




Simba Amazon Redshift ODBC Driver

Installation and Configuration Guide

Simba Technologies Inc.

Version 1.3.7
September 27, 2017



Copyright © 2017 Simba Technologies Inc. All Rights Reserved.

Information in this document is subject to change without notice. Companies, names and data used in examples herein are fictitious unless otherwise noted. No part of this publication, or the software it describes, may be reproduced, transmitted, transcribed, stored in a retrieval system, decompiled, disassembled, reverse-engineered, or translated into any language in any form by any means for any purpose without the express written permission of Simba Technologies Inc.

Trademarks

Simba, the Simba logo, SimbaEngine, and Simba Technologies are registered trademarks of Simba Technologies Inc. in Canada, United States and/or other countries. All other trademarks and/or servicemarks are the property of their respective owners.

Contact Us

Simba Technologies Inc.
938 West 8th Avenue
Vancouver, BC Canada
V5Z 1E5

Tel: +1 (604) 633-0008

Fax: +1 (604) 633-0004

www.simba.com

About This Guide

Purpose

The *Simba Amazon Redshift ODBC Driver Installation and Configuration Guide* explains how to install and configure the Simba Amazon Redshift ODBC Driver. The guide also provides details related to features of the driver.

Audience

The guide is intended for end users of the Simba Amazon Redshift ODBC Driver, as well as administrators and developers integrating the driver.

Knowledge Prerequisites

To use the Simba Amazon Redshift ODBC Driver, the following knowledge is helpful:

- Familiarity with the platform on which you are using the Simba Amazon Redshift ODBC Driver
- Ability to use the data source to which the Simba Amazon Redshift ODBC Driver is connecting
- An understanding of the role of ODBC technologies and driver managers in connecting to a data source
- Experience creating and configuring ODBC connections
- Exposure to SQL

Document Conventions

Italics are used when referring to book and document titles.

Bold is used in procedures for graphical user interface elements that a user clicks and text that a user types.

Monospace font indicates commands, source code, or contents of text files.

Note:

A text box with a pencil icon indicates a short note appended to a paragraph.

! Important:

A text box with an exclamation mark indicates an important comment related to the preceding paragraph.

Table of Contents

About the Simba Amazon Redshift ODBC Driver	7
Windows Driver	8
Windows System Requirements	8
Installing the Driver on Windows	8
Creating a Data Source Name on Windows	9
Configuring SSL Verification on Windows	11
Configuring Authentication on Windows	11
Configuring Data Type Options on Windows	18
Configuring Additional Options on Windows	19
Configuring TCP Keepalives on Windows	20
Configuring Logging Options on Windows	22
Verifying the Driver Version Number on Windows	24
macOS Driver	25
macOS System Requirements	25
Installing the Driver on macOS	25
Verifying the Driver Version Number on macOS	26
Linux Driver	27
Linux System Requirements	27
Installing the Driver Using the Tarball Package	27
Configuring the ODBC Driver Manager on Non-Windows Machines	29
Specifying ODBC Driver Managers on Non-Windows Machines	29
Specifying the Locations of the Driver Configuration Files	30
Configuring ODBC Connections on a Non-Windows Machine	32
Creating a Data Source Name on a Non-Windows Machine	32
Configuring a DSN-less Connection on a Non-Windows Machine	35
Configuring SSL Verification on a Non-Windows Machine	37
Configuring IAM Authentication on a Non-Windows Machine	38
Configuring Query Processing Modes on a Non-Windows Machine	40
Configuring TCP Keepalives on a Non-Windows Machine	41
Configuring Logging Options on a Non-Windows Machine	42
Testing the Connection on a Non-Windows Machine	44
Using a Connection String	46
DSN Connection String Example	46

Installation and Configuration Guide

DSN-less Connection String Examples	46
Features	49
Query Processing Modes	49
TCP Keepalives	50
Data Types	50
Security and Authentication	52
Driver Configuration Options	53
Configuration Options Appearing in the User Interface	53
Configuration Options Having Only Key Names	70
Third-Party Trademarks	74
Third-Party Licenses	75

About the Simba Amazon Redshift ODBC Driver

The Simba Amazon Redshift ODBC Driver enables Business Intelligence (BI), analytics, and reporting on data that is stored in Amazon Redshift. The driver complies with the ODBC 3.80 data standard and adds important functionality such as Unicode, as well as 32- and 64-bit support for high-performance computing environments on all platforms.

ODBC is one of the most established and widely supported APIs for connecting to and working with databases. At the heart of the technology is the ODBC driver, which connects an application to the database. For more information about ODBC, see the *Data Access Standards Glossary*: <http://www.simba.com/resources/data-access-standards-library>. For complete information about the ODBC specification, see the *ODBC API Reference*: [http://msdn.microsoft.com/en-us/library/windows/desktop/ms714562\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/ms714562(v=vs.85).aspx).

The Simba Amazon Redshift ODBC Driver is available for Microsoft® Windows®, Linux, and macOS platforms.

The *Installation and Configuration Guide* is suitable for users who are looking to access data residing within Redshift from their desktop environment. Application developers might also find the information helpful. Refer to your application for details on connecting via ODBC.

**Note:**

For information about how to use the driver in various BI tools, see the *Simba ODBC Drivers Quick Start Guide for Windows*: http://cdn.simba.com/docs/ODBC_QuickstartGuide/content/quick_start/intro.htm.

Windows Driver

Windows System Requirements

Install the driver on client machines where the application is installed. Each machine that you install the driver on must meet the following minimum system requirements:

- One of the following operating systems:
 - Windows 7, 8.1, or 10
 - Windows Server 2008 or later
- 75 MB of available disk space
- Visual C++ Redistributable for Visual Studio 2013 installed (with the same bitness as the driver that you are installing).
You can download the installation packages at <https://www.microsoft.com/en-ca/download/details.aspx?id=40784>.

To install the driver, you must have Administrator privileges on the machine.

Installing the Driver on Windows

On 64-bit Windows operating systems, you can execute both 32- and 64-bit applications. However, 64-bit applications must use 64-bit drivers, and 32-bit applications must use 32-bit drivers. Make sure that you use the version of the driver that matches the bitness of the client application:

- `SimbaAmazonRedshiftODBC32.msi` for 32-bit applications
- `SimbaAmazonRedshiftODBC64.msi` for 64-bit applications

You can install both versions of the driver on the same machine.

To install the Simba Amazon Redshift ODBC Driver on Windows:

1. Depending on the bitness of your client application, double-click to run **SimbaAmazonRedshiftODBC32.msi** or **SimbaAmazonRedshiftODBC64.msi**.
2. Click **Next**.
3. Select the check box to accept the terms of the License Agreement if you agree, and then click **Next**.
4. To change the installation location, click **Change**, then browse to the desired folder, and then click **OK**. To accept the installation location, click **Next**.
5. Click **Install**.
6. When the installation completes, click **Finish**.


7. If you received a license file through email, then copy the license file into the `\lib` subfolder of the installation folder you selected above. You must have Administrator privileges when changing the contents of this folder.


Creating a Data Source Name on Windows

Typically, after installing the Simba Amazon Redshift ODBC Driver, you need to create a Data Source Name (DSN).

Alternatively, for information about DSN-less connections, see [Using a Connection String](#) on page 46.

To create a Data Source Name on Windows:

1. Open the ODBC Administrator:
 - If you are using Windows 7 or earlier, click **Start**  > **All Programs** > **Simba Amazon Redshift Driver 1.3** > **ODBC Administrator**.
 - Or, if you are using Windows 8 or later, on the Start screen, type **ODBC administrator**, and then click the **ODBC Administrator** search result.

 **Note:**

Make sure to select the ODBC Data Source Administrator that has the same bitness as the client application that you are using to connect to Redshift.

2. In the ODBC Data Source Administrator, click the **Drivers** tab, and then scroll down as needed to confirm that the Simba Amazon Redshift ODBC Driver appears in the alphabetical list of ODBC drivers that are installed on your system.
3. Choose one:
 - To create a DSN that only the user currently logged into Windows can use, click the **User DSN** tab.
 - Or, to create a DSN that all users who log into Windows can use, click the **System DSN** tab.

 **Note:**

It is recommended that you create a System DSN instead of a User DSN. Some applications load the data using a different user account, and might not be able to detect User DSNs that are created under another user account.

4. Click **Add**.
5. In the Create New Data Source dialog box, select **Simba Amazon Redshift ODBC Driver** and then click **Finish**. The Simba Amazon Redshift ODBC Driver

DSN Setup dialog box opens.

6. In the **Data Source** field, type a name for your DSN.
7. In the **Server** field, type the endpoint of the server hosting the database that you want to access.
8. In the **Port** field, type the number of the TCP port that the server uses to listen for client connections.

 **Note:**

The default port used by Redshift is 5439.

9. In the **Database** field, type the name of the database that you want to access.
10. Set the options in the **Authentication** area to configure standard or IAM authentication. For more information, see [Configuring Authentication on Windows](#) on page 11.
11. Encrypt your credentials by selecting one of the following:
 - If the credentials are used only by the current Windows user, select **Current User Only**.
 - Or, if the credentials are used by all users on the current Windows machine, select **All Users Of This Machine**.
12. To configure client-server verification over SSL, click **SSL Options**. For more information, see [Configuring SSL Verification on Windows](#) on page 11.
13. To configure advanced driver options, click **Additional Options**. For more information, see [Configuring Additional Options on Windows](#) on page 19.
14. To configure logging behavior for the driver, click **Logging Options**. For more information, see [Configuring Logging Options on Windows](#) on page 22.
15. To configure how the driver returns and displays data, click **Data Type Options**. For more information, see [Configuring Data Type Options on Windows](#) on page 18.
16. To test the connection, click **Test**. Review the results as needed, and then click **OK**.

 **Note:**

If the connection fails, then confirm that the settings in the Simba Amazon Redshift ODBC Driver DSN Setup dialog box are correct. Contact your Redshift server administrator as needed.


17. To save your settings and close the Simba Amazon Redshift ODBC Driver DSN Setup dialog box, click **OK**.
18. To close the ODBC Data Source Administrator, click **OK**.

Configuring SSL Verification on Windows

If you are connecting to a Redshift server that has Secure Sockets Layer (SSL) enabled, then you can configure the driver to connect to an SSL-enabled socket. When connecting to a server over SSL, the driver supports identity verification between the client and the server.

To configure SSL verification on Windows:

1. To access the SSL options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **SSL Options**.
2. In the **Authentication Mode** list, select the appropriate SSL mode.

 **Note:**

For information about SSL support in Amazon Redshift, see the topic *Connect Using SSL* in the Amazon Redshift Management Guide at <http://docs.aws.amazon.com/redshift/latest/mgmt/connecting-ssl-support.html#connect-using-ssl>.

3. To use the System Trust Store for SSL certificates, select the **Use System Trust Store** check box.
4. If you selected **Use System Trust Store**, choose one of the following options:
 - To check the validity of the certificate's trust chain, select the **Check Certificate Revocation** checkbox.
 - To accept self-signed certificates, select the **Allow Self-signed Server Certificate** checkbox.
5. To specify an SSL certificate, select the **Enable Custom SSL CA Root Certificate** check box, and then, in the **Path** field, specify the full path to the certificate file.
6. To save your settings and close the dialog box, click **OK**.
7. To save your settings and close the Simba Amazon Redshift ODBC Driver DSN Setup dialog box, click **OK**.

Configuring Authentication on Windows

Redshift databases require authentication. You can configure the driver to provide your credentials and authenticate the connection to the database, or to use a profile or credentials service.

The driver supports the following authentication methods:

- Standard authentication using your database user name and password (see [Using Standard Authentication](#) on page 12)
- IAM authentication using a profile (see [Using an IAM Profile](#) on page 12)
- IAM authentication using IAM credentials (see [Using IAM Credentials](#) on page 14)
- IAM authentication using Active Directory Federation Services (AD FS) (see [Using Active Directory Federation Services \(AD FS\)](#) on page 14)
- IAM authentication using PingFederate service (see [Using PingFederate Service on Windows](#) on page 15)
- IAM authentication using Okta service (see [Using Okta Service](#) on page 16)

For more information on IAM Roles and authentication, see http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2.html.

Follow the appropriate set of steps below to configure authentication for your connection.

Using Standard Authentication

You can configure the driver to authenticate your connection using your Redshift user name and password.

To configure standard authentication on Windows:

1. To access the standard authentication options, open the **ODBC Data Source Administrator** where you created the DSN, select the DSN, and then click **Configure**.
2. If **Auth Type** is not already set to **Standard**, click the drop down and select it.
3. In the **Username** field, type your user name for accessing your Redshift account.
4. In the **Password** field, type the password corresponding to the user name you typed.
5. Encrypt your credentials by selecting one of the following:
 - If the credentials are used only by the current Windows user, select **Current User Only**.
 - Or, if the credentials are used by all users on the current Windows machine, select **All Users Of This Machine**.
6. To save your settings and close the dialog box, click **OK**.

Using an IAM Profile

You can configure the driver to authenticate your connection through IAM authentication using the credentials stored in a chained roles profile.

 **Note:**

- The default location for the credentials file that contains profiles is `~/.aws/Credentials`. The `AWS_SHARED_CREDENTIALS_FILE` environment variable can be used to point to a different credentials file.
- If any of the information requested in the following steps is already a part of the profile you intend to use, that field can be left blank. If you have the default profile configured on your local machine, you only need to set the **Auth Type** to **AWS Profile**.

To configure IAM authentication using a profile on Windows:

1. To access the IAM authentication options, open the **ODBC Data Source Administrator** where you created the DSN, select the DSN, and then click **Configure**.
2. From the **Auth Type** drop-down list select **AWS Profile**.
3. In the **Username** field, type the user name for accessing your IDP Server.
4. In the **Password** field, type the password corresponding to the user name you typed.
5. Encrypt your credentials by selecting one of the following:
 - If the credentials are used only by the current Windows user, select **Current User Only**.
 - Or, if the credentials are used by all users on the current Windows machine, select **All Users Of This Machine**.

 **Note:**

The Cluster ID and Region fields are optional if the Server field is used.


6. In the **Cluster ID** field, type the ID for the Redshift server cluster.
7. In the **Region** field, type the region for the Redshift server cluster.
8. In the **DbUser** field, type the ID you want the Redshift user to use or have.
9. If the ID you entered in the DbUser field doesn't already exist in your Redshift account:
 - To create it, select the **User Auto Create** check box.
 - In the **DbGroups** field, type the names of any user groups you want the DbUser added to, separated by commas .
10. In the **AWS Profile** field, type the name of the profile that contains your connection settings.
11. To save your settings and close the dialog box, click **OK**.

Using IAM Credentials

You can configure the driver to authenticate your connection through IAM authentication using IAM credentials.

To configure IAM authentication using IAM on Windows:

1. To access the IAM authentication options, open the **ODBC Data Source Administrator** where you created the DSN, select the DSN, and then click **Configure**.
2. Click the **Auth Type** drop down and select **AWS IAM Credentials**.

 **Note:**

The Cluster ID and Region fields are optional if the Server field is used.

3. In the **Cluster ID** field, type the ID for the Redshift server cluster.
4. In the **Region** field, type the region for the Redshift server cluster.
5. In the **DbUser** field, type the ID you want the Redshift user to use or have.
6. If the ID you entered in the DbUser field doesn't already exist in your Redshift account:
 - To create it, select the **User Auto Create** check box.
 - In the **DbGroups** field, type the names of any user groups you want the DbUser added to, separated by commas .
7. In the **AccessKeyID** field, type your Redshift access key ID.
8. In the **SecretAccessKey** field, type your Redshift secret key.
9. If you are using an IAM Role, in the **SessionToken** field, type your temporary session token.
10. To save your settings and close the dialog box, click **OK**.

Using Active Directory Federation Services (AD FS)

You can configure the driver to authenticate your connection through IAM authentication using the credentials stored in AD FS .

To configure IAM authentication using AD FS on Windows:

1. To access the IAM authentication options, open the **ODBC Data Source Administrator** where you created the DSN, select the DSN, and then click **Configure**.
2. In the Authentication area, click the **Auth Type** drop down and select **Identity Provider: AD FS**.

3. Choose one of the following options:
 - To log in using Windows Integrated Authentication, leave the **Username** and **Password** fields blank.
 - Or, to log in without using integrated authentication:
 - a. In the **Username** field, type the user name associated with your AD FS account.
 - b. In the **Password** field, type the password associated with your AD FS user name.
4. Encrypt your credentials by selecting one of the following:
 - If the credentials are used only by the current Windows user, select **Current User Only**.
 - Or, if the credentials are used by all users on the current Windows machine, select **All Users Of This Machine**.

 **Note:**

The Cluster ID and Region fields are optional if the Server field is used.

5. In the **Cluster ID** field, type the ID for the Redshift server cluster.
6. In the **Region** field, type the region for the Redshift server cluster.
7. In the **DbUser** field, type the ID you want the Redshift user to use or have.
8. If the ID you entered in the DbUser field doesn't already exist in your Redshift account:
 - To create it, select the **User Auto Create** check box.
 - In the **DbGroups** field, type the names of any user groups you want the DbUser added to, separated by commas .
9. In the **IdP Host** field, type the address of the service host.
10. In the **IdP Port** field, type the port number the service listens at.
11. If you don't want to verify the SSL certificate of the IDP server, select the **SSL Insecure** check box.
12. In the **Preferred Role** field, type the name or ID for the IAM Role you want the user to assume when logged in to Redshift.
13. To save your settings and close the dialog box, click **OK**.

Using PingFederate Service on Windows

You can configure the driver to authenticate your connection through IAM authentication using the credentials stored in the PingFederate service.

To configure IAM authentication using PingFederate service on Windows:

1. To access the IAM authentication options, open the **ODBC Data Source Administrator** where you created the DSN, select the DSN, and then click **Configure**.
2. In the Authentication area, click the **Auth Type** drop down and select **Identity Provider: PingFederate**.
3. In the **Username** field, type the user name associated with your Ping account.
4. In the **Password** field, type the password associated with your Ping user name.

 **Note:**

The Cluster ID and Region fields are optional if the Server field is used.

5. In the **Cluster ID** field, type the ID for the Redshift server cluster.
6. In the **Region** field, type the region for the Redshift server cluster.
7. In the **DbUser** field, type the ID you want the Redshift user to use or have.
8. If the ID you entered in the DbUser field doesn't already exist in your Redshift account:
 - To create it, select the **User Auto Create** check box.
 - In the **DbGroups** field, type the names of any user groups you want the DbUser added to, separated by commas .
9. In the **IdP Host** field, type the address of the service host.
10. In the **IdP Port** field, type the port number the service listens at.
11. In the **Preferred Role** field, type the name or ID for the IAM Role you want the user to assume when logged in to Redshift.
12. If you don't want to verify the SSL certificate of the IDP server, select the **SSL Insecure** check box.
13. To save your settings and close the dialog box, click **OK**.

Using Okta Service

You can configure the driver to authenticate your connection through IAM authentication using the credentials stored in Okta.

To configure IAM authentication using Okta on Windows:

1. To access the IAM authentication options, open the **ODBC Data Source Administrator** where you created the DSN, select the DSN, and then click **Configure**.
2. In the Authentication area, click the **Auth Type** drop down and select **Identity Provider Okta**.
3. In the **Username** field, type the user name associated with your Okta account.
4. In the **Password** field, type the password associated with your Okta user name. If you are using a profile, this may be optional.

5. Encrypt your credentials by selecting one of the following:
 - If the credentials are used only by the current Windows user, select **Current User Only**.
 - Or, if the credentials are used by all users on the current Windows machine, select **All Users Of This Machine**.

 **Note:**

The Cluster ID and Region fields are optional if the Server field is used.

6. In the **Cluster ID** field, type the ID for the Redshift server cluster.
7. In the **Region** field, type the region for the Redshift server cluster.
8. In the **DbUser** field, type the ID you want the Redshift user to use or have.
9. If the ID you entered in the DbUser field doesn't already exist in your Redshift account:
 - To create it, select the **User Auto Create** check box.
 - In the **DbGroups** field, type the names of any user groups you want the DbUser added to, separated by commas .
10. In the **IdP Host** field, type the address of the service host.
11. In the **Preferred Role** field, type the name or ID for the IAM Role you want the user to assume when logged in to Redshift.
12. In the **Okta app ID** field, type the Okta-supplied ID associated with your Redshift application.
13. To save your settings and close the dialog box, click **OK**.

Using Access Keys

You can configure the [driver] to authenticate your connection through IAM authentication using an access key ID and a secret access key. If you are using temporary credentials, then you must also provide a session token.

 **Note:**

Temporary credentials are only valid for a limited amount of time. Contact your Redshift server administrator to get temporary credentials.

To configure IAM authentication using access keys on Windows:

1. To access the IAM authentication options, open the ODBC Data Source Administrator where you created the DSN, select the DSN, and then click **Configure**.
2. In the Authentication area, select the **IAM** tab.
3. Select the **Enable IAM Authentication** check box.

4. In the AccessKeyId field, type your access key ID for connecting to Redshift.
5. In the SecretAccessKey field, type the secret access key corresponding to your access key ID.
6. Encrypt your secret access key by selecting one of the following:
 - If the secret access key is used only by the current Windows user, select **Current User Only**.
 - Or, if the secret access key is used by all users on the current Windows machine, select **All Users Of This Machine**.
7. To save your settings and close the dialog box, click **OK**.

Using the Default Credential Provider Chain

You can configure the [driver] to authenticate your connection through IAM authentication using the credentials stored in the default credential provider chain (DefaultAWSCredentialsProviderChain). For information about the default credential provider chain, see “DefaultAWSCredentialsProviderChain” in the AWS SDK for Java : <http://docs.aws.amazon.com/AWSJavaSDK/latest/javadoc/com/amazonaws/auth/DefaultAWSCredentialsProviderChain.html>

For information about other authentication methods that the driver supports, see [Configuring Authentication on Windows](#).

To configure IAM authentication using a profile on Windows:

1. To access the IAM authentication options, open the ODBC Data Source Administrator where you created the DSN, select the DSN, and then click **Configure**.
2. In the Authentication area, select the **IAM** tab.
3. Select the **Enable IAM Authentication** check box.
4. Make sure that the AccessKeyId, SecretAccessKey, Profile, and Plugin fields are empty.

 **Note:**

If any of these fields contain values, then the driver attempts to authenticate using those values instead.


5. To save your settings and close the dialog box, click **OK**.

Configuring Data Type Options on Windows

You can configure data type options to modify how the driver displays or returns some data types.

To configure data type options on Windows:

1. To access data type options, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **Data Type Options**.
2. To enable the driver to return data as Unicode character types, select the **Use Unicode** check box.

 **Note:**

When the **Use Unicode** check box is selected, the driver does the following:

- Returns SQL_WCHAR instead of SQL_CHAR.
- Returns SQL_WVARCHAR instead of SQL_VARCHAR.
- Returns SQL_WLONGVARCHAR instead of SQL_LONGVARCHAR.

3. To configure the driver to return Boolean columns as SQL_VARCHAR instead of SQL_BIT, select the **Show Boolean Column As String** check box.
4. To configure the driver to return Text columns as SQL_LONGVARCHAR instead of SQL_VARCHAR, select the **Text as LongVarChar** check box.
5. In the **Max Varchar** field, type the maximum data length for VarChar columns.
6. In the **Max LongVarChar** field, type the maximum data length for LongVarChar columns.
7. To save your settings and close the Data Type Configuration dialog box, click **OK**.


Configuring Additional Options on Windows

You can configure additional options to modify the behavior of the driver.

To configure additional options on Windows:

1. To access advanced options, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **Additional Options**.
2. Specify how the driver processes queries by doing one of the following:
 - To return query results one row at a time, select **Single Row Mode**.
 - To return a specific number of rows at a time, select **Use Declare/Fetch** and then, in the **Cache Size** field, type the number of rows.
 - To enable the driver to have multiple queries active on the same connection, select **Use Multiple Statements**. The ODBC application may interleave calls to ODBC statements, but all queries are still sent and executed sequentially.

- To return the entire query result, select **Retrieve Entire Result Into Memory**.

 **Note:**


Use **Single Row Mode** if you plan to query large results and you do not want to retrieve the entire result into memory. Disabling **Single Row Mode** increases performance, but can result in out-of-memory errors.

3. To configure the driver to have only one active query at a time per connection, select **Enforce Single Statement Mode**.
4. To configure the driver to recognize table type information from the data source, select the **Enable Table Types** checkbox. For more information, see [Enable Table Types](#) on page 57.
5. If you are connecting through a proxy server, then select the **Enable HTTP Proxy Connection** check box and then do the following:
 - a. In the **Proxy Server** field, type the host name or IP address of the proxy server.
 - b. In the **Proxy Port** field, type the number of the TCP port that the proxy server uses to listen for client connections.
6. To save your settings and close the Advanced Options dialog box, click **OK**.
7. To save your settings and close the Simba Amazon Redshift ODBC Driver DSN Setup dialog box, click **OK**.

Configuring TCP Keepalives on Windows

By default, the Simba Amazon Redshift ODBC Driver is configured to use TCP keepalives to prevent connections from timing out. Settings such as how frequently the driver sends TCP keepalive packets are based on the operating system defaults. You can configure the TCP keepalive settings or disable the feature by modifying the appropriate values in the Windows Registry.

To configure TCP keepalives on Windows:

1. Choose one:
 - If you are using Windows 7 or earlier, click **Start** , then type **regedit** in the Search field, and then click **regedit.exe** in the search results.
 - Or, if you are using Windows 8 or later, on the Start screen, type **regedit**, and then click the **regedit** search result.
2. Select the appropriate registry key for the bitness of your driver:
 - If you are using the 32-bit driver on a 64-bit machine, then select the following registry key, where *[YourDSN]* is the DSN for which you want to configure keepalives:

HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ODBC\ODBC.INI[YourDSN]

- Otherwise, select the following registry key, where [YourDSN] is the DSN for which you want to configure keepalives:

HKEY_LOCAL_MACHINE\SOFTWARE\ODBC\ODBC.INI[YourDSN]

3. To specify the interval of inactivity before the driver sends a TCP keepalive packet, configure the **KeepAliveldle** value by doing the following:
 - a. If the **KeepAliveldle** value does not already exist, create it. Select **Edit > New > String Value**, type **KeepAliveldle** as the name of the value, and then press **Enter**.
 - b. Select the **KeepAliveldle** value, and then Select **Edit > Modify**.
 - c. In the Edit String dialog box, in the **Value Data** field, type the number of seconds of inactivity before the driver sends a TCP keepalive packet.

 **Note:**

To use the system default, in the **Value Data** field, type **0**.

- d. Click **OK**.
4. To specify the number of TCP keepalive packets that can be lost before the connection is considered broken, configure the **KeepAliveCount** value. To do this, follow the procedure above, but type **KeepAliveCount** for the value name, and in the **Value Data** field, type the number of keepalive packets that can be lost.

 **Note:**

To use the system default, in the **Value Data** field, type **0**.


5. To specify the interval of time between each retransmission of a keepalive packet, configure the **KeepAliveInterval** value. To do this, follow the procedure above, but type **KeepAliveInterval** for the value name, and in the **Value Data** field, type the number of seconds to wait between each retransmission.

 **Note:**

To use the system default, in the **Value Data** field, type **0**.

6. Close the Registry Editor.

To disable TCP keepalives:

1. Choose one:
 - If you are using Windows 7 or earlier, click **Start** , then type **regedit** in the Search field, and then click **regedit.exe** in the search results.
 - Or, if you are using Windows 8 or later, on the Start screen, type **regedit**, and then click the **regedit** search result.

2. Select the appropriate registry key for the bitness of your driver:
 - If you are using the 32-bit driver on a 64-bit machine, then select the following registry key, where *[YourDSN]* is the DSN for which you want to configure keepalives:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\ODBC\ODBC.INI\YourDSN
```
 - Otherwise, select the following registry key, where *[YourDSN]* is the DSN for which you want to configure keepalives:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ODBC\ODBC.INI\YourDSN
```
3. If the **KeepAlive** value does not already exist, create it. Select **Edit > New > String Value**, then type **KeepAlive** as the name of the value, and then press **Enter**.
4. Select the **KeepAlive** value, and then click **Edit > Modify**.
5. In the Edit String dialog box, in the **Value Data** field, type **0**.
6. Click **OK**.
7. Close the Registry Editor.

 **Note:**

To enable TCP keepalives after disabling them, set `KeepAlive` to 1.

Configuring Logging Options on Windows

To help troubleshoot issues, you can enable logging. In addition to functionality provided in the Simba Amazon Redshift ODBC Driver, the ODBC Data Source Administrator provides tracing functionality.

! Important:

Only enable logging or tracing long enough to capture an issue. Logging or tracing decreases performance and can consume a large quantity of disk space.

The settings for logging apply to every connection that uses the Simba Amazon Redshift ODBC Driver, so make sure to disable the feature after you are done using it.

To enable driver logging on Windows:

1. To access logging options, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **Logging Options**.
2. From the **Log Level** drop-down list, select the logging level corresponding to the amount of information that you want to include in log files:

Logging Level	Description
OFF	Disables all logging.
FATAL	Logs severe error events that lead the driver to abort.
ERROR	Logs error events that might allow the driver to continue running.
WARNING	Logs events that might result in an error if action is not taken.
INFO	Logs general information that describes the progress of the driver.
DEBUG	Logs detailed information that is useful for debugging the driver.
TRACE	Logs all driver activity.

3. In the **Log Path** field, specify the full path to the folder where you want to save log files.
4. Click **OK**.
5. Restart your ODBC application to make sure that the new settings take effect.

The Simba Amazon Redshift ODBC Driver produces two log files at the location you specify in the Log Path field, where *[DriverName]* is the name of the driver:

- A *[DriverName]_driver.log* file that logs driver activity that is not specific to a connection.
- A *[DriverName]_connection_[Number].log* for each connection made to the database, where *[Number]* is a number that identifies each log file. This file logs driver activity that is specific to the connection.

If you enable the `UseLogPrefix` connection property, the driver prefixes the log file name with the user name associated with the connection and the process ID of the application through which the connection is made. For more information, see [UseLogPrefix](#) on page 73.

To disable driver logging on Windows:


1. Open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **Logging Options**.
2. From the **Log Level** drop-down list, select **LOG_OFF**.


3. Click **OK**.
4. Restart your ODBC application to make sure that the new settings take effect.

Verifying the Driver Version Number on Windows

If you need to verify the version of the Simba Amazon Redshift ODBC Driver that is installed on your Windows machine, you can find the version number in the ODBC Data Source Administrator.

To verify the driver version number on Windows:

1. Open the ODBC Administrator:
 - If you are using Windows 7 or earlier, click **Start**  > **All Programs** > **Simba Amazon Redshift Driver 1.3** > **ODBC Administrator**.
 - Or, if you are using Windows 8 or later, on the Start screen, type **ODBC administrator**, and then click the **ODBC Administrator** search result.

 **Note:**

Make sure to select the ODBC Data Source Administrator that has the same bitness as the client application that you are using to connect to Redshift.

2. Click the **Drivers** tab and then find the Simba Amazon Redshift ODBC Driver in the list of ODBC drivers that are installed on your system. The version number is displayed in the **Version** column.

macOS Driver

macOS System Requirements

Install the driver on client machines where the application is installed. Each machine that you install the driver on must meet the following minimum system requirements:

- macOS version 10.9, 10.10, or 10.11
- 215 MB of available disk space
- iODBC 3.52.7 or later

Installing the Driver on macOS

The Simba Amazon Redshift ODBC Driver is available for macOS as a `.dmg` file named `SimbaAmazonRedshiftODBC.dmg`. The driver supports both 32- and 64-bit client applications.

To install the Simba Amazon Redshift ODBC Driver on macOS:

1. Double-click **SimbaAmazonRedshiftODBC.dmg** to mount the disk image.
2. Double-click **SimbaAmazonRedshiftODBC.pkg** to run the installer.
3. In the installer, click **Continue**.
4. On the Software License Agreement screen, click **Continue**, and when the prompt appears, click **Agree** if you agree to the terms of the License Agreement.
5. Optionally, to change the installation location, click **Change Install Location**, then select the desired location, and then click **Continue**.

 **Note:**

By default, the driver files are installed in the `/Library/simba/amazonredshiftdbc` directory.

6. To accept the installation location and begin the installation, click **Install**.
7. When the installation completes, click **Close**.
8. If you received a license file through email, then copy the license file into the `/lib` subfolder in the driver installation directory. You must have root privileges when changing the contents of this folder.

For example, if you installed the driver to the default location, you would copy the license file into the `/Library/simba/amazonredshiftdbc/lib` folder.

Next, configure the environment variables on your machine to make sure that the ODBC driver manager can work with the driver. For more information, see [Configuring the ODBC Driver Manager on Non-Windows Machines](#) on page 29.

Verifying the Driver Version Number on macOS

If you need to verify the version of the Simba Amazon Redshift ODBC Driver that is installed on your macOS machine, you can query the version number through the Terminal.

To verify the driver version number on macOS:

- At the Terminal, run the following command:

```
pkgutil --info com.simba.redshiftdbc
```

The command returns information about the Simba Amazon Redshift ODBC Driver that is installed on your machine, including the version number.

Linux Driver

The Linux driver is available as an RPM file and as a tarball package.

Linux System Requirements

Install the driver on client machines where the application is installed. Each machine that you install the driver on must meet the following minimum system requirements:

- One of the following distributions:
 - Red Hat® Enterprise Linux® (RHEL) 6 or 7
 - CentOS 6 or 7
 - SUSE Linux Enterprise Server (SLES) 11 or 12
 - Debian 7 or 8
 - Ubuntu 14.04 or 16.04
- Distribution must support C++11
- GCC 4.9 or later
- 150 MB of available disk space
- One of the following ODBC driver managers installed:
 - iODBC 3.52.7 or later
 - unixODBC 2.3.0 or later

To install the driver, you must have root access on the machine.

Installing the Driver Using the Tarball Package

The Simba Amazon Redshift ODBC Driver is available as a tarball package named `SimbaRedshiftODBC-[Version].[Release]-Linux.tar.gz`, where `[Version]` is the version number of the driver and `[Release]` is the release number for this version of the driver. The package contains both the 32-bit and 64-bit versions of the driver.

On 64-bit editions of Linux, you can execute both 32- and 64-bit applications. However, 64-bit applications must use 64-bit drivers, and 32-bit applications must use 32-bit drivers. Make sure that you use the version of the driver that matches the bitness of the client application. You can install both versions of the driver on the same machine.

To install the Simba Amazon Redshift ODBC Driver using the tarball package:

1. Log in as the root user, and then navigate to the folder containing the tarball package.
2. Run the following command to extract the package and install the driver:

```
tar --directory=/opt -zxvf [TarballName]
```

Where *[TarballName]* is the name of the tarball package containing the driver.

The Simba Amazon Redshift ODBC Driver files are installed in the `opt/simba/amazonredshiftdbc` directory.

3. If you received a license file through email, then copy the license file into the `opt/simba/amazonredshiftdbc/lib/32` or `opt/simba/amazonredshiftdbc/lib/64` folder, depending on the version of the driver that you installed. You must have root privileges when changing the contents of this folder.

Next, configure the environment variables on your machine to make sure that the ODBC driver manager can work with the driver. For more information, see [Configuring the ODBC Driver Manager on Non-Windows Machines](#) on page 29.

Configuring the ODBC Driver Manager on Non-Windows Machines

To make sure that the ODBC driver manager on your machine is configured to work with the Simba Amazon Redshift ODBC Driver, do the following:

- Set the library path environment variable to make sure that your machine uses the correct ODBC driver manager. For more information, see [Specifying ODBC Driver Managers on Non-Windows Machines](#) on page 29.
- If the driver configuration files are not stored in the default locations expected by the ODBC driver manager, then set environment variables to make sure that the driver manager locates and uses those files. For more information, see [Specifying the Locations of the Driver Configuration Files](#) on page 30.

After configuring the ODBC driver manager, you can configure a connection and access your data store through the driver. For more information, see [Configuring ODBC Connections on a Non-Windows Machine](#) on page 32.

Specifying ODBC Driver Managers on Non-Windows Machines

You need to make sure that your machine uses the correct ODBC driver manager to load the driver. To do this, set the library path environment variable.

macOS

If you are using a macOS machine, then set the `DYLD_LIBRARY_PATH` environment variable to include the paths to the ODBC driver manager libraries. For example, if the libraries are installed in `/usr/local/lib`, then run the following command to set `DYLD_LIBRARY_PATH` for the current user session:

```
export DYLD_LIBRARY_PATH=$DYLD_LIBRARY_PATH:/usr/local/lib
```

For information about setting an environment variable permanently, refer to the macOS shell documentation.

Linux

If you are using a Linux machine, then set the `LD_LIBRARY_PATH` environment variable to include the paths to the ODBC driver manager libraries. For example, if the libraries are installed in `/usr/local/lib`, then run the following command to set `LD_LIBRARY_PATH` for the current user session:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/lib
```

For information about setting an environment variable permanently, refer to the Linux shell documentation.

Specifying the Locations of the Driver Configuration Files

By default, ODBC driver managers are configured to use hidden versions of the `odbc.ini` and `odbcinst.ini` configuration files (named `.odbc.ini` and `.odbcinst.ini`) located in the home directory, as well as the `simba.amazonredshiftoDBC.ini` file in the `lib` subfolder of the driver installation directory. If you store these configuration files elsewhere, then you must set the environment variables described below so that the driver manager can locate the files.

If you are using iODBC, do the following:

- Set `ODBCINI` to the full path and file name of the `odbc.ini` file.
- Set `ODBCINSTINI` to the full path and file name of the `odbcinst.ini` file.
- Set `SIMBAAMAZONREDSHIFTODBCINI` to the full path and file name of the `simba.amazonredshiftoDBC.ini` file.

 **Note:**

If you acquired the driver from a vendor other than Simba, you need to replace `SIMBA` with the name of your vendor.

If you are using unixODBC, do the following:

- Set `ODBCINI` to the full path and file name of the `odbc.ini` file.
- Set `ODBCSYSINI` to the full path of the directory that contains the `odbcinst.ini` file.
- Set `SIMBAAMAZONREDSHIFTODBCINI` to the full path and file name of the `simba.amazonredshiftoDBC.ini` file.

 **Note:**

If you acquired the driver from a vendor other than Simba, you need to replace `SIMBA` with the name of your vendor.

For example, if your `odbc.ini` and `odbcinst.ini` files are located in `/usr/local/odbc` and your `simba.amazonredshiftoDBC.ini` file is located in `/etc`, then set the environment variables as follows:

For iODBC:

```
export ODBCINI=/usr/local/odbc/odbc.ini
export ODBCINSTINI=/usr/local/odbc/odbcinst.ini
export
SIMBAAMAZONREDSHIFTODBCINI=/etc/simba.amazonredshiftodbc.ini
```

For unixODBC:

```
export ODBCINI=/usr/local/odbc/odbc.ini
export ODBCYSINI=/usr/local/odbc
export
SIMBAAMAZONREDSHIFTODBCINI=/etc/simba.amazonredshiftodbc.ini
```

To locate the `simba.amazonredshiftodbc.ini` file, the driver uses the following search order:

1. If the `SIMBAAMAZONREDSHIFTODBCINI` environment variable is defined, then the driver searches for the file specified by the environment variable.
2. The driver searches the directory that contains the driver library files for a file named `simba.amazonredshiftodbc.ini`.
3. The driver searches the current working directory of the application for a file named `simba.amazonredshiftodbc.ini`.
4. The driver searches the home directory for a hidden file named `.simba.amazonredshiftodbc.ini` (prefixed with a period).
5. The driver searches the `/etc` directory for a file named `simba.amazonredshiftodbc.ini`.

Configuring ODBC Connections on a Non-Windows Machine

The following sections describe how to configure ODBC connections when using the Simba Amazon Redshift ODBC Driver on non-Windows platforms:

- [Creating a Data Source Name on a Non-Windows Machine](#) on page 32
- [Configuring a DSN-less Connection on a Non-Windows Machine](#) on page 35
- [Configuring SSL Verification on a Non-Windows Machine](#) on page 37
- [Configuring Query Processing Modes on a Non-Windows Machine](#) on page 40
- [Configuring TCP Keepalives on a Non-Windows Machine](#) on page 41
- [Configuring Logging Options on a Non-Windows Machine](#) on page 42
- [Testing the Connection on a Non-Windows Machine](#) on page 44

Creating a Data Source Name on a Non-Windows Machine

When connecting to your data store using a DSN, you only need to configure the `odbc.ini` file. Set the properties in the `odbc.ini` file to create a DSN that specifies the connection information for your data store. For information about configuring a DSN-less connection instead, see [Configuring a DSN-less Connection on a Non-Windows Machine](#) on page 35.

If your machine is already configured to use an existing `odbc.ini` file, then update that file by adding the settings described below. Otherwise, copy the `odbc.ini` file from the `Setup` subfolder in the driver installation directory to the home directory, and then update the file as described below.

To create a Data Source Name on a non-Windows machine:

1. In a text editor, open the `odbc.ini` configuration file.

 **Note:**

If you are using a hidden copy of the `odbc.ini` file, you can remove the period (.) from the start of the file name to make the file visible while you are editing it.

2. In the `[ODBC Data Sources]` section, add a new entry by typing a name for the DSN, an equal sign (=), and then the name of the driver.

For example, on a macOS machine:


```
[ODBC Data Sources]
Sample DSN=Simba Amazon Redshift ODBC Driver
```

As another example, for a 32-bit driver on a Linux machine:

```
[ODBC Data Sources]
Sample DSN=Simba Amazon Redshift ODBC Driver 32-bit
```

3. Create a section that has the same name as your DSN, and then specify configuration options as key-value pairs in the section:
 - a. Set the `Driver` property to the full path of the driver library file that matches the bitness of the application.

For example, on a macOS machine:

```
Driver=/Library/simba/amazonredshiftodbc/lib/libamazonredshiftodbc_sbu.dylib
```

As another example, for a 32-bit driver on a Linux machine:

```
Driver=/opt/simba/amazonredshiftodbc/lib/32/libamazonredshiftodbc_sb32.so
```

- b. Set the `Server` property to the endpoint of the server, and then set the `Port` property to the number of the TCP port that the server uses to listen for client connections.

For example:

```
Server=testserver.abcabcabcabc.us-west-2.redshift.amazonaws.com
Port=5439
```

- c. Set the `Database` property to the name of the database that you want to access.

For example:

```
Database=TestDB
```


- d. To configure authentication, set the `UID` property to an appropriate user name for accessing the Redshift server, and set the `PWD` property to the password corresponding to the user name you provided.

For example:

```
UID=simba
```

```
PWD=simba123
```

- e. Optionally, modify how the driver runs queries and retrieves results into memory. For more information, see [Configuring Query Processing Modes on a Non-Windows Machine](#) on page 40.
 - f. Optionally, modify the TCP keepalive settings that the driver uses to prevent connections from timing out. For more information, see [Configuring TCP Keepalives on a Non-Windows Machine](#) on page 41.
 - g. Optionally, set additional key-value pairs as needed to specify other optional connection settings. For detailed information about all the configuration options supported by the Simba Amazon Redshift ODBC Driver, see [Driver Configuration Options](#) on page 53.
4. Save the `odbc.ini` configuration file.

 **Note:**

If you are storing this file in its default location in the home directory, then prefix the file name with a period (.) so that the file becomes hidden. If you are storing this file in another location, then save it as a non-hidden file (without the prefix), and make sure that the `ODBCINI` environment variable specifies the location. For more information, see [Specifying the Locations of the Driver Configuration Files](#) on page 30.

For example, the following is an `odbc.ini` configuration file for macOS containing a DSN that connects to Redshift:

```
[ODBC Data Sources]
Sample DSN=Simba Amazon Redshift ODBC Driver
[Sample DSN]
Driver=/Library/simba/amazonredshiftdbc/lib/libamazonredshiftdbc_sbu.dylib
Host=192.168.222.160
Port=5432
Database=TestDB
UID=simba
PWD=simba123
```

As another example, the following is an `odbc.ini` configuration file for a 32-bit driver on a Linux machine, containing a DSN that connects to Redshift:

```
[ODBC Data Sources]
Sample DSN=Simba Amazon Redshift ODBC Driver 32-bit
[Sample DSN]
```

```
Driver=/opt/simba/amazonredshiftdbc/lib/32/libamazonredshif
todbcb_sb32.so
Host=192.168.222.160
Port=5432
Database=TestDB
UID=simba
PWD=simba123
```

You can now use the DSN in an application to connect to the data store.

Configuring a DSN-less Connection on a Non-Windows Machine

To connect to your data store through a DSN-less connection, you need to define the driver in the `odbcinst.ini` file and then provide a DSN-less connection string in your application.

If your machine is already configured to use an existing `odbcinst.ini` file, then update that file by adding the settings described below. Otherwise, copy the `odbcinst.ini` file from the `Setup` subfolder in the driver installation directory to the home directory, and then update the file as described below.

To define a driver on a non-Windows machine:

1. In a text editor, open the `odbcinst.ini` configuration file.

Note:

If you are using a hidden copy of the `odbcinst.ini` file, you can remove the period (.) from the start of the file name to make the file visible while you are editing it.

2. In the `[ODBC Drivers]` section, add a new entry by typing a name for the driver, an equal sign (=), and then `Installed`.

For example:

```
[ODBC Drivers]
Simba Amazon Redshift ODBC Driver=Installed
```

3. Create a section that has the same name as the driver (as specified in the previous step), and then specify the following configuration options as key-value pairs in the section:
 - a. Set the `Driver` property to the full path of the driver library file that matches the bitness of the application.

For example, on a macOS machine:

```
Driver=/Library/simba/amazonredshiftodbc/lib/libamazonredshiftodbc_sbu.dylib
```

As another example, for a 32-bit driver on a Linux machine:

```
Driver=/opt/simba/amazonredshiftodbc/lib/32/libamazonredshiftodbc_sb32.so
```

- b. Optionally, set the `Description` property to a description of the driver.

For example:

```
Description=Simba Amazon Redshift ODBC Driver
```

4. Save the `odbcinst.ini` configuration file.

 **Note:**

If you are storing this file in its default location in the home directory, then prefix the file name with a period (.) so that the file becomes hidden. If you are storing this file in another location, then save it as a non-hidden file (without the prefix), and make sure that the `ODBCINSTINI` or `ODBCSYSINI` environment variable specifies the location. For more information, see [Specifying the Locations of the Driver Configuration Files](#) on page 30.

For example, the following is an `odbcinst.ini` configuration file for macOS:

```
[ODBC Drivers]
Simba Amazon Redshift ODBC Driver=Installed
[Simba Amazon Redshift ODBC Driver]
Description=Simba Amazon Redshift ODBC Driver
Driver=/Library/simba/amazonredshiftodbc/lib/libamazonredshiftodbc_sbu.dylib
```

As another example, the following is an `odbcinst.ini` configuration file for both the 32- and 64-bit drivers on Linux:

```
[ODBC Drivers]
Simba Amazon Redshift ODBC Driver 32-bit=Installed
Simba Amazon Redshift ODBC Driver 64-bit=Installed
[Simba Amazon Redshift ODBC Driver 32-bit]
Description=Simba Amazon Redshift ODBC Driver (32-bit)
Driver=/opt/simba/amazonredshiftodbc/lib/32/libamazonredshiftodbc_sb32.so
```

```
[Simba Amazon Redshift ODBC Driver 64-bit]
Description=Simba Amazon Redshift ODBC Driver (64-bit)
Driver=/opt/simba/amazonredshiftoDBC/lib/64/libamazonredshif
toDBC_sb64.so
```

You can now connect to your data store by providing your application with a connection string where the `Driver` property is set to the driver name specified in the `odbcinst.ini` file, and all the other necessary connection properties are also set. For more information, see "DSN-less Connection String Examples" in [Using a Connection String](#) on page 46.

For instructions about configuring specific connection features, see the following:

- [Configuring Query Processing Modes on a Non-Windows Machine](#) on page 40
- [Configuring TCP Keepalives on a Non-Windows Machine](#) on page 41

For detailed information about all the connection properties that the driver supports, see [Driver Configuration Options](#) on page 53.

Configuring SSL Verification on a Non-Windows Machine

If you are connecting to a Redshift server that has Secure Sockets Layer (SSL) enabled, then you can configure the driver to connect to an SSL-enabled socket. When connecting to a server over SSL, the driver supports identity verification between the client and the server.

You can set the connection properties described below in a connection string or in a DSN (in the `odbc.ini` file). Settings in the connection string take precedence over settings in the DSN.

To configure SSL verification on a non-Windows machine:

1. Set the `SSLMode` property to the appropriate SSL mode.

 **Note:**

For information about SSL support in Amazon Redshift, see the topic *Connect Using SSL* in the Amazon Redshift Management Guide at <http://docs.aws.amazon.com/redshift/latest/mgmt/connecting-ssl-support.html#connect-using-ssl>.

2. To use the System Trust Store for SSL certificates, set `UseSystemTrustStore` to 1.

3. If you set `UseSystemTrustStore` above, choose one of the following options:
 - To check the validity of the certificate's trust chain, the `CheckCertRevocation` to 1.
 - To accept self-signed certificates, check the `CheckCertRevocation` to 1.
4. To specify an SSL certificate, set the `SSLCertPath` property to the full path and file name of the certificate file.

Configuring IAM Authentication on a Non-Windows Machine

If you are connecting to a Redshift server using IAM authentication, then you can configure the driver accordingly.

You can set the connection properties described below in a connection string or in a DSN (in the `odbc.ini` file). Settings in the connection string take precedence over settings in the DSN.

For more information on IAM Roles and authentication, see http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2.html.

Using IAM Credentials

To configure IAM Authentication using IAM credentials on a non-Windows machine:

1. Set the `iam` property to 1.
2. Set the `ClusterID` property to the name of the cluster you are connecting to.
3. Set the `Region` property to the region your cluster is in.
4. Set the `DbUser` property to the name of your Redshift user.
5. If the user doesn't currently exist in the Redshift instance
 - Set the `AutoCreate` property to `true`
 - Set the `DbGroups` property to a list of any groups you want the user to be a member of, separated by commas.
6. Set the `AccessKeyId` property to the access key of the user or role you are using to authenticate into Redshift.
7. Set the `SecretAccessKey` property to the secret key associated with your user or role.
8. If you are using an IAM Role to authenticate, set the `SessionToken` property to the temporary token for your Redshift instance.

Using an IAM Profile

Note:

- The default location for the credentials file that contains profiles is `~/.aws/Credentials`. The `AWS_SHARED_CREDENTIALS_FILE` environment variable can be used to point to a different credentials file.
- If any of the information requested in the following steps is already a part of the profile you intend to use, that field can be left blank. If you have the default profile configured on your local machine, you only need to set the **Auth Type to AWS Profile**.

To configure IAM Authentication using an AWS profile on a non-Windows machine:

1. Set the `iam` property to `1`.
2. Set `User` property to your Redshift user name.
3. Set `Password` property to your Redshift user name password.
4. Set the `ClusterID` property to the name of the cluster you are connecting to.
5. Set the `Region` property to the region your cluster is in.
6. Set the `DbUser` property to the name of your Redshift user.
7. If the user doesn't currently exist in the Redshift instance
 - Set the `AutoCreate` property to `true`
 - Set the `DbGroups` property to a list of any groups you want the user to be a member of, separated by commas.
8. Set the `Profile` property to the name of the profile that contains your credentials.

Using an Identity Provider Service

To configure IAM Authentication using an identity provider service on a non-Windows machine:

1. Set the `iam` property to `1`.
2. Set `User` property to your the user name of your identity provider service account.
3. Set `Password` property to the password of your identity provider service account.
4. Set the `ClusterID` property to the name of the cluster you are connecting to.
5. Set the `Region` property to the region your cluster is in.
6. Set the `DbUser` property to the name of your Redshift user.

7. If the user doesn't currently exist in the Redshift instance
 - Set the `AutoCreate` property to `true`
 - Set the `DbGroups` property to a list of any groups you want the user to be a member of, separated by commas.
8. Set the `preferred_role` property to the role you want the user to have when logged in to the data source.
9. If you are using Okta, set the `app_id` property to the Okta-provided unique ID associated with your Redshift applicaiton.

Configuring Query Processing Modes on a Non-Windows Machine

To optimize driver performance, you can modify how the driver runs queries and retrieves results into memory. For example, you can configure the driver to return entire query results into memory all at once, or one row at a time. Use a query processing mode that prevents queries from consuming too much memory, based on the expected result size of your queries and the specifications of your system.

Note:

Use Single Row Mode if you plan to query large results and you do not want to retrieve the entire result into memory. Using the other query processing modes increases performance, but can result in out-of-memory errors.

You can set the connection properties described below in a connection string or in a DSN (in the `odbc.ini` file). Settings in the connection string take precedence over settings in the DSN.

Enabling Single Row Mode

You can configure the driver to return query results one row at a time.

To enable Single Row Mode:

1. Set the `SingleRowMode` property to `1`.
2. Make sure that the `UseDeclareFetch` property is set to `0` or not set.

Enabling Declare/Fetch Mode

You can configure the driver to return a specific number of rows at a time.

To enable Declare/Fetch Mode:

1. Set the `UseDeclareFetch` property to 1.
2. Set the `Fetch` property to the number of rows that the driver returns at a time.

Enabling Retrieve Entire Result Mode

You can configure the driver to return entire query results into memory.

To enable Retrieve Entire Result Mode:

- Make sure that the `SingleRowMode`, `UseDeclareFetch`, and `UseMultipleStatements` properties are set to 0 or not set.

Enabling Multiple Statements Mode

You can enable the driver to have multiple queries active on the same connection. The ODBC application may interleave calls to ODBC statements, but all queries are still sent and executed sequentially. When using this mode, the driver returns all the query results into memory.

To enable Multiple Statements Mode:

1. Set the `UseMultipleStatements` property to 1.
2. Make sure that the `SingleRowMode` and `UseDeclareFetch` properties are set to 0 or not set.

Configuring TCP Keepalives on a Non-Windows Machine

By default, the Simba Amazon Redshift ODBC Driver is configured to use TCP keepalives to prevent connections from timing out. Settings such as how frequently the driver sends TCP keepalive packets are based on the operating system defaults.

You can set the connection properties described below in a connection string or in a DSN (in the `odbc.ini` file). Settings in the connection string take precedence over settings in the DSN.

To configure TCP keepalives on a non-Windows machine:

1. Set the `KeepAliveIdle` property to the number of seconds of inactivity before the driver sends a TCP keepalive packet.
2. Set the `KeepAliveCount` property to the number of keepalive packets that can be lost before the connection is considered broken.

3. Set the `KeepAliveInterval` property to the number of seconds to wait before each retransmission of a keepalive packet.

 **Note:**

To use the system default for `KeepAliveIdle`, `KeepAliveCount`, or `KeepAliveInterval`, set the property to 0.

To disable TCP keepalives:

- Set the `KeepAlive` property to 0.

 **Note:**

To enable TCP keepalives after disabling them, remove the `KeepAlive` property or set it to 1.

Configuring Logging Options on a Non-Windows Machine

To help troubleshoot issues, you can enable logging in the driver.

! Important:

Only enable logging long enough to capture an issue. Logging decreases performance and can consume a large quantity of disk space.

Logging is configured through driver-wide settings in the `simba.amazonredshiftodbc.ini` file, which apply to all connections that use the driver.


To enable logging on a non-Windows machine:

1. Open the `simba.amazonredshiftodbc.ini` configuration file in a text editor.
2. To specify the level of information to include in log files, set the `LogLevel` property to one of the following numbers:

LogLevel Value	Description
0	Disables all logging.

LogLevel Value	Description
1	Logs severe error events that lead the driver to abort.
2	Logs error events that might allow the driver to continue running.
3	Logs events that might result in an error if action is not taken.
4	Logs general information that describes the progress of the driver.
5	Logs detailed information that is useful for debugging the driver.
6	Logs all driver activity.

3. Set the `LogPath` key to the full path to the folder where you want to save log files.
4. Set the `LogFileCount` key to the maximum number of log files to keep.

 **Note:**

After the maximum number of log files is reached, each time an additional file is created, the driver deletes the oldest log file.

5. Set the `LogFileSize` key to the maximum size of each log file in megabytes (MB).

 **Note:**

After the maximum file size is reached, the driver creates a new file and continues logging.

6. Optionally, to prefix the log file name with the user name and process ID associated with the connection, set the `UseLogPrefix` property to 1.
7. Save the `simba.amazonredshiftodbc.ini` configuration file.
8. Restart your ODBC application to make sure that the new settings take effect.

The Simba Amazon Redshift ODBC Driver produces two log files at the location you specify using the `LogPath` key, where `[DriverName]` is the name of the driver:

- A `[DriverName]_driver.log` file that logs driver activity that is not specific to a connection.
- A `[DriverName]_connection_[Number].log` for each connection made to the database, where `[Number]` is a number that identifies each log file. This file logs driver activity that is specific to the connection.

If you set the `UseLogPrefix` property to 1, then each file name is prefixed with `[UserName]_[ProcessID]_`, where `[UserName]` is the user name associated with the connection and `[ProcessID]` is the process ID of the application through which the connection is made.

To disable logging on a non-Windows machine:

1. Open the `simba.amazonredshiftodbc.ini` configuration file in a text editor.
2. Set the `LogLevel` key to 0.
3. Save the `simba.amazonredshiftodbc.ini` configuration file.
4. Restart your ODBC application to make sure that the new settings take effect.

Testing the Connection on a Non-Windows Machine

To test the connection, you can use an ODBC-enabled client application. For a basic connection test, you can also use the test utilities that are packaged with your driver manager installation. For example, the iODBC driver manager includes simple utilities called `iodbctest` and `iodbctestw`. Similarly, the unixODBC driver manager includes simple utilities called `isql` and `iusql`.

Using the iODBC Driver Manager

You can use the `iodbctest` and `iodbctestw` utilities to establish a test connection with your driver. Use `iodbctest` to test how your driver works with an ANSI application, or use `iodbctestw` to test how your driver works with a Unicode application.

Note:

There are 32-bit and 64-bit installations of the iODBC driver manager available. If you have only one or the other installed, then the appropriate version of `iodbctest` (or `iodbctestw`) is available. However, if you have both 32- and 64-bit versions installed, then you need to make sure that you are running the version from the correct installation directory.

For more information about using the iODBC driver manager, see <http://www.iodbc.org>.


To test your connection using the iODBC driver manager:

1. Run **iodbctest** or **iodbctestw**.
2. Optionally, if you do not remember the DSN, then type a question mark (?) to see a list of available DSNs.
3. Type the connection string for connecting to your data store, and then press ENTER. For more information, see [Using a Connection String](#) on page 46.

If the connection is successful, then the `SQL>` prompt appears.

Using the unixODBC Driver Manager

You can use the `isql` and `iusql` utilities to establish a test connection with your driver and your DSN. `isql` and `iusql` can only be used to test connections that use a DSN. Use `isql` to test how your driver works with an ANSI application, or use `iusql` to test how your driver works with a Unicode application.

 **Note:**

There are 32-bit and 64-bit installations of the unixODBC driver manager available. If you have only one or the other installed, then the appropriate version of `isql` (or `iusql`) is available. However, if you have both 32- and 64-bit versions installed, then you need to make sure that you are running the version from the correct installation directory.

For more information about using the unixODBC driver manager, see <http://www.unixodbc.org>.


To test your connection using the unixODBC driver manager:

- Run `isql` or `iusql` by using the corresponding syntax:

- `isql [DataSourceName]`
- `iusql [DataSourceName]`

`[DataSourceName]` is the DSN that you are using for the connection.

If the connection is successful, then the `SQL>` prompt appears.

 **Note:**

For information about the available options, run `isql` or `iusql` without providing a DSN.

Using a Connection String

For some applications, you might need to use a connection string to connect to your data source. For detailed information about how to use a connection string in an ODBC application, refer to the documentation for the application that you are using.

The connection strings in the following sections are examples showing the minimum set of connection attributes that you must specify to successfully connect to the data source. Depending on the configuration of the data source and the type of connection you are working with, you might need to specify additional connection attributes. For detailed information about all the attributes that you can use in the connection string, see [Driver Configuration Options](#) on page 53.

DSN Connection String Example

The following is an example of a connection string for a connection that uses a DSN:

```
DSN=[DataSourceName]
```

[DataSourceName] is the DSN that you are using for the connection.

You can set additional configuration options by appending key-value pairs to the connection string. Configuration options that are passed in using a connection string take precedence over configuration options that are set in the DSN.

DSN-less Connection String Examples

Some applications provide support for connecting to a data source using a driver without a DSN. To connect to a data source without using a DSN, use a connection string instead.

! Important:

When you connect to the data store using a DSN-less connection string, the driver does not encrypt your credentials.

The placeholders in the examples are defined as follows, in alphabetical order:

- *[DatabaseName]* is the database that you want to access.
- *[DbGroups]* is the datasource security group or groups you want your user associated with (for more details see [Configuring IAM Authentication on a Non-Windows Machine](#) on page 38)

- *[PortNumber]* is the number of the TCP port that the Redshift server uses to listen for client connections.
- *[PPort]* is the number of the TCP port that the proxy server uses to listen for client connections.
- *[PServer]* is the IP address or host name of the proxy server to which you are connecting.
- *[Region]* is the region code for your Redshift server cluster
- *[SecretAccessKey]* is the secret key you use with your IAM access key and credentials.
- *[Server]* is the endpoint of the Redshift server to which you are connecting.
- *[YourCluster]* is the cluster ID for your Redshift server.
- *[YourAccessKey]* is the IAM access key you use with your IAM credentials.
- *[YourPassword]* is the password corresponding to your user name.
- *[YourUserName]* is the user name that you use to access the Redshift server.
- *[YourUserID]* is the user ID you use with your IAM credentials.

Connecting to a Redshift Server Directly

The following is the format of a DSN-less connection string for a basic connection to a Redshift server:

```
Driver=Simba Amazon Redshift ODBC Driver;  
Server=[Server];Port=[PortNumber];  
Database=[DatabaseName];UID=[YourUserName];  
PWD=[YourPassword];
```

For example:

```
Driver=Simba Amazon Redshift ODBC Driver;  
Server=testserver.abcabcabcabc.us-west-  
2.redshift.amazonaws.com;Port=5439;Database=TestDB;  
UID=simba;PWD=simba;
```

Connecting to a Redshift Server Through a Proxy Server

The following is the format of a DSN-less connection string for connecting to a Redshift server through a proxy server:

```
Driver=Simba Amazon Redshift ODBC Driver;  
Server=[Server];Port=[PortNumber];  
Database=[DatabaseName];UID=[YourUserName];  
PWD=[YourPassword];ProxyHost=[PServer];ProxyPort=[PPort];
```

For example:

```
Driver=Simba Amazon Redshift ODBC Driver;  
Server=testserver.abcabcabcabc.us-west-  
2.redshift.amazonaws.com;Port=5439;Database=TestDB;  
UID=simba;PWD=simba;ProxyHost=192.168.222.160;  
ProxyPort=8000;
```

Connecting to a Redshift Server using IAM User Credentials

The following is the format of a DSN-less connection string for connecting to a Redshift server using an IAM user profile:

```
Driver=Simba Amazon Redshift ODBC Driver;  
Server=[Server];Port=[PortNumber];  
Database=[DatabaseName];IAM=1;ClusterID=  
[YourCluster];Region=[YourRegion];DbUser=  
[YourUserID];AutoCreate=true;DbGroups=  
[DatabaseUserGroup];AccessKeyId=  
[YourAccessKey];SecretAccessKey=[YourSecretKey];
```

For example:

```
Driver=Simba Amazon Redshift ODBC  
Driver;Server=testserver.abcabcabcabc.us-west-  
2.redshift.amazonaws.com;Port=5439;Database=TestDB;IAM=1;Clu  
sterID=j-12345678;Region=na-central-  
1;DbUser=simba;AutoCreate=true;DbGroups=ViewOnly;AccessKeyId  
=AKIAIOSFODNN7EXAMPLE;SecretAccessKey=wJalrXUtnFEMI/K7MDENG/  
bPxrFiCYEXAMPLEKEY;
```


Features

For more information on the features of the Simba Amazon Redshift ODBC Driver, see the following:

- [Query Processing Modes](#) on page 49
- [TCP Keepalives](#) on page 50
- [Data Types](#) on page 50
- [Security and Authentication](#) on page 52

Query Processing Modes

To support performance tuning, the Simba Amazon Redshift ODBC Driver provides different query processing modes that you can configure to modify how the driver runs queries and retrieves results into memory. The following query processing modes are available:

- **Single Row Mode:** The driver returns query results one row at a time.
- **Declare/Fetch Mode:** The driver returns a user-specified number of rows at a time.
- **Retrieve Entire Result Mode:** The driver returns the entire query result into memory.
- **Multiple Statements Mode:** The driver can have multiple queries active on the same connection. The ODBC application may interleave calls to ODBC statements, but all queries are still sent and executed sequentially. When using this mode, the driver returns all the query results into memory.

By default, the driver returns the entire query result into memory.

Use a query processing mode that prevents queries from consuming too much memory, considering the expected result size of your queries and the specifications of your system.

If the Enforce Single Statement Mode option is set, the driver allows only one active statement at a time for each connection. This applies to Single Row, Declare/Fetch and Retrieve Entire Result modes. If you attempt to set both Enforce Single Statement and Multiple Statements modes, Multiple Statements Mode will take precedence.

For information about configuring how the driver processes queries, see [Configuring Additional Options on Windows](#) on page 19 if you are using the Windows version of the driver, or see [Configuring Query Processing Modes on a Non-Windows Machine](#) on page 40 if you are using a non-Windows version of the driver.

TCP Keepalives


By default, the Simba Amazon Redshift ODBC Driver is configured to use TCP keepalives to verify the status of a connection and prevent it from timing out. After you connect to a Redshift server, the driver automatically sends keepalive packets to the server. If the server does not respond, then the driver returns an indication that the connection is broken.

For information about configuring settings for TCP keepalives when using the Windows driver, see [Configuring TCP Keepalives on Windows](#) on page 20. For information about configuring settings for TCP keepalives when using the Linux or macOS driver, see [Configuring TCP Keepalives on a Non-Windows Machine](#) on page 41.

Data Types

The Simba Amazon Redshift ODBC Driver supports many common data formats, converting between Redshift data types and SQL data types.

The table below lists the supported data type mappings.

 **Note:**

If the Use Unicode option (the `UseUnicode` key) is enabled, then the driver returns `SQL_WCHAR` instead of `SQL_CHAR`, and `SQL_WVARCHAR` instead of `SQL_VARCHAR`.

Redshift Type	SQL Type
BigInt	SQL_BIGINT
Boolean	SQL_BIT If the Show Boolean Column As String option (the <code>BoolsAsChar</code> key) is enabled, then <code>SQL_VARCHAR</code> is returned instead.

Redshift Type	SQL Type
Char	SQL_CHAR <ul style="list-style-type: none"> • If the length of the column is greater than the Max Varchar (<code>MaxVarchar</code>) setting, then SQL_LONGVARCHAR is returned instead. • If the Use Unicode option (the <code>UseUnicode</code> key) is enabled, then SQL_WCHAR is returned instead. • If the Use Unicode option (the <code>UseUnicode</code> key) is enabled and the column length is greater than the Max Varchar (<code>MaxVarchar</code>) setting, then SQL_WLONGVARCHAR is returned instead.
Date	SQL_TYPE_DATE
Decimal	SQL_NUMERIC
Double Precision	SQL_DOUBLE
Integer	SQL_INTEGER
Real	SQL_REAL
SmallInt	SQL_SMALLINT
Text	SQL_VARCHAR <ul style="list-style-type: none"> • If the Use Unicode option (the <code>UseUnicode</code> key) is enabled, then SQL_WVARCHAR is returned instead. • If the Text As LongVarChar option (the <code>TextAsLongVarchar</code> key) is enabled, then SQL_LONGVARCHAR is returned instead. • If both options are enabled, then SQL_WLONGVARCHAR is returned instead.
Timestamp	SQL_TYPE_TIMESTAMP or SQL_TIMESTAMP (ODBC 2.0)

Redshift Type	SQL Type
VarChar	SQL_VARCHAR <ul style="list-style-type: none"> • If the length of the column is greater than the Max Varchar (<code>MaxVarchar</code>) setting, then SQL_LONGVARCHAR is returned instead. • If the Use Unicode option (the <code>UseUnicode</code> key) is enabled, then SQL_WVARCHAR is returned instead. • If the Use Unicode option (the <code>UseUnicode</code> key) is enabled and the column length is greater than the Max Varchar (<code>MaxVarchar</code>) setting, then SQL_WLONGVARCHAR is returned instead.

Security and Authentication

To protect data from unauthorized access, Redshift data stores require all connections to be authenticated using user credentials. Some data stores also require connections to be made over the Secure Sockets Layer (SSL) protocol, either with or without one-way authentication. The Simba Amazon Redshift ODBC Driver provides full support for these authentication protocols.

Note:

In this documentation, "SSL" refers to both TLS (Transport Layer Security) and SSL (Secure Sockets Layer). The driver supports TLS 1.1 and 1.2. The SSL version used for the connection is the highest version that is supported by both the driver and the server.

The driver provides a mechanism that enables you to authenticate your connection using your Redshift user name and password. For detailed configuration instructions, see [Creating a Data Source Name on Windows](#) on page 9 or [Creating a Data Source Name on a Non-Windows Machine](#) on page 32.

Additionally, the driver supports SSL connections with or without one-way authentication. If the server has an SSL-enabled socket, then you can configure the driver to connect to it.

It is recommended that you enable SSL whenever you connect to a server that is configured to support it. SSL encryption protects data and credentials when they are transferred over the network, and provides stronger security than authentication alone. For information about configuring SSL settings, see [Configuring SSL Verification on Windows](#) on page 11 or [Configuring SSL Verification on a Non-Windows Machine](#) on page 37.

Driver Configuration Options

Driver Configuration Options lists the configuration options available in the Simba Amazon Redshift ODBC Driver alphabetically by field or button label. Options having only key names, that is, not appearing in the user interface of the driver, are listed alphabetically by key name.

When creating or configuring a connection from a Windows machine, the fields and buttons described below are available in the following dialog boxes:

- Simba Amazon Redshift ODBC Driver DSN Setup
- Additional Options
- Data Type Configuration
- SSL Options
- Logging Options

When using a connection string or configuring a connection from a Linux or macOS machine, use the key names provided below.

Configuration Options Appearing in the User Interface

The following configuration options are accessible via the Windows user interface for the Simba Amazon Redshift ODBC Driver, or via the key name when using a connection string or configuring a connection from a Linux or macOS computer:

- [Allow Self-Signed Server Certificate](#) on page 54
- [Authentication Mode](#) on page 65
- [Cache Size](#) on page 55
- [Cluster Identifier](#) on page 55
- [Custom SSL Certificate Path](#) on page 56
- [Database](#) on page 56
- [DbUser](#) on page 56
- [DbGroups](#) on page 57
- [Enable Table Types](#) on page 57
- [Encrypt Password](#) on page 57
- [Enforce Single Statement Mode](#)
- [Port](#) on page 62
- [Profile Name](#) on page 62
- [Proxy Port](#) on page 62
- [Proxy Server](#) on page 63
- [Retrieve Entire Result Into Memory](#) on page 63
- [Server](#) on page 64
- [SessionToken](#) on page 66
- [SecretAccessKey](#) on page 63
- [Show Boolean Column As String](#) on page 64
- [Single Row Mode](#) on page 64
- [SSL Insecure](#) on page 66

- on page 58
- [IdP Host](#) on page 58
- [IdP Port](#) on page 58
- [Log Level](#) on page 59
- [Log Path](#) on page 60
- [Max LongVarChar](#) on page 60
- [Max Varchar](#) on page 60
- [Okta App ID](#) on page 61
- [Password](#) on page 61
- [Preferred Role](#) on page 61
- [Text As LongVarChar](#) on page 66
- [Use Declare/Fetch](#) on page 67
- [Use Multiple Statements](#) on page 67
- [Use System Trust Store](#) on page 68
- [Use Unicode](#) on page 68
- [User](#) on page 69
- [User Auto Create](#) on page 69

AccessKeyID

Key Name	Default Value	Required
AccessKeyID	None	No.

Description

The IAM access key for the user or role. If this is specified IAMSecretAccessKey must also be specified.

Allow Self-Signed Server Certificate

Key Name	Default Value	Required
AllowSelfSigned ServerCert	Clear (0)	No

Description

This option specifies whether the driver allows self-signed certificates from the server.

- Enabled (1): The driver authenticates the Redshift server even if the server is using a self-signed certificate.
- Disabled (0): The driver does not allow self-signed certificates from the server.



Note:

This setting is applicable only when SSL is enabled.

Cache Size

Key Name	Default Value	Required
Fetch	100	Yes, if Declare/Fetch Mode is enabled.

Description

The number of rows that the driver returns when Declare/Fetch Mode is enabled. For more information, see [Use Declare/Fetch](#) on page 67.

Check Certificate Revocation

Key Name	Default Value	Required
CheckCertRevocation	Clear (0)	No

Description

This option specifies whether the driver checks to see if a certificate has been revoked while retrieving a certificate chain from the Windows Trust Store.

This option is only applicable if you are using a CA certificate from the Windows Trust Store (see [Use System Trust Store](#) on page 68).

- Enabled (1): The driver checks for certificate revocation while retrieving a certificate chain from the Windows Trust Store.
- Disabled (0): The driver does not check for certificate revocation while retrieving a certificate chain from the Windows Trust Store.



Note:

This option is only available on Windows.

Cluster Identifier

Key Name	Default Value	Required
ClusterID	None	No

Description

The name of the Redshift cluster you want to connect to.

Custom SSL Certificate Path

Key Name	Default Value	Required
SSLCertPath	The location of the driver DLL file.	No

Description

The full path of the file containing the root certificate for verifying the server.

If this option is not set, then the driver looks in the folder that contains the driver DLL file.

Database

Key Name	Default Value	Required
Database	None	Yes

Description

The name of the Redshift database that you want to access.

DbUser

Key Name	Default Value	Required
DbUser	None	No

Description

The user ID you wish to use with your Redshift account. You can use an ID that does not currently exist if you have enabled **User Auto Create**. See [Configuring Authentication on Windows](#) on page 11 or [Creating a Data Source Name on a Non-Windows Machine](#) on page 32.

DbGroups

Key Name	Default Value	Required
DbGroups	None	No

Description

A comma-separated list of existing database group names that a DbUser will join for the current session. If not specified, defaults to PUBLIC.

Enable Table Types

Key Name	Default Value	Required
EnableTableTypes	Clear (0)	No

Description

This option specifies whether the driver recognizes table type information from the data source. By default, the driver only recognizes a single, generic table type.

- **Clear (0):** All tables returned from the data source have the generic type TABLE.
- **Selected (1):** The driver recognizes the following table types: TABLE, SYSTEM TABLE, and GLOBAL TEMPORARY.

Encrypt Password

Key Name	Default Value	Required
N/A	All Users Of This Machine	No

Description

This option specifies how the driver encrypts the credentials that are saved in the DSN:

- **Current User Only:** The credentials are encrypted, and can only be used by the current Windows user.
- **All Users Of This Machine:** The credentials are encrypted, but can be used by any user on the current Windows machine.

! Important:

This option is available only when you configure a DSN using the Simba Amazon Redshift ODBC Driver DSN Setup dialog box in the Windows driver. When you connect to the data store using a connection string, the driver does not encrypt your credentials.

Enforce Single Statement Mode

Key Name	Default Value	Required
<code>enforceSingleStatement</code>	0	No

Description

This option specifies if the driver can have more than one active query at a time per connection.

- Enabled (1): The driver can have only one active query at a time.
- Disabled (0): The driver can have multiple active queries if the Use Multiple Statements option (the `UseMultipleStatements` property) is enabled. For more information, see [Use Multiple Statements](#) on page 67.

IdP Host

Key Name	Default Value	Required
<code>idp_host</code>	None	No

Description

The IDP host you are using to authenticate into Redshift.

IdP Port

Key Name	Default Value	Required
<code>idp_port</code>	None	No

Description

The port for an IDP. This can be omitted if it is part of a specified user profile.

Log Level

Key Name	Default Value	Required
LogLevel	OFF (0)	No

Description

Use this property to enable or disable logging in the driver and to specify the amount of detail included in log files.

! Important:

- Only enable logging long enough to capture an issue. Logging decreases performance and can consume a large quantity of disk space.
- This option is not supported in connection strings. To configure logging for the Windows driver, you must use the Logging Options dialog box. To configure logging for a non-Windows driver, you must use the `simba.amazonredshiftodbc.ini` file.

Set the property to one of the following values:

- OFF (0): Disable all logging.
- FATAL (1): Logs severe error events that lead the driver to abort.
- ERROR (2): Logs error events that might allow the driver to continue running.
- WARNING (3): Logs events that might result in an error if action is not taken.
- INFO (4): Logs general information that describes the progress of the driver.
- DEBUG (5): Logs detailed information that is useful for debugging the driver.
- TRACE (6): Logs all driver activity.

When logging is enabled, the driver produces two log files at the location you specify in the Log Path (`LogPath`) property, where `[DriverName]` is the name of the driver:

- A `[DriverName]_driver.log` file that logs driver activity that is not specific to a connection.
- A `[DriverName]_connection_[Number].log` for each connection made to the database, where `[Number]` is a number that identifies each log file. This file logs driver activity that is specific to the connection.

If you enable the `UseLogPrefix` connection property, the driver prefixes the log file name with the user name associated with the connection and the process ID of the application through which the connection is made. For more information, see [UseLogPrefix](#) on page 73.

Log Path

Key Name	Default Value	Required
LogPath	None	Yes, if logging is enabled.

Description

The full path to the folder where the driver saves log files when logging is enabled.

! Important:

This option is not supported in connection strings. To configure logging for the Windows driver, you must use the Logging Options dialog box. To configure logging for a non-Windows driver, you must use the `simba.amazonredshiftodbc.ini` file.

Max LongVarChar

Key Name	Default Value	Required
MaxLongVarChar	8190	No

Description

The maximum data length for LongVarChar columns.

Max Varchar

Key Name	Default Value	Required
MaxVarchar	255	No

Description

The maximum data length for VarChar columns.

Okta App ID

Key Name	Default Value	Required
app_id	None	No.

Description

The Okta-provided unique ID associated with your Redshift application.

Password

Key Name	Default Value	Required
PWD OR Password	None	Yes if Username has been set.

Description

The password corresponding to the user name that you provided in the User field (the Username or UID key).

Preferred Role

Key Name	Default Value	Required
preferred_role	None	No

Description

The role you want to assume during the connection to Redshift.

Profile Name

Key Name	Default Value	Required
profile	None	No.

Description

The name of the user Profile you are using to authenticate into Redshift.



Note:

The default location for the credentials file that contains profiles is `~/.aws/Credentials`. The `AWS_SHARED_CREDENTIALS_FILE` environment variable can be used to point to a different credentials file.

Port

Key Name	Default Value	Required
Port	5439	Yes

Description

The TCP port that the Redshift server uses to listen for client connections.

Proxy Port

Key Name	Default Value	Required
ProxyPort	None	Yes, if connecting through a proxy server.

Description

The number of the port that the proxy server uses to listen for client connections.

Proxy Server

Key Name	Default Value	Required
ProxyHost	None	Yes, if connecting through a proxy server.

Description

The host name or IP address of a proxy server that you want to connect through.

Retrieve Entire Result Into Memory

Key Name	Default Value	Required
N/A	Selected	No

Description

When this option is enabled, the driver returns the entire query result into memory instead of returning results in chunks or single rows at a time.

When using keys to set driver options, you can enable this option by setting the `SingleRowMode`, `UseDeclareFetch`, and `UseMultipleStatements` keys to 0.

Note:

By default, the driver is configured to use this query processing mode.

SecretAccessKey

Key Name	Default Value	Required
SecretAccessKey	None	No.

Description

The IAM secret key for the user or role. If this is specified, `IAMAccessKeyID` must also be specified.

SessionToken

Key Name	Default Value	Required
SessionToken	None	No.

Description

The temporary IAM session token associated with the IAM role you are using to authenticate.

Server

Key Name	Default Value	Required
Server	None	Yes

Description

The endpoint of the Redshift server.

Show Boolean Column As String

Key Name	Default Value	Required
BoolsAsChar	Clear (0)	No

Description

This option specifies the SQL data type that the driver uses to return Boolean data.

- Enabled (1): The driver returns Boolean columns as SQL_VARCHAR data with a length of 5.
- Disabled (0): The driver returns Boolean columns as SQL_BIT data.

Single Row Mode

Key Name	Default Value	Required
SingleRowMode	Cleared (0)	No

Description

When this option is enabled (1), the driver uses Single Row Mode and returns query results one row at a time. Enable this option if you plan to query large results and do not want to retrieve the entire result into memory.

When using keys to set driver options, make note of the following:

- If `SingleRowMode` and `UseDeclareFetch` are both set to 0, then the driver retrieves the entire query result into memory.
- If `UseDeclareFetch` is set to 1, then it takes precedence over `SingleRowMode`.
- If `SingleRowMode` is set to 1 and `UseDeclareFetch` is set to 0, then `SingleRowMode` takes precedence over `UseMultipleStatements`.

Authentication Mode

Key Name	Default Value	Required
<code>SSLMode</code>	<code>requireprefer</code>	No

Description

The SSL certificate verification mode to use when connecting to Redshift. The following values are possible:

- **verify-full**: Connect only using SSL, a trusted certificate authority, and a server name that matches the certificate.
- **verify-ca**: Connect only using SSL and a trusted certificate authority.
- **require**: Connect only using SSL.
- **prefer**: Connect using SSL if available. Otherwise, connect without using SSL.
- **allow**: By default, connect without using SSL. If the server requires SSL connections, then use SSL.
- **disable**: Connect without using SSL.

Note:

For information about SSL support in Amazon Redshift, see "Connect Using SSL" in the *Amazon Redshift Management Guide*:

<http://docs.aws.amazon.com/redshift/latest/mgmt/connecting-ssl-support.html#connect-using-ssl>.

SSL Insecure

Key Name	Default Value	Required
ssl_insecure	false	No

Description

- `true`: The driver will not check the authenticity of the IDP server certificate.
- `false`: The driver will check the authenticity of the IDP server certificate.

Text As LongVarChar

Key Name	Default Value	Required
TextAsLongVarchar	0	No

Description

This option specifies the SQL data type that the driver uses to return Text data. The returned data type is also affected by the Use Unicode option (the `UseUnicode` key). For more information, see [Use Unicode](#) on page 68.

- Enabled (1): The driver returns Text columns as `SQL_LONGVARCHAR` data. If the Use Unicode option (the `UseUnicode` key) is also enabled, then the driver returns `SQL_WLONGVARCHAR` data instead.
- Disabled (0): The driver returns Text columns as `SQL_VARCHAR` data. If the Use Unicode option (the `UseUnicode` key) is also enabled, then the driver returns `SQL_WVARCHAR` data instead.

SessionToken

Key Name	Default Value	Required
SessionToken	None	No.

Description

The temporary IAM session token associated with the IAM role you are using to authenticate.

Use Declare/Fetch

Key Name	Default Value	Required
UseDeclareFetch	Cleared (0)	No

Description

When this option is enabled (1), the driver uses Declare/Fetch Mode and returns a specific number of rows at a time. To set the number of rows, configure the Cache Size option (the `Fetch` key).

When using keys to set driver options, make note of the following:

- If `UseDeclareFetch` is set to 1, then it takes precedence over `SingleRowMode` and `UseMultipleStatements`.
- If `UseDeclareFetch` is set to 0 and `SingleRowMode` is set to 1, then the driver returns query results one row at a time.
- If `UseDeclareFetch` and `SingleRowMode` are both set to 0, then the driver retrieves the entire query result into memory.

Use Multiple Statements

Key Name	Default Value	Required
UseMultipleStatements	0	No

Description

When this option is enabled (1), the driver can have multiple queries active on the same connection. The ODBC application may interleave calls to ODBC statements, but all queries are still sent and executed sequentially. The driver returns all the query results into memory.

When this option is disabled (0), the driver runs queries one at a time.

When using keys to set driver options, make note of the following:

- If `UseDeclareFetch` is set to 1, then it takes precedence over `UseMultipleStatements`.
- If `UseDeclareFetch` is set to 0 and `SingleRowMode` is set to 1, then `SingleRowMode` takes precedence over `UseMultipleStatements`.

Use System Trust Store

Key Name	Default Value	Required
UseSystemTrustStore	Selected (1)	No

Description

This option specifies whether to use a CA certificate from the system trust store, or from a specified PEM file.

- Enabled (1): The driver verifies the connection using a certificate in the system trust store.
- Disabled (0): The driver verifies the connection using a specified PEM file.

Note:

This option is only available on Windows.

Use Unicode

Key Name	Default Value	Required
UseUnicode	1	No

Description

This option specifies whether the driver returns Redshift data as Unicode or regular SQL types.

When this option is enabled (1), the driver returns data as Unicode character types:

- SQL_WCHAR is returned instead of SQL_CHAR.
- SQL_WVARCHAR is returned instead of SQL_VARCHAR.
- SQL_WLONGVARCHAR is returned instead of SQL_LONGVARCHAR.

When this option is disabled (0), the driver returns data as regular SQL types:

- SQL_CHAR is returned instead of SQL_WCHAR.
- SQL_VARCHAR is returned instead of SQL_WVARCHAR.
- SQL_LONGVARCHAR is returned instead of SQL_WLONGVARCHAR.

For detailed information about how the driver returns Redshift data as SQL types, see [Data Types](#) on page 50.

User

Key Name	Default Value	Required
UID OR User	None	No

Description

The user name that you use to access the Redshift server.

If you are using keys to set driver options, `UID` takes precedence over `Username`.

If you are using IAM authentication, can be used in the following ways:

- If the connection uses a credential provider plugin, this will be the user name for the `idp_host` server. In this case the information can be included in a user profile and may not be required for the connection URL.
- If your connection does not use a credential provider, this is used as the user name for your data source or `UID`.

If this value is defined in multiple places, the preference order will be: `DbUser` > `user` > `UID`.

User Auto Create

Key Name	Default Value	Required
<code>AutoCreate</code>	<code>false</code>	No.

Description

- `true`: If user specified by either `DbUser` or `UID` does not exist, a new user with that name should be created.
- `false`: New users will not be created by the driver. If the specified user does not exist the authentication fails.

Configuration Options Having Only Key Names

The following configuration options do not appear in the Windows user interface for the Simba Amazon Redshift ODBC Driver. They are accessible only when you use a connection string or configure a connection on macOS or Linux.

- [cafile](#) on page 71
- [Driver](#) on page 70
- [KeepAlive](#) on page 71
- [KeepAliveCount](#) on page 71
- [KeepAliveInterval](#) on page 72
- [KeepAliveTime](#) on page 72
- [Locale](#) on page 72
- [plugin_name](#) on page 73
- [UseLogPrefix](#) on page 73

Driver

Key Name	Default Value	Required
Driver	Simba Amazon Redshift ODBC Driver when installed on Windows, or the absolute path of the driver shared object file when installed on a non-Windows machine.	Yes

Description

On Windows, the name of the installed driver (Simba Amazon Redshift ODBC Driver).

On other platforms, the name of the installed driver as specified in `odbcinst.ini`, or the absolute path of the driver shared object file.

cafile

Key Name	Default Value	Required
cafile	None	No

Description

The file path to the CA certificate file.

iam

Key Name	Default Value	Required
iam	0	No

Description

Indicated if IAM authentication should be used.

- 0: A standard authentication method will be used.
- 1: An IAM authentication method will be used.

KeepAlive

Key Name	Default Value	Required
KeepAlive	1	No

Description

When this option is enabled (1), the driver uses TCP keepalives to prevent connections from timing out.

When this option is disabled (0), the driver does not use TCP keepalives.

KeepAliveCount

Key Name	Default Value	Required
KeepAliveCount	0	No

Description

The number of TCP keepalive packets that can be lost before the connection is considered broken.

When this key is set to 0, the driver uses the system default for this setting.

KeepAliveTime

Key Name	Default Value	Required
KeepAliveTime	0	No

Description

The number of seconds of inactivity before the driver sends a TCP keepalive packet.

When this key is set to 0, the driver uses the system default for this setting.

KeepAliveInterval

Key Name	Default Value	Required
KeepAliveInterval	0	No

Description

The number of seconds between each TCP keepalive retransmission.

When this key is set to 0, the driver uses the system default for this setting.

Locale

Key Name	Default Value	Required
Locale	en-US	No

Description

The locale to use for error messages.

plugin_name

Key Name	Default Value	Required
plugin_name	None	No

Descriptions

This is a string that identifies the specific credentials provider plugin class. Three services are supported:

- `adfs` (Active Directory Federation Service)
- `ping` (PingFederate Service)
- `okta` (Okta Service)

UseLogPrefix

Key Name	Default Value	Required
UseLogPrefix	0	No

Description

This option specifies whether the driver includes a prefix in the names of log files so that the files can be distinguished by user and application.

- 1: The driver prefixes log file names with the user name and process ID associated with the connection that is being logged.

For example, if you are connecting as a user named "jdoe" and using the driver in an application with process ID 7836, the generated log file would be named `jdoe_7836_SimbaRedshiftODBCDriver.log` and `jdoe_7836_SimbaRedshiftODBCDriver_connection_[Number].log`, where **[Number]** is a number that identifies each connection-specific log file.

- 0: The driver does not include the prefix in log file names.

Third-Party Trademarks

Linux is the registered trademark of Linus Torvalds in Canada, United States and/or other countries.

Mac, macOS, Mac OS, and OS X are trademarks or registered trademarks of Apple, Inc. or its subsidiaries in Canada, United States and/or other countries.

Microsoft, MSDN, Windows, Windows Server, Windows Vista, and the Windows start button are trademarks or registered trademarks of Microsoft Corporation or its subsidiaries in Canada, United States and/or other countries.

Red Hat, Red Hat Enterprise Linux, and CentOS are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in Canada, United States and/or other countries.

SUSE is a trademark or registered trademark of SUSE LLC or its subsidiaries in Canada, United States and/or other countries.

Amazon Redshift, Amazon, and Redshift are trademarks or registered trademarks of Amazon Web Services, Inc. or its subsidiaries in Canada, United States and/or other countries.

All other trademarks are trademarks of their respective owners.

Third-Party Licenses

The licenses for the third-party libraries that are included in this product are listed below.

CityHash License

Copyright (c) 2011 Google, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

CityHash, by Geoff Pike and Jyrki Alakuijala

<http://code.google.com/p/cityhash/>

cURL License

COPYRIGHT AND PERMISSION NOTICE

Copyright (c) 1996 - 2015, Daniel Stenberg, daniel@haxx.se.

All rights reserved.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE

AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

dtoa License

The author of this software is David M. Gay.

Copyright (c) 1991, 2000, 2001 by Lucent Technologies.

Permission to use, copy, modify, and distribute this software for any purpose without fee is hereby granted, provided that this entire notice is included in all copies of any software which is or includes a copy or modification of this software and in all copies of the supporting documentation for such software.

THIS SOFTWARE IS BEING PROVIDED "AS IS", WITHOUT ANY EXPRESS OR IMPLIED WARRANTY. IN PARTICULAR, NEITHER THE AUTHOR NOR LUCENT MAKES ANY REPRESENTATION OR WARRANTY OF ANY KIND CONCERNING THE MERCHANTABILITY OF THIS SOFTWARE OR ITS FITNESS FOR ANY PARTICULAR PURPOSE.

Expat License

Copyright (c) 1998, 1999, 2000 Thai Open Source Software Center Ltd

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NOINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

ICU License - ICU 1.8.1 and later**COPYRIGHT AND PERMISSION NOTICE**

Copyright (c) 1995-2014 International Business Machines Corporation and others

All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

All trademarks and registered trademarks mentioned herein are the property of their respective owners.

OpenSSL License

Copyright (c) 1998-2016 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"

4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License

Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)

All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The

SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

"This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)"

The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related :-).

4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement:

"This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The licence and distribution terms for any publically available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution licence [including the GNU Public Licence.]

PostgreSQL Database Management System License

(formerly known as Postgres, then as Postgres95)

Portions Copyright (c) 1996-2015, The PostgreSQL Global Development Group

Portions Copyright (c) 1994, The Regents of the University of California

Permission to use, copy, modify, and distribute this software and its documentation for any purpose, without fee, and without a written agreement is hereby granted, provided that the above copyright notice and this paragraph and the following two paragraphs appear in all copies.

IN NO EVENT SHALL THE UNIVERSITY OF CALIFORNIA BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, INCLUDING LOST PROFITS, ARISING OUT OF THE USE OF THIS SOFTWARE AND ITS DOCUMENTATION, EVEN IF THE UNIVERSITY OF CALIFORNIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE UNIVERSITY OF CALIFORNIA SPECIFICALLY DISCLAIMS ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE SOFTWARE PROVIDED HEREUNDER IS ON AN "AS IS" BASIS, AND THE UNIVERSITY OF CALIFORNIA HAS NO OBLIGATIONS TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

Stringencoders License

Copyright 2005, 2006, 2007

Nick Galbreath -- nickg [at] modp [dot] com

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of the modp.com nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This is the standard "new" BSD license:

<http://www.opensource.org/licenses/bsd-license.php>

UPX License

-----BEGIN PGP SIGNED MESSAGE-----

```

○○○○○      ○○○ ○○○○○○○○○.   ○○○○○○○○ ○○○○○
`888'      `8' `888 `Y88. `8888 d8'
 888      8 888 .d88' Y888..8P
 888      8 888○○○88P' `8888'
 888      8 888 .8PY888.
`88.     .8' 888 d8' `888b
`YbodP'   ○888○      ○888○ ○88888○

```

The Ultimate Packer for eXecutables

Copyright (c) 1996-2000 Markus Oberhumer & Laszlo Molnar

<http://wildsau.idv.uni-linz.ac.at/mfx/upx.html>

<http://www.nexus.hu/upx>

<http://upx.tsx.org>

PLEASE CAREFULLY READ THIS LICENSE AGREEMENT, ESPECIALLY IF YOU PLAN TO MODIFY THE UPX SOURCE CODE OR USE A MODIFIED UPX VERSION.

ABSTRACT

=====

UPX and UCL are copyrighted software distributed under the terms of the GNU General Public License (hereinafter the "GPL").

The stub which is imbedded in each UPX compressed program is part of UPX and UCL, and contains code that is under our copyright. The terms of the GNU General Public License still apply as compressing a program is a special form of linking with our stub.

As a special exception we grant the free usage of UPX for all executables, including commercial programs. See below for details and restrictions.

COPYRIGHT

=====

UPX and UCL are copyrighted software. All rights remain with the authors.

UPX is Copyright (C) 1996-2000 Markus Franz Xaver Johannes Oberhumer

UPX is Copyright (C) 1996-2000 Laszlo Molnar

UCL is Copyright (C) 1996-2000 Markus Franz Xaver Johannes Oberhumer

GNU GENERAL PUBLIC LICENSE

=====

UPX and the UCL library are free software; you can redistribute them and/or modify them under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

UPX and UCL are distributed in the hope that they will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; see the file COPYING.

SPECIAL EXCEPTION FOR COMPRESSED EXECUTABLES

=====

The stub which is imbedded in each UPX compressed program is part of UPX and UCL, and contains code that is under our copyright. The terms of the GNU General Public License still apply as compressing a program is a special form of linking with our stub.

Hereby Markus F.X.J. Oberhumer and Laszlo Molnar grant you special permission to freely use and distribute all UPX compressed programs (including commercial ones), subject to the following restrictions:

1. You must compress your program with a completely unmodified UPX version; either with our precompiled version, or (at your option) with a self compiled version of the unmodified UPX sources as distributed by us.
2. This also implies that the UPX stub must be completely unmodified, i.e. the stub imbedded in your compressed program must be byte-identical to the stub that is produced by the official unmodified UPX version.
3. The decompressor and any other code from the stub must exclusively get used by the unmodified UPX stub for decompressing your program at program startup. No portion of the stub may get read, copied, called or otherwise get used or accessed by your program.

ANNOTATIONS

=====

- You can use a modified UPX version or modified UPX stub only for programs that are compatible with the GNU General Public License.

- We grant you special permission to freely use and distribute all UPX compressed programs. But any modification of the UPX stub (such as, but not limited to, removing our copyright string or making your program non-decompressible) will immediately revoke your right to use and distribute a UPX compressed program.

- UPX is not a software protection tool; by requiring that you use the unmodified UPX version for your proprietary programs we make sure that any user can decompress your program. This protects both you and your users as nobody can hide malicious code - any program that cannot be decompressed is highly suspicious by definition.

- You can integrate all or part of UPX and UCL into projects that are compatible with the GNU GPL, but obviously you cannot grant any special exceptions beyond the GPL for our code in your project.

- We want to actively support manufacturers of virus scanners and similar security software. Please contact us if you would like to incorporate parts of UPX or UCL into such a product.

Markus F.X.J. Oberhumer Laszlo Molnar

markus.oberhumer@jk.uni-linz.ac.at

ml1050@cdata.tvnet.hu

Linz, Austria, 25 Feb 2000

-----BEGIN PGP SIGNATURE-----

Version: 2.6.3ia

Charset: noconv

iQCVAwUBOLaLS210fyLu8beJAQFYVAP/ShzENWKLTVedLCjZbDcwaBEHfUVcrG
MI
wE7frMkbWT2zmkdv9hW90WmjMhOBu7yhUplvN8BKotLioIEnZmLCYu8AGCwr5wB
f
dfLoClxnzFTgQv5axF1awp4RwCUH3hf4cDrOVqmAsWXKPHtm4hx96jF6L4oHhJx
OO03+ojZdO8=
=CS52

-----END PGP SIGNATURE-----

zlib License

Copyright (C) 1995-2013 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly

Mark Adler

jloup@gzip.org

madler@alumni.caltech.edu

Apache License, Version 2.0

The following notice is included in compliance with the Apache License, Version 2.0 and is applicable to all software licensed under the Apache License, Version 2.0.

Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing

lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. **Grant of Copyright License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. **Grant of Patent License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
4. **Redistribution.** You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
 - (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and
 - (b) You must cause any modified files to carry prominent notices stating that You changed the files; and
 - (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
 - (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least

one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. **Submission of Contributions.** Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. **Trademarks.** This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. **Disclaimer of Warranty.** Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
8. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. **Accepting Warranty or Additional Liability.** While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

This product includes software that is licensed under the Apache License, Version 2.0 (listed below):

AWS SDK for C++

Copyright © 2015 Amazon.com, Inc. or its affiliates. All Rights Reserved.

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.