

## Version 1.1.0

Version 1.1.0 was the initial release of the Amazon Redshift JDBC Driver.

# 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>.