



# Magnitude Simba SDK

---

**Build a C++ ODBC Driver in 5 Days (Linux/Unix)**

Version 10.2.2

October 2022

## Copyright

This document was released in October 2022.

Copyright ©2014-2022 Magnitude Software, Inc., an insightsoftware company. All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without prior written permission from Magnitude, Inc.

The information in this document is subject to change without notice. Magnitude, Inc. strives to keep this information accurate but does not warrant that this document is error-free.

Any Magnitude product described herein is licensed exclusively subject to the conditions set forth in your Magnitude license agreement.

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

All other company and product names mentioned herein are used for identification purposes only and may be trademarks or registered trademarks of their respective owners.

Information about the third-party products is contained in a third-party-licenses.txt file that is packaged with the software.

## Contact Us

Magnitude Software, Inc.

[www.magnitude.com](http://www.magnitude.com)

## Table of Contents

---

About this Guide .....	5
Simba SDK Overview .....	8
ODBC Standards .....	8
The Simba SDK Solution .....	8
About the QuickStart Sample Connector .....	9
Day One .....	13
Install the Simba SDK .....	13
Build the Sample ODBC Connector .....	14
Configure the Connector and Data Source .....	15
Connect to the Data Source .....	21
Set up a Custom ODBC Connector Project .....	23
Configure Your Custom Connector and Data Source .....	24
Debug Your Custom Connector .....	26
Enable Logging .....	27
Day Two .....	29
Set the Configuration Branding .....	29
Set Connector Properties .....	29
Set Logging Details .....	30
Check Connection Settings .....	31
Establish a Connection .....	33
Day Three .....	34
Create and Return Metadata Sources .....	34
Day Four .....	40
Enable Data Retrieval .....	40
Day Five .....	43
Rebrand Error Messages .....	43
Rebrand the Custom ODBC Connector .....	44
Reference .....	45
Driver Managers .....	45
Locating the Configuration Files .....	46

---

Data Retrieval .....	48
Server Configuration .....	50
How to Add Schema Support .....	51
Install the Evaluation License .....	51
Troubleshooting .....	52
Contact Us .....	53
Third-Party Trademarks .....	54

## About this Guide

### Purpose

This guide explains how to use the Magnitude Simba SDK to create a custom ODBC connector for a data store that is not SQL-capable. It explains how to customize the QuickStart sample connector, which is included with the Simba SDK.

Using this sample connector is the quickest and easiest way to create a custom ODBC connector. At the end of five days, you will have a read-only connector that connects to your data store. This custom ODBC connector can be used as the foundation for a commercial DSI implementation.

#### Note:

An online version of this guide is located at <http://www.simba.com/resources/sdk/documentation>.

### Advantages of Using the Simba SDK

The ODBC specification defines a rich interface that allows any ODBC-enabled application to connect to a data store. In order to implement a connector that supports this specification, developers have to understand all the complexities of error checking, session management, and data conversion, then design their code in a robust and efficient manner. Developers must also understand how to optimize data retrieval in order to get maximum performance when connecting to large and complex data stores.

For data stores that do not support SQL, the Simba SDK provides an SQL parser and an execution engine. Developers can use these features to translate SQL queries to a custom API that the data store understands.

The Simba SDK, developed by experts in the field, is a complete implementation of the ODBC specification. It exposes an easy-to-use SDK that allows you to create a robust and efficient connector for your data store.

### Build a Custom ODBC Connector in Five Days

Over the course of five days, this guide explains how to accomplish the following tasks:

1. Set up the development environment and build the sample connector.
2. Use the sample connector as a template to create a custom ODBC connector.

3. Make a connection to the data store.
4. Retrieve metadata.
5. Work with columns.
6. Retrieve data.
7. Rename and rebrand the custom ODBC connector.

In the QuickStart connector, the areas of code that require modification are marked with “TODO” messages and a short explanation. Some of these changes customize the connector for your specific data store, while other changes rename the connector for your company or product.

## Audience

The guide is intended for developers who want to use the Simba SDK to build a connector for a data store that is not SQL-capable.

## 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:

Indicates a short note appended to a paragraph.

### IMPORTANT:

Indicates an important comment related to the preceding paragraph.

## Knowledge Prerequisites

To use the Simba SDK to build a custom ODBC connector, the following knowledge is helpful:

- Familiarity with the C++ programming language.
- Ability to use the data store to which the connector you are developing will connect.
- An understanding of the role of ODBC technologies and driver managers in

connecting to a data store.

- Exposure to SQL.

## Variables Used in this Document

The following variables are used in this document:

Variable	Description
<code>[INSTALL_DIR]</code>	<p>Installation directory for the SimbaEngine X SDK.</p> <p><b>Default value on Windows platforms:</b> <code>C:\SimbaTechnologies\SimbaEngineSDK\10.2</code></p> <p><b>Default value on Linux, Unix, and macOS platforms:</b> <code>[UNTAR_DIR]/SimbaEngineSDK/10.2</code></p>
<code>[UNTAR_DIR]</code>	<p>Directory where the SimbaEngine X SDK distributable was untarred.</p>
<code>[JDBC_VERSION]</code>	<p>The version of JDBC that your driver supports.</p> <p>You can use the SimbaEngine X SDK to build a driver for version 4.2 and 4.3, or a hybrid version.</p> <p>Possible values of <code>[JDBC_VERSION]</code> are 42 and 43, and Hybrid.</p>

## Simba SDK Overview

Applications, such as Crystal Reports and Tableau, use connectors to connect to data stores from which they read and write data. Applications support the ODBC protocol to enable connection with any connector that also supports ODBC. A connector exposes the ODBC protocol to the application and another API, such as SQL or a custom API, to the data store.

### **Note:**

This guide explains how to create an ODBC connector for a data store that is not SQL-capable. To create an ODBC connector for a data store that is SQL-capable, see [Build an ODBC Connector for SQL-Capable Data Sources in 5 Days](#).

## ODBC Standards

ODBC is one of the most established and widely-supported APIs for connecting to and working with databases. A main component of this technology is the ODBC connector, which connects an application to the database.

For a brief description of the ODBC standard, see <http://www.simba.com/resources/data-access-standards-library#!odbc>.


For complete information on the ODBC 3.80 specification, see the ODBC Programmer's Reference at [http://msdn.microsoft.com/en-us/library/ms714177\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/ms714177(v=vs.85).aspx).

## The Simba SDK Solution

Connectors based on the Magnitude Simba SDK leverage its error checking, session management, data conversion, optimization, and other low-level implementation details. The Simba SDK uses ODBC to communicate with the driver manager and a simple API (called the Data Store Interface API or DSI API) to communicate with the data store. The DSI API defines the primitive operations needed to access a data store.

The figure below shows a typical ODBC stack:





SDK developers create an implementation of a DSI (also known as a DSI Implementation or DSII) that applications use to access the particular data store in the process of executing an SQL statement. In the final executable, the components from Simba SDK take responsibility for meeting the data access standards while the custom DSI implementation takes responsibility for accessing the data store and translating it to the DSI API.

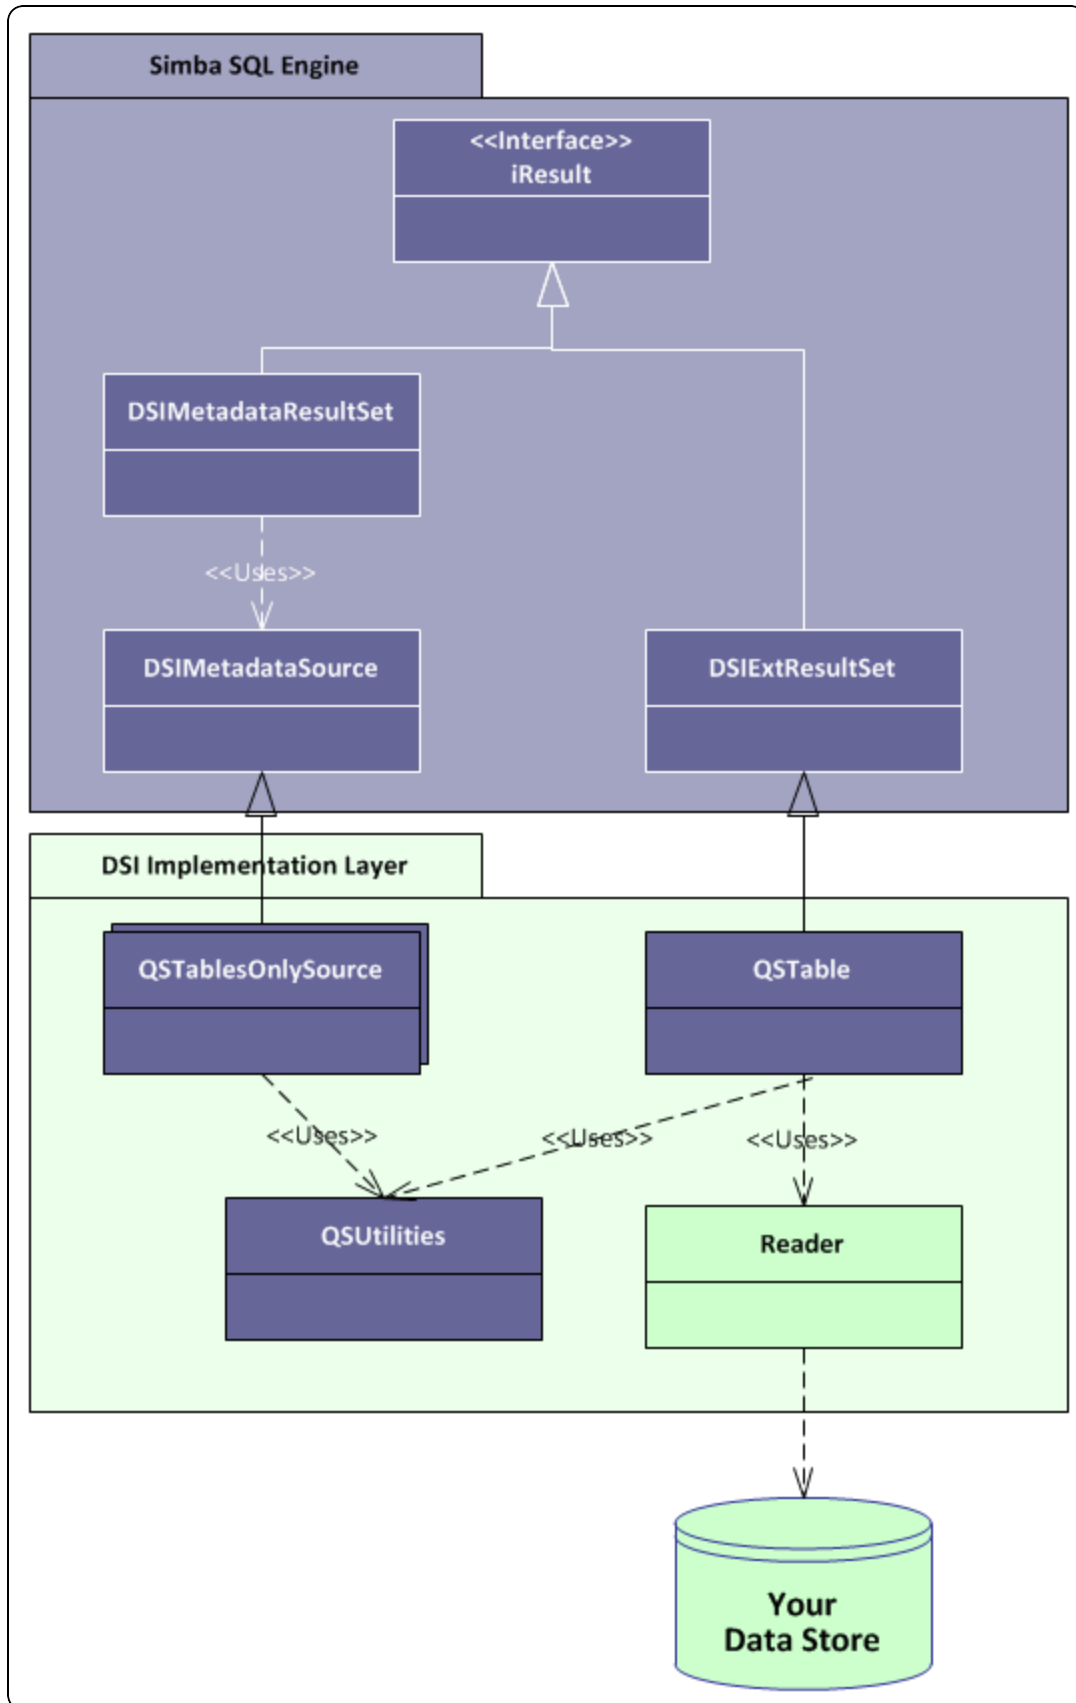
ODBC applications, such as Tableau or Microsoft Excel, use this executable when connecting to the data store in the process of executing an SQL statement.

### About the QuickStart Sample Connector

The Simba SDK includes a sample connector that you can use as a template to create a custom ODBC connector for data stores that are not SQL-capable. The QuickStart connector is a sample DSI implementation of an ODBC connector, written in C++, that reads files in tabbed Unicode text format. The Simba SQLEngine is used to perform SQL processing, allowing the QuickStart connector to retrieve data from the sample text files, which are not a SQL-capable data source.

Using the QuickStart sample connector to prototype a DSI implementation for a custom data store helps developers understand how the Simba SDK works. By removing the shortcuts and simplifications implemented in the QuickStart connector, you can use it as the foundation for a commercial DSI implementation and create a custom ODBC connector for a data store that is not SQL-capable.

The UML diagram below shows a typical design pattern for a DSI implementation:



Notice the circular pattern of class relationships, headed by `iResult` and anchored by `QSUtilities`.

The `iResult` class is responsible for retrieving column data and maintaining a cursor across result rows. The `QSUtilities` class contains platform-specific utility functions.

To implement data retrieval, your `Reader` class interacts directly with your data store to retrieve data, and then deliver the data to the `QSTable` class on demand. The `Reader` class should manage caching, buffering, paging and all the other techniques to speed data access.

As a starting point, to make your connector work properly with Microsoft Excel you can add metadata access by implementing the `QSTypeInfoMetadataSource` class and using the `DSIExtMetadataHelper` class. The `DSIExtMetadataHelper` class is responsible for iterating through tables and stored procedures so the engine can generate catalog function metadata.

## Day One

The Day One instructions explain how to install the Simba SDK, compile the sample ODBC connector, and review the configuration information created at compile time.

After the sample ODBC connector is successfully compiled, it is used to retrieve data from the data source that is included with the Simba SDK. The sample ODBC connector is then used to create the framework for a custom ODBC connector, which is renamed and used to retrieve sample data.

At the end of the day, you will have compiled, built and tested your custom ODBC connector.

### Install the Simba SDK

The Simba SDK is provided as `SimbaEngineSDK [BUILD].tar.gz`, where `[BUILD]` is the build number and platform, for example:

```
SimbaEngineSDK_Release_Linux-x86_10.2.0.1000.tar.gz
```

This file is a `.tar` file that has been compressed using `gzip`.

The Simba SDK installation package includes:

- The Simba SDK implementation in all supported platforms.
- Sample connectors for most supported platforms.
- PDF versions of the documentation.
- An HTML version of the C++ and Java API.

### To install the Simba SDK:

1. Uninstall any previous versions of the Simba SDK.
2. Ensure the evaluation license is installed. See [Install the Evaluation License](#).
3. Open a command prompt and navigate to the directory where the Simba SDK will be installed.
4. Copy the `SimbaEngineSDK [BUILD].tar.gz` file to the Simba SDK installation directory, where `[BUILD]` is the build number and platform.
  - Ensure you copy the correct version of the Simba SDK for your platform. To determine your machine version, type `uname -m` at the command prompt.
5. Uncompress the file by typing the following:

```
gunzip SimbaEngineSDK[BUILD].tar.gz
```

6. Extract the `.tar` file by typing the following:

```
tar -xvf SimbaEngineSDK[BUILD].tar
```

The Simba SDK is installed.

[Install the Evaluation License](#)

[Driver Managers](#)

## Build the Sample ODBC Connector

You can use the sample makefile to build the QuickStart connector. The sample makefile automatically detects the required settings based on your operating system, machine bitness, and compiler. For more information about different makefile options, see *Compiling Your Connector* in the guide [Developing Connectors for Data Stores Without SQL](#).

**To build the Simba SDK QuickStart sample connector:**

The sample connectors included with the Simba SDK are installed in the folder `[INSTALL_DIR]/SimbaEngineSDK/10.2/Examples`, where `[INSTALL_DIR]` is the installation directory. The sample connectors include sample makefiles.

In the following instructions, replace `[INSTALL_DIR]` with the Simba SDK installation directory.

1. Set the **SIMBAENGINE\_DIR** environment variable:  

```
export SIMBAENGINE_DIR=[INSTALL_DIR]/SimbaEngineSDK/10.2/DataAccessComponents
```
2. Set the **SIMBAENGINE\_THIRDPARTY\_DIR** environment variable:  

```
export SIMBAENGINE_THIRDPARTY_DIR=[INSTALL_DIR]/SimbaEngineSDK/10.2/DataAccessComponents/ThirdParty
```
3. Change to the following directory:  

```
[INSTALL_DIR]/SimbaEngineSDK/10.2/Examples/Source/Quickstart/Source
```
4. Type `./mk.sh MODE=debug` to run the makefile for the debug target.

### Important:

Do not use the makefile directly. Use the `mk.sh` script instead.

The resulting library, `libQuickStart<BITNESS>.so`, is put in the following directory:

```
[INSTALL_  
DIR]/SimbaEngineSDK/  
10.2  
/Examples/Source/  
QuickStart/Bin/<BUILD>/<RELEASE|DEBUG><BITNESS>.
```

Where `<BUILD>` is a combination of your operating system, machine bitness, and compiler, `<RELEASE|DEBUG>` is either `release` or `debug`, and `<BITNESS>` is `32`, `64`, or `3264`.

### Example Build Location

```
[INSTALL_  
DIR]/SimbaEngineSDK/10.2/Examples/Source/QuickStart/Bin/Linux_  
x86_gcc/debug64/libQuickStart64.so
```

*Compiling Your Connector* in the guide [Developing Connectors for Data Stores Without SQL](#)

## Configure the Connector and Data Source

You configure the Data Source Name (DSN) in the `odbc.ini` file and the connector in the `odbcinst.ini` file. The driver manager uses information in these files to locate and load ODBC connectors. You configure connector-specific information, such as log file location, in the `.simba.QuickStart.ini` file.

### Driver Managers

On Linux and Unix platforms, the Simba SDK supports both the iODBC and the unixODBC driver managers.

Unlike Windows platforms, Linux, Unix, and macOS platforms do not come with a driver manager as part of the operating system. You need to install your own driver manager before loading a connector and accessing data stores. For more information on installing a driver manager, see [Driver Managers](#).

### Using the Sample Configuration Files

If the `odbc.ini` and `odbcinst.ini` files already exist on your machine, you can add the QuickStart connector entries to these files. If these files do not exist on your machine, you can create them and add the required entries. The Simba SDK includes sample files in the following directory:

```
[INSTALL_DIR]/SimbaEngineSDK/10.2/Documentation/Setup
```

where `[INSTALL_DIR]` is the Simba SDK installation directory.

#### Important:

- The sample `odbc.ini` and `odbcinst.ini` files show how to configure DSNs and connector entries for 32-bit connectors on Linux platforms. To configure a 64-bit connector, a debug connector, or a connector for platforms other than Linux, you will need to make the changes shown below.
- The sample `.simba.quickstart.ini` file is hidden, so you may need to configure your window or terminal listing to view hidden files.

### Configuring an ODBC Data Source in the `odbc.ini` File

ODBC Data Sources are defined in the `.odbc.ini` configuration file. If this file does not exist, you must create it in the correct location.

#### To create and configure the `odbc.ini` file:

1. To see if the `.odbc.ini` file already exists in the home directory, run the following command:

```
ls -al ~ | grep .odbc.ini
```

If no response is returned, the `.odbc.ini` file does not exist in the home directory.

2. To see if the file already exists in another location, check if the `ODBCINI` environment variable is set. If it is set, then check the specified location to see if the file exists.
3. If the file does not exist, copy the sample `odbc.ini` file to the desired directory. This example shows how to copy the file to the user's home directory, but you can choose another directory:

```
cp [INSTALL_
DIR]/SimbaEngineSDK/10.2/Documentation/Setup/odbc.ini
~/odbc.ini
```

4. Using a text editor, open the `~/odbc.ini` configuration file.



**Note:**

The “.” before `odbc.ini` in `~/.odbc.ini` causes the copied file to be hidden. To open the file, you may need to configure your text editor to show hidden files.

5. Replace every instance of `[INSTALLDIR]` with the installation location of the Simba SDK.
6. Make sure there is an entry in the `[ODBC Data Sources]` section that defines the Data Source Name (DSN). The `[ODBC Data Sources]` section is used to specify the available data sources.

**Example:**

```
[ODBC Data Sources]
QuickstartDSII=QuickstartDSIIDriver
```

7. Make sure there is a section with a name that matches the data source name (DSN). This section specifies the configuration options as key-value pairs.

Be sure to use the path to the connector that you built in the previous step, as shown below. By default, the sample `odbcinst.ini` file contains a path to the pre-built connectors that ship with the Simba SDK.

**Example: 32-bit Connector**

```
[QuickstartDSII]
Description=Sample 32-bit SimbaEngine QuickStart DSII
DBF=[INSTALL_
DIR]/SimbaEngineSDK/10.2/Examples/Databases/Quickstart
Driver=[INSTALL_DIR]/SimbaEngineSDK/10.2/Examples/
Source/Quickstart/Bin/Linux_x86_
gcc/debug32/libQuickstart32.so
```

Note: replace `[INSTALL_DIR]` with the installation directory.

**Example: 64-bit Connector**

```
[QuickstartDSII]
Description=Sample 64-bit SimbaEngine QuickStart DSII
DBF=[INSTALL_
DIR]/SimbaEngineSDK/10.2/Examples/Databases/Quickstart/
Driver=[INSTALL_DIR]/SimbaEngineSDK/10.2/Examples/
Source/Quickstart/Bin/Linux_x86_
gcc/debug64/libQuickstart64.so
Locale=en-US
```

Note: replace [INSTALL\_DIR] with the installation directory.

**⚠ Important:**

Take care not to include characters such as spaces or tabs at the end of a value. This can cause failure to connect.

## Configuring an ODBC Connector in the `odbcinst.ini` File

ODBC Connectors are defined in the `odbcinst.ini` configuration file. This configuration is optional, because connectors can be specified directly in the `odbc.ini` configuration file as discussed in the previous section. Before defining the connector, see if the `odbcinst.ini` file exists.

### To configure a connector in the `.odbcinst.ini` file:

1. To see if the `.odbcinst.ini` file exists in your home directory, type the following command:

```
ls -al ~ | grep .odbcinst.ini
```

If the file does not exist, then nothing is returned at the command prompt.

2. To see if the file already exists in another location, check if the `ODBCINSTINI` environment variable is set. If it is set, then check the specified location to see if the file exists.
3. If the file does not exist, copy the sample `odbcinst.ini` file to the desired directory. This example shows how to copy the file to the user's home directory, but you can choose another directory as described in [Configure the Connector and Data Source](#):

```
cp [INSTALL_DIR]/SimbaEngineSDK/10.2/Documentation/Setup/odbcinst.ini ~/.odbcinst.ini
```

where `[INSTALL_DIR]` is the Simba SDK installation directory.

4. Open the `.odbcinst.ini` configuration file in a text editor.
5. Replace every instance of `[INSTALLDIR]` with the installation location of the Simba SDK.
6. Add a new entry to the `[ODBC Drivers]` section. The `[ODBC Drivers]` section is used to specify the available connectors. Enter the connector name and the value `Installed`. The connector name should be used for the "Driver" value in the data source definition instead of the connector shared library name.

**Example:**

```
[ODBC Drivers]
QuickstartDSIIDriver=Installed
```

7. Add a new section with a name that matches the new connector name. This section will contain the configuration options specified as key-value pairs.

Be sure to use the path to the connector that you built in the previous step, as shown below. By default, the sample `odbcinst.ini` file contains a path to the pre-built connectors that ship with the Simba SDK.

**Example - 32-bit Linux:**

```
[QuickstartDSIIDriver]
APILevel=1
ConnectFunctions=YYY
Description=Sample 32-bit SimbaEngine QuickStart DSII
Driver=[INSTALL_DIR]/SimbaEngineSDK/10.2/Examples/
Source/Quickstart/Bin/Linux_x86_
gcc/debug32/libQuickstart32.so
DriverODBCVer=03.80
SQLLevel=1
```

Note: replace `[INSTALL_DIR]` with the installation directory.

**Configuring the QuickStart Connector in the `simba.ini` File**

Configuration options for connectors created using the Simba SDK appear in the connector's configuration file, which by default is called `.simba.drivervname.ini`. This file specifies connector-wide settings. The Simba SDK searches for the connector `.ini` file in the following locations, in the specified order:

1. The path, including the file name, specified using the **SIMBAINI** environment variable.
2. The connector directory, as a non-hidden `.ini` file.
3. The directory that the client application is launched from.
4. In `$HOME`, as a hidden `.ini` file
5. In `/etc/` as a non-hidden `.ini` file

**Note:**

Examples in this section assume that the `.ini` file exists in the user's home directory.

## To configure the QuickStart Connector in the `.simba.quickstart.ini`:

1. To see if the `.simba.quickstart.ini` file already exists in your home directory, type the following command:

```
ls -al ~ | grep .simba.quickstart.ini
```

If the file does not exist, then nothing is returned at the command prompt.

2. If the file does not exist, copy the sample `.simba.quickstart.ini` file to the desired directory. Note that the sample `.simba.quickstart.ini` file is hidden. This example shows how to copy the file to the user's home directory, but you can choose another directory:

```
cp [INSTALL_DIR]/SimbaEngineSDK/10.2/Documentation/  
Setup/.simba.quickstart.ini ~/.simba.quickstart.ini
```

3. Open the `~/.simba.quickstart.ini` configuration file in a text editor.
4. Replace every instance of `[INSTALLDIR]` with the installation location of the Simba SDK.
5. Edit the `DriverManagerEncoding` setting:

### **Note:**

- This step is optional, because the Simba SDK automatically detects the type and version of the driver manager.
- If you are unsure which driver manager you are using, or where the driver manager is installed, contact your system administrator or see [Driver Managers](#) for more information.

- If you are using the iODBC ODBC driver manager set the `DriverManagerEncoding` setting to `UTF-32`.
  - Or, if you are using the unixODBC ODBC driver manager, check which setting to use:
    - a. At a command prompt, type `odbc_config --cflags`.
    - b. If you see the `"DSQL_WCHART_CONVERT"` flag, then set `DriverManagerEncoding` to `UTF-32`.
    - c. Otherwise, set `DriverManagerEncoding` to `UTF-16`.  
For more information about your ODBC driver manager, consult your system administrator or your ODBC Driver Manager documentation.
6. Edit the `ErrorMessagesPath` setting to replace `[INSTALLDIR]` with your install directory.
  7. Set the `ODBCInstLib` to the absolute path of the `ODBCInst` library for the Driver Manager that you are using.

**Note:**

This step is optional, because the Simba SDK automatically detects the type and version of the ODBCInst library. Set this value only if you want to override the value that is automatically detected.

The ODBCInst library is a part of the driver manager but is used by the connector to read values from the odbc.ini file. The value of this key is the absolute path of the ODBCInst library and depends on which driver manager you are using. For example, for the iODBC Driver Manager this would be `<driver manager dir>/lib/libiodbcinst.so` (notice the "i" after the lib) and for unixODBC this would be `<driver manager dir>/lib/libodbcinst.so`

**Example: the iODBC Driver Manager**

(notice the i after lib)

```
ODBCInstLib=<driver manager dir>/lib/libiodbcinst.so
```

**Example: the unixODBC Driver Manager**

```
ODBCInstLib=<driver manager dir>/lib/libodbcinst.so
```

8. Save the file.

[Locating the Configuration Files](#)[Driver Managers](#)[Troubleshooting](#)

## Connect to the Data Source

In order for the QuickStart sample connector to connect to the sample database successfully, the library path environment variable must include references to the OpenSSL and ICU libraries.

OpenSSL is used by SimbaClient for ODBC. Your custom ODBC connector may not require this library.

**Note:**

- You must have a driver manager installed. See [Driver Managers](#).
- You must use a 64-bit driver manager with a 64-bit connector, or a 32-bit driver manager with a 32-bit connector.

**To add the ICU and OpenSSL libraries:**

- Type the following at the command line, replacing [INSTALLDIR], <PLATFORM>, and <CONFIGURATION>. Note these libraries are supported for all versions of Unix and Linux, not just for CentOS:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:  
[INSTALLDIR]/SimbaEngineSDK/  
10.2  
/DataAccessComponents/ThirdParty/icu/53.1.x/<  
PLATFORM>/<CONFIGURATION>/lib:  
[INSTALLDIR]/SimbaEngineSDK/  
10.2  
/DataAccessComponents/ThirdParty/openssl/1.1.0/<  
PLATFORM>/<CONFIGURATION>/lib
```

**Note:**

This document uses the LD\_LIBRARY\_PATH environment variable. Some Unix platforms use a different environment variable:

- Use SHLIB\_PATH on HP/UX
- Use LIBPATH on AIX

**To test connecting the QuickStart sample connector to the sample data source:**

This procedure uses the iodbctest utility that is included with the iODBC driver manager. For help, see [Troubleshooting](#).

1. At the command prompt, type `iodbctest`.
2. At the prompt that says “Enter ODBC connect string”, type `?` to show the list of DSNs and Drivers.

The list contains QuickStartDSII DSN.

3. To connect to your data source, type: `DSN=QuickStartDSII`

The prompt `SQL>` appears.

4. Type a SQL command to query the database. For example:

```
SELECT * FROM PRODUCT
```

The SQL results are returned.

5. To quit `iodbctest`, type `quit` at the prompt.

Once the QuickStart sample connector project is built, it can be used to create a custom connector project.

## [Troubleshooting](#)

### Set up a Custom ODBC Connector Project

Once the QuickStart project has been built and tested, you can create a custom project for your ODBC connector.

#### Important:

It is very important that you create your own project directory. You might be tempted to simply modify the sample project files, but we strongly recommend that you create your own project directory. If you simply modify the sample project files:

- All your changes will be lost when you install a new version of the SDK.
- You will lose your frame of reference for debugging.  
There may be times, for debugging purposes, that you will need to see if the same error occurs using the sample connectors. If you have modified the sample connectors, this won't be possible.

#### To create a custom ODBC connector project based on the QuickStart sample connector:

1. Ensure you are working in a window or shell where the `SIMBAENGINE_DIR` and `SIMBAENGINE_THIRDPARTY_DIR` environment variables are set, as explained in [Build the Sample ODBC Connector](#).
2. Copy the `Quickstart` directory to create a new top-level directory for your custom ODBC connector project. Be sure to copy hidden files and symlinks too, for example:

```
mkdir [INSTALL_
DIR]/SimbaEngineSDK/10.2/Examples/Source/MyQuickstart
cp -a [INSTALL_
DIR]/SimbaEngineSDK/10.2/Examples/Source/Quickstart/.
[INSTALL_
DIR]/SimbaEngineSDK/10.2/Examples/Source/MyQuickstart
```

where `[INSTALL_DIR]` is the Simba SDK installation directory.

This new directory is referred to as the `[Project]` directory in the following steps.

3. In the `[Project]/Source` directory, open the `GNUmakefile` in a text editor.
4. Replace the `target.driver` target name with the name of your custom connector.

#### Example:

Replace this line: `target.driver = libQuickstart${BITS}.${SO}`

With this line: `target.driver = libMyQuickstart${BITS}.${SO}`

5. Save and close the file.
6. In the `[Project]/Source` directory, run the following command to build your custom ODBC connector:

```
./mk.sh MODE=debug
```

Your custom ODBC connector project is built.

*Compiling Your Connector* in the guide [Developing Connectors for Data Stores Without SQL](#)

### [Troubleshooting](#)

## Configure Your Custom Connector and Data Source

The Simba SDK supports both the iODBC and the unixODBC driver managers. The driver manager uses configuration files to define and configure ODBC data sources and connectors. The `odbc.ini` file defines ODBC data sources, or DSNs, and the `odbcinst.ini` file defines ODBC connectors. The `.simba.quickstart.ini` file defines connector-specific properties such as logging.

#### **Note:**

The examples in this section assume your custom ODBC connector is a 32-bit connector named `MyQuickStart`.



**To configure the custom connector and data source:**

1. In the user's home directory, open the `.odbc.ini` configuration file in a text editor.
2. Make sure there is an entry in the `[ODBC Data Sources]` section that defines the data source name (DSN) of your custom ODBC connector.

**Example:**

```
[ODBC Data Sources]
MyQuickstartDSII=MyQuickstartDSIIDriver
```

3. Make sure there is a section with a name that matches the data source name (DSN).

**Example: 32-bit Linux**

```
[MyQuickStartDSII]
Description=64-bit MyQuickStart DSII
DBF=[INSTALL_
DIR]/SimbaEngineSDK/10.2/Examples/Databases/QuickStart
Driver=
[INSTALLDIR]/SimbaEngineSDK/
10.2/Examples/Builds/Bin/Linux_x86_
gcc/release32/libMyQuickStart32.so
```

Note: replace `[INSTALL_DIR]` with the installation directory.

4. Open the `.odbcinst.ini` configuration file in a text editor.
5. Add a new entry to the `[ODBC Drivers]` section.

**Example:**

```
[ODBC Drivers]
MyQuickstartDSIIDriver=Installed
```

6. Add a new section with a name that matches the new connector name.

**Example: 32-bit Linux**

```
[MyQuickstartDSIIDriver]
Driver=
[INSTALLDIR]/SimbaEngineSDK/
10.2/Examples/Builds/Bin/Linux_x86_
gcc/release32/libMyQuickStart32.so
```

Your custom ODBC connector and data source are configured.

[Configure the Connector and Data Source](#)

## Debug Your Custom Connector

Using a debugger to step through the custom connector code can give developers a better understanding of the connector's functionality. This section explains how to use the `iodbctest` application to connect to your custom ODBC connector, then use the `gdb` debugger to step through the connector code.

### To debug your custom ODBC connector:

1. Follow the instructions in [Connect to the Data Source](#) to use `iodbctest` to connect to the custom connector. Use the name of the custom connector instead of the QuickStart connector.
2. To quit `iodbctest`, type `quit` at the command prompt.
3. To start the debugger, type `gdb iodbctest`.
4. Type `break Simba::DSI::DSIDriverFactory`.

This will set a breakpoint at the `DSIDriverFactory()` function in the `Main_Unix.cpp` file. This is a good breakpoint to start with, because this function runs as soon as the driver manager loads the ODBC connector. To set a different breakpoint, view the source code in your custom project:

```
[INSTALL_DIR]/SimbaEngineSDK/10.2/Examples/Source/[your project name]/Source/
```

5. Type `run DSN=MySimba SDKDSII`, replacing `MySimba SDKDSII` with the DSN for your custom ODBC connector.

The program runs until the breakpoint is encountered.

#### **Note:**

When using the `gdb` debugger with an application, the ODBC connector is not loaded until the application is running and a connection is made. This means that breakpoints can be set either before or after the connector is loaded, depending on which breakpoint you want to hit.

This step verifies that the custom connector, based on the QuickStart project, is correctly installed and configured, and that the development system is properly set up.

## Summary of Day One

You have successfully completed the following tasks:

- Built and tested the QuickStart sample connector. This verifies that your installation and development environment are properly configured.

- Created, built, and tested a custom connector project by copying the QuickStart connector.  
You can use this project as a framework to create your custom ODBC connector.

## ODBC Troubleshooting: How to Enable Driver-manager Tracing

### [Troubleshooting](#)

## Enable Logging

You can turn on logging for your custom ODBC connector. By setting the log level to `Trace`, you can gain a better understanding of how your custom ODBC connector works.

**To enable logging in your custom ODBC connector:**

1. Open the `.simba.quickstart.ini` file in the user's home directory.

### **Note:**

Your custom ODBC connector is not yet completely rebranded, so configuration information is read from the `.simba.quickstart.ini` file. The name of this file is set using `#define DRIVER_LINUX_BRANDING` in the `Main_Unix.cpp` file.

2. Set the `LogLevel` to 6 for trace level, or another level if you prefer.
3. Set the `LogPath` to the directory to use for log files.

### **Example:**

```
LogLevel=6
LogPath=/usr/tmp/myquickstart/logs
```

The log files are created the next time the connector is used.

## Log File Format

The log files have the following format, where `[Message]` is optional:

```
[Date] [Log Level] [Thread ID] [Class] [Message]
```

### **Example:**

```
Jun 15 14:05:12.017 INFO 9864
ConnectionSettings::LoadSettings: ConnString setting: "DSN" =
"MyQuickstartDSII"
```

## Summary of Day One

You have successfully completed the following tasks:

- Built and tested the QuickStart sample connector.  
This verifies that your installation and development environment are properly configured.
- Created, built, and tested a custom connector project by copying the QuickStart connector.  
You can use this project as a framework to create your custom ODBC connector.

<http://www.simba.com/resources/sdk/knowledge-base/enable-logging-in-odbc/>

<http://www.simba.com/resources/sdk/knowledge-base/simbaengine-logging/>

## Day Two

Day Two instructions explain how to customize your ODBC connector, enable logging, and establish a connection to your data store.

### Set the Configuration Branding

The `DSIDriverFactory()` implementation in `Main_Unix.cpp` is the main entry point that is called from Simba's ODBC layer to create an instance of the DSI implementation. This method is called as soon as the Driver Manager calls `LoadLibrary()` on the ODBC connector shared object.

To construct the connector singleton:

1. In your custom ODBC connector project, open the file `Main_Unix.cpp`.
2. Navigate to the line **TODO #1: Construct connector singleton**.
3. Look at the `DSIDriverFactory()` implementation, and locate the following line of code:  

```
SimbaSettingReader::SetConfigurationBranding(DRIVER_LINUX_BRANDING);
```
4. In the default implementation, `DRIVER_LINUX_BRANDING` defines the string `"simba.QuickStart.ini"`. This is the name of the `.ini` file that specifies the connector settings.
5. Change this string to the name of a configuration file reflecting the name of your connector or company.
6. Save the `Main_Unix.cpp` file.
7. Update the name of the `"simba.QuickStart.ini"` file to match `DRIVER_LINUX_BRANDING`.

### Set Connector Properties

To set connector properties:

1. Open the file `QSDriver.cpp` file and navigate to the line **TODO #2 Set the connector properties**.
2. Go to the method `SetDriverPropertyValues()`, where the general properties for the connector are set. Change the properties described below:

Property	Description
DSI_DRIVER_DRIVER_NAME	Set this property to the name of the connector (the same name used to replace QuickStartDSII in Day One). This is the connector name that is shown to the application.
DSI_DRIVER_STRING_DATA_ENCODING	Optional. The encoding of char data from the perspective of the data store. Depending on the character sets, this property may need to be changed.
DSI_DRIVER_WIDE_STRING_DATA_ENCODING	Optional. The encoding of wide character data from the perspective of the data store. Depending on the character sets, this property may need to be changed.

## Set Logging Details

This section explains how to set the connector-wide and connection-wide logging.

### To set logging details:

1. Open the file `QSDriver.cpp` and navigate to the line **TODO #3 Set the connector-wide logging details**.
2. Change the connector log's file name.
3. Open the file `QSConnection.cpp` and navigate to the line **TODO #4 Set the connection-wide logging details**.
4. Change the connection log's file name. Under default behavior this is routed to the connector logs. If you want to log connections differently you need to override the `QSConnection::GetLog()` function.
5. Click **Save All**.

**Note:**

By default, the QuickStart connector maintains two kinds of log files: one for all connector-based calls and one for each connection created. Update these sections if you do not require such fine granularity in logging.

For more information about how to enable logging, see [Developing Connectors for Data Stores Without SQL](#).

## Check Connection Settings

When the Simba ODBC layer is given a connection string from an ODBC-enabled application, the Simba ODBC layer parses the connection string into key-value pairs. The entries in the connection string and the DSN are then sent to the `QSConnection::UpdateConnectionSettings()` function for validation.

If entries of the connection string overlap entries from the DSN, then the connection string will override parameters from the DSN. To pass additional parameters to your DSII, simply add new parameters to the connection string, or add new entries to the DSN entry. These values will automatically be picked up by the SDK and passed through for use by your DSII.

`UpdateConnectionSettings()` receives all the incoming connection settings that are specified in the DSN that was used to establish the connection. The role of this function is to ensure that all of the required, and any optional, settings are present. Note that actual data validation of the settings should be done in the `Connect()` function.

**Example:**

The connection string “DSN=QuickStart;UID=user;” will be broken down into key value pairs and passed in via the `DSIConnSettingRequestMap` parameter. In this case that map would contain two entries: `{DSN, QuickStart}` and `{UID, user}`. If a DSN was specified, then the DSN value is removed from the map and any entries that are stored in the preconfigured DSN are inserted into the map. Once the map has been created with all the key-value pairs from the connection string and DSN, this map is passed down to the DSII.

**To check the connection settings for the custom connector:**

1. Open the file `QSConnection.cpp` and navigate to the line **TODO #5 Check Connection Settings**.
2. Modify the `UpdateConnectionSettings()` function to validate that the settings (key-value pairs) in the `DSIConnSettingRequestMap` are sufficient

to create a connection. Any settings that are not present should be added to the `DSIConnSettingResponseMap` parameter.

We recommend using the `VerifyRequiredSetting()` or `VerifyOptionalSetting()` functions to perform this verification. These functions also add missing settings to `DSIConnSettingResponseMap`.

**Note:**

The connection settings listed in `UpdateConnectionSettings()` are specific to the QuickStart connector. A custom connector will require different settings.

**Example - QuickStart connector:**

The QuickStart connector verifies that the settings contained in `in_connectionSettings` are sufficient to create a connection by using the following code:

```
VerifyRequiredSetting(QS_DBF_KEY, in_connectionSettings,
out_connectionSettings);
VerifyOptionalSetting(QS_DEFAULT_MAX_COLUMN_SIZE_KEY, in_
connectionSettings, out_connectionSettings);
VerifyOptionalSetting(QS_USE_CUSTOM_STATES_KEY, in_
connectionSettings, out_connectionSettings);
VerifyOptionalSetting(QS_USE_CACHING_KEY, in_
connectionSettings, out_connectionSettings);
VerifyOptionalSetting(QS_LOCALE, in_connectionSettings,
out_connectionSettings);
```

The QuickStart connector requires a single key in the DSN, `DBF`. The other keys are optional.

3. If any required values are missing, the connector will either fail to connect, or will call `PromptDialog()`, depending on the connection settings. If all required values exist, then `Connect()` will be called.
4. If any of the values received are invalid, then the code should throw an `ErrorException` seeded with `DIAG_INVALID_AUTH_SPEC`.

**Manually verifying the connection settings**

Settings can also be verified manually. If the entries within `in_connectionSettings` are not sufficient to create a connection, then the connector can ask for additional information from the ODBC-enabled application by manually specifying the additional, required settings in `out_connectionSettings`. If there are no further entries required, simply leave `out_connectionSettings` empty.



For more information on ODBC connections, see the Knowledge Base article *DSII Connection Process for ODBC* at <http://www.simba.com/resources/sdk/knowledge-base/dsii-connection-process-for-odbc>.

## Establish a Connection

The Simba SDK calls `UpdateConnectionSettings()` before calling `QSCConnection::Connect()`. Once `QSCConnection::UpdateConnectionSettings()` returns `out_connectionSettings` without any required settings—if there are only optional settings, a connection can still occur—then the Simba ODBC layer calls `QSCConnection::Connect()`, passing in all the connection settings received from the application.

During `Connect()`, the connector should have all the settings necessary to make a connection as verified by `UpdateConnectionSettings()`. You can use the utility functions `GetRequiredSetting()` and `GetOptionalSetting()` to request the required and optional settings for your connection, and attempt to make an actual connection.

### To establish a connection:

1. Open the file `QSCConnection.cpp` and navigate to the line **TODO #6 Establish A Connection**.
2. Look at the code that authenticates the user against your data store using the information provided within the `in_connectionSettings` parameter. Use `GetRequiredSetting()` and `GetOptionalSetting()` to access the settings in the map.
3. Add validation to your custom ODBC connector. If authentication fails, throw an error. Note that the sample ODBC connector does not perform validation.

The user is now authenticated against your data store.

## Summary of Day Two

You have successfully authenticated the user against your data store and established a connection.

## Day Three

The Day Three instructions explain how to return the data used to pass catalog information back to the ODBC-enabled application.

### Create and Return Metadata Sources

Your custom ODBC connector uses metadata sources, provided by the Simba SDK, to handle SQL catalog functions.

#### Overview of SQL Catalog Functions

ODBC applications need to understand the structure of a data store in order to execute SQL queries against it. This information is provided using catalog functions. For example, an application might request a result set containing information about all the tables in the data store, or all the columns in a particular table. Each catalog function returns data as a result set.

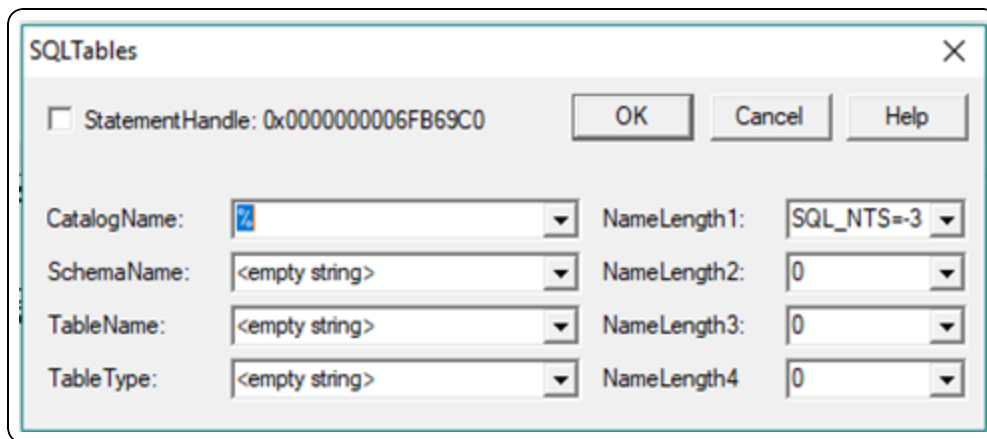
Most ODBC-enabled applications require a connector to implement the following catalog functions. You may wish to implement additional catalog functions in your custom connector.


Catalog Function	Description
SQLGetTypeInfo	Returns information about data types supported by the data source.
SQLTables (CATALOG_ONLY)	If CatalogName is SQL_ALL_CATALOGS and SchemaName and TableName are empty strings, the result set contains a list of valid catalogs for the data source. (All columns except the TABLE_CAT column contain NULLs.)
SQLTables (SCHEMA_ONLY)	If SchemaName is SQL_ALL_SCHEMAS and CatalogName and TableName are empty strings, the result set contains a list of valid schemas for the data source. (All columns except the TABLE_SCHEM column contain NULLs.)

Catalog Function	Description
SQLTables (TABLE_TYPE_ONLY)	If TableType is SQL_ALL_TABLE_TYPES and CatalogName, SchemaName, and TableName are empty strings, the result set contains a list of valid table types for the data source. (All columns except the TABLE_TYPE column contain NULLs.)
SQLTables	Returns the list of table, catalog, or schema names, and table types, stored in a specific data source.
SQLColumns	Returns a list of columns in one or more tables.

### Example: Using Catalog Functions with the QuickStartconnector

1. In the ODBC Test application, connect to the QuickStart connector.
2. To send the SQLTables (CATALOG\_ONLY) catalog function, select **Catalog > SQLTables**.
3. Enter SQL\_ALL\_CATALOGS for the **CatalogName**, then select the correct value for **NameLength1**. For example:



4. Click **OK**.
5. Select  to retrieve the results.

The following list of valid catalogs for the QuickStart data source are returned: "TABLE\_QUALIFIER", "TABLE\_OWNER", "TABLE\_NAME", "TABLE\_TYPE", "REMARKS"

For more information on SQL catalog functions, see [https://msdn.microsoft.com/en-us/library/ms713520\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/ms713520(v=vs.85).aspx).

## Implementing Metadata Sources to Handle Catalog Functions

SQL catalog functions are represented in the DSI by metadata sources: there is one metadata source for each of the catalog functions.

`QSDatabase::MakeNewMetadataTable()` is responsible for creating the metadata sources. Metadata sources are used to return the catalog metadata about your data store to the ODBC application for the ODBC catalog functions.

Open the file `QSDatabase.cpp` and navigate to the line **TODO #7 Create and return your Metadata Sources**.

There is one metadata source for each of the catalog functions. For example, when the application calls `SQLColumns()`, a `DSI_COLUMNS_METADATA` source is created to return the list of columns in one or more tables in the data store.

Each ODBC catalog function is mapped to a unique `DSIMetadataTableId`, which is then mapped to an underlying `MetadataSource` that the connector implements and returns. Each `MetadataSource` instance is responsible for the following:

1. Creating a data structure that holds the data relevant for the custom data store: `Constructor`
2. Navigating the structure on a row-by-row basis: `Move()`
3. Retrieving data: `GetData()` (See [Data Retrieval](#) for a brief overview of data retrieval).

### Required Metadata Sources

All custom ODBC connectors must implement the following metadata sources, as they are required by ODBC applications:

Metadata Source	Description
<code>DSI_TABLES_METADATA</code>	List of all tables defined in the data source. This source is constructed via the <code>QSMetadataHelper</code> and SQL Engine.
<code>DSI_CATALOGONLY_METADATA</code>	List of all catalogs defined in the data source, if catalogs are supported. This source is constructed via the <code>QSMetadataHelper</code> and SQL Engine.

Metadata Source	Description
DSI_SCHEMA_METADATA	List of all schemas defined in the data source. This source is constructed via the QSMetadataHelper and SQL Engine. QuickStartDSII does not support schemas.
DSI_TABLETYPEONLY_METADATA	List of all table types (TABLE,VIEW,SYSTEM) defined within the data source. This source is constructed via the QSMetadataHelper and SQL Engine.
DSI_COLUMNS_METADATA	List of all columns defined across all tables in the data source. This source is constructed via the QSMetadataHelper and SQL Engine.
DSI_TYPE_INFO_METADATA	List of the supported types by the data source. This means the actual types that can be stored in the data source, not necessarily the types that can be returned by the connector. For instance, a conversion may result in a type being returned that is not stored in the data source.

Most catalog types are created using the metadata helper in conjunction with the SQL Engine.

### Handling DSI\_TYPE\_INFO\_METADATA

The underlying ODBC catalog function `SQLGetTypeInfo` is handled as follows:

- When called with `DSI_TYPE_INFO_METADATA`, `QSDatabase::MakeNewMetadataTable()` will return an instance of `QTypeInfoMetadataSource()`.
- The QuickStart sample connector exposes support for all data types, but due to its underlying file format, it is constrained to support only the following types:
  - `SQL_BIGINT`
  - `SQL_DECIMAL`
  - `SQL_LONGVARCHAR`
  - `SQL_REAL`
  - `SQL_SMALLINT`
  - `SQL_TYPE_TIME`
  - `SQL_WCHAR`
  - `SQL_CHAR`
  - `SQL_INTEGER`

- SQL\_TYPE\_DATE
- SQL\_VARCHAR
- SQL\_BIT
- SQL\_DOUBLE
- SQL\_LONGVARCHAR
- SQL\_NUMERIC
- SQL\_TINYINT
- SQL\_TYPE\_TIMESTAMP
- SQL\_WVARCHAR

3. For your connector, you may need to change the types returned and the parameters for the types in `QSTypeInfoMetadataSource::PrepareType()`. Change the passed in `SqlTypeInfo` object to modify the parameters of the types that are supported.

### Handling the Other MetadataSources

The other ODBC catalog functions, including `SQLTables (CATALOG_ONLY)`, `SQLTables (TABLE_TYPE_ONLY)`, `SQLTables (SCHEMA_ONLY)`, `SQLTables` and `SQLColumns`, are handled automatically by the metadata helper class.

When these functions are called with any other `DSIMetadataTableId`, `QSDatabaseEngine::MakeNewMetadataTable()` should return `null`. Returning `null` tells the Simba SDK to use the metadata helper class returned via `QSDatabaseEngine::CreateMetadataHelper()` along with some default `MetadataSources` to create the data source metadata. You can also choose to return a `DSIMetadataSource` if you don't want to use the metadata helper.

Change the following methods:

- `QSMetadataHelper::QSMetadataHelper()`  
The example constructor retrieves a list of the tables in the data source. Modify this method to load the tables defined in your data store.
- `QSMetadataHelper::GetNextTable()`  
In the QuickStart connector, this method returns the next table in the data source. Modify this method to retrieve the next table from your data store.

The `DSIExtMetadataHelper` class works by retrieving the identifying information for each table and then opening the table via `QSDatabaseEngine::OpenTable()`. After `QSTable` is implemented, the correct metadata will be returned for all of the tables and columns in your data source.

You can now retrieve type metadata from your data store.

**i Tip:**

On Linux, Unix, and macOS platforms, this metadata is available using the `datatypes` command in the `iodbctest` utility. As well, the `SQLTables` catalog function is available using the `tables` command.

You can use these commands to test your implementation of Day Three.

*Fetching Metadata for Catalog Functions in [Developing Connectors for Data Stores Without SQL](#)*

**Summary of Day Three**

Your custom ODBC connector can now return type metadata. You can use a ODBC-enabled application to connect to your connector and retrieve type metadata from within your data store

## Day Four

Day Four instructions explain how to enable data retrieval from within the connector.

### Enable Data Retrieval

`QSDataEngine::OpenTable()` is the entry point where the Simba SQL Engine requests that tables involved in the query be opened. It is called during the preparation of a SQL statement.

#### **Note:**

The SQL Engine component of the Simba SDK allows applications to execute SQL commands on data stores that are not SQL-capable.

Open the file `QSDataEngine.cpp` and navigate to the line **TODO #8: Open A Table** to go to the relevant section of code.

`QSTable` is an implementation of `DSIExtSimpleResultSet`, an abstract class that provides basic forward-only result set traversal. The main role of `QSTable` is to translate the stored data from your native data format into SQL Data types.

The QuickStart sample connector is implemented for Tabbed Unicode Files. The sample connector translates the text from UTF16-LE strings into the SQL Data types defined for each column within the configuration dialog.

In the QuickStart connector, `QSTable` uses a `TabbedUnicodeFileReader`, which provides an interface to navigate between lines within a Unicode text file. This class preprocesses each row in the file to determine the starting file offset of each column in the row. Its `GetData` method takes a `columnIndex` and uses it to calculate the exact position in the file where the column's data resides. The method repositions the file and retrieves the data as if from a byte-buffer. See [Data Retrieval](#) for a brief overview of data retrieval.

The following sections explain how to implement data retrieval in your custom ODBC connector.

### Modify the OpenTable Method

The `QSDataEngine::OpenTable()` method is called during the preparation of a SQL statement. Modify this method to check that the supplied catalog, schema and table names are valid and correspond to a table defined in your data store. If the inputs are not valid, return `null` to indicate that the table does not exist. If the inputs are valid, return a new instance of `QSTable`.



## Modify QSTable

This section tells you how to modify `QSTable` so that it can work with your data store.

### Return the catalog, schema and table names for your table

Make the following changes to your custom ODBC connector so that it can work with your data store:

- `QSTable::QSTable()`: Modify the constructor to take in the catalog, schema and table names and save them in member variables.
- `QSTable::GetCatalogName()`: Returns `QS_CATALOG`
- `QSTable::GetSchemaName()`: Returns `simba_wstring()` (because it does not support schemas)
- `QSTable::GetTableName()`: Returns `m_tableName`

### Return the columns defined for your table

Modify `QSTable::InitializeColumns()` so that, for each column defined in the table, you define a `DSIResultSetColumn` in terms of SQL types.

**Example: pseudo code for a custom `QSTable::InitializeColumns()`**

For Each Defined Column

```
{
    AutoPtr<DSIColumnMetadata> columnMetadata (
        new DSIColumnMetadata());
    columnMetadata->m_catalogName = m_catalogName;
    columnMetadata->m_schemaName = m_schemaName;
    columnMetadata->m_tableName = m_tableName;
    columnMetadata->m_name = //column name
    columnMetadata->m_label = //localized column name
    columnMetadata->m_unnamed = false;
    columnMetadata->m_charOrBinarySize = //the length in
    bytes
    columnMetadata->m_nullable = DSI_NULLABLE;
    // Change the first parameter of this method to the
    SQL
    // type that maps to your data store type.
    SqlTypeMetadata* sqlTypeMetadata =
    SqlTypeMetadataFactory::MakeNewSqlTypeMetadata (
    SQL_WVARCHAR, TDW_BUFFER_OWNED);
    columns->AddColumn (
    new DSIResultSetColumn (
    sqlTypeMetadata,
```

```
        columnMetadata.Detach() );  
    }  
    m_columns.Attach(columns.Detach() );
```

## Implement Navigation

The methods listed below are responsible for navigating a data structure containing information about one table in the data store, and retrieving data from the table. Modify these methods for your data store:

- `QSTable::MoveToBeforeFirstRow()`
- `QSTable::MoveToNextRow()`
- `TabbedUnicodeFileReader::GetData()`

In your custom implementation:

- It is best to implement a streaming interface for the data in the table within your data store.
- Provide the ability to navigate forward from one table row to the next.
- Provide the ability to navigate across columns within the row.
- Provide the ability to read the data associated with the current row and column combination.

## Modify the DoCloseCursor Method

`QSTable::DoCloseCursor()` is a callback method called from Simba SQL Engine to indicate that data retrieval has completed, and that you may now do any tasks related to closing the connection to your data store.

On Linux and UNIX platforms, lists of catalogs, schemas, tables and types are available using the `qualifiers`, `owners`, `tables` and `types` commands in the `iodbctest` utility.

## Summary of Day Four

You can now execute queries and retrieve data from your data store. You can use any ODBC-enabled application to execute queries and see the results returned from your data store.

You can also run `SQLTables()` and `SQLColumns()` from within `ODBCTest32.exe` (Unicode) and see the correct metadata returned.

## Day Five

Day Five instructions explain how to rebrand your custom ODBC connector.

### Rebrand Error Messages

Error messages sent by the connector are visible to applications and customers. In the QuickStart sample connector, error messages are branded with `QuickStart`, `QS`, and `Simba`. This section explains how to rebrand the error messages to reflect the custom connector name and the company name.

All the error messages used within the DSI implementation are stored in a file called `QSMessages.xml`.

#### To configure error messages:

1. Rename the `QSMessages.xml` file to reflect the name of your company or your custom ODBC connector.
2. Open the file `QSDriver.cpp` and navigate to the **TODO #9 Register the QSMessages.xml file for handling by DSIMessageSource** message to go to the relevant section of code.
3. Update the line associated with the **TODO** to match the new name of the `QSMessages.xml` file.
4. Open the `QSMessages.xml` file and change all instances of the following items:
  - Change the letters `QS` to an appropriate two-letter abbreviation.
  - Change the word `QuickStart` to an appropriate name for your custom connector.
5. For each exception thrown within the custom DSI implementation, change the parameters to match your custom connector name. This rebrands the error messages to reflect the name of your connector.
6. Open the file `QSDriver.cpp` and navigate to the **TODO #10 Set the vendor name, which will be prepended to error messages** message to jump to the relevant section of code.
7. The vendor name is prepended to all error messages that are visible to applications. As explained in the code comments, change the vendor name from `Simba` to an appropriate name for your company.

#### How can I update the vendor name in the Tableau Datasource Connection (TDC) file?

A TDC file contains configuration information that will be applied to any Tableau connection that matches the database vendor name and connector name described in

the TDC file.

### To set the vendor name for your custom ODBC connector:

1. Ensure you have set your vendor name as described in **TODO #10**.
2. In the class that extends `DSIConnection` (`QSCConnection` in our sample QuickStart Connector), set the property `DSI_CONN_DBMS_NAME` to your vendor name.

#### Example:

```
SetProperty(DSI_CONN_DBMS_NAME,  
AttributeData::MakeNewWStringAttributeData  
("YourVendorName"));
```

By default, the value of the `DSI_CONN_DBMS_NAME` property is `TEXT`.

3. Set the vendor name in the TDC file by following the instructions on Tableau's website.

These steps allow Tableau to match the vendor name in the TDC file with the associated `SQLGetInfo()` property it queries the connector for.

## Rebrand the Custom ODBC Connector

All the TODOs in the QuickStart sample connector project are finished, and the custom connector is rebranded and retrieving data from your data store. To complete the custom connector, add the following functionality:

1. Rename all files and classes in the project to have the two-letter abbreviation chosen as part of **TODO #9**.
2. Create a connector configuration dialog. This dialog is presented to the user when they create a new ODBC DSN or configure an existing one. Note that the QuickStart connector project for Linux and UNIX platforms does not contain an example ODBC configuration dialog.

## Conclusion

You have written a custom ODBC connector that can be used by ODBC-enabled applications to query and retrieve data from a custom data store. The custom ODBC connector is renamed and rebranded for your company and product.

## Reference

This section contains more information that you may find useful when developing your sample ODBC driver.

### Driver Managers

Unlike Windows machines, most Linux, Unix, and macOS installations do not come with a driver manager as part of the operating system. You must install your own driver manager before you can compile and test your connector under Linux, Unix, or macOS. The following driver managers are supported by the Simba SDK:

Driver Manager	Download Location	Simba SDK Support
iODBC	<a href="http://www.iodbc.org">www.iodbc.org</a>	<ul style="list-style-type: none"> <li>Linux and Unix SDK</li> <li>MacOS SDK</li> </ul>
UnixODBC	<a href="http://www.unixodbc.org">www.unixodbc.org</a>	<ul style="list-style-type: none"> <li>Linux and Unix SDK</li> </ul>

#### **Note:**

This document uses the iODBC driver manager as an example, because it is supported by the Simba SDK on all Linux, Unix, and macOS platforms. It also contains an ODBC test utility.

### How do I know where my driver manager is installed?

If you did not install the driver manager yourself, you can look for it in typical installation directories. The driver manager must be installed to a directory that is on the library path:

- `LD_LIBRARY_PATH` on most Linux platforms
- `SHLIB_PATH` on HP/UX
- `LIBPATH` on AIX
- `DYLD_LIBRARY_PATH` on macOS

The iODBC driver manager is often installed to `/usr/lib` or `/usr/local/lib`. If you do not know where your driver manager is installed, try searching those directories for libraries containing the name `libiodb`.

## Locating the Configuration Files

The driver manager loads configuration information from the first `odbc.ini` and `odbcinst.ini` files that it finds. The Simba SDK loads connector-specific configuration information from the first `.simba.QuickStart.ini` file that it finds. Having multiple copies of these configuration files is not uncommon on development or customer machines, but can lead to confusion when a connector is not loaded as expected.

### Important:

Though we recommend only having one copy of each of the configuration files on your or your customer's machine, multiple versions may exist in different locations. The driver manager and the Simba SDK will load the first version of the file that they find.

## Locating the odbc.ini File

The driver manager looks for this file in the following locations, in the order specified:

1. If the **ODBCINI** environment variable is set, the driver manager looks in the directory specified by this variable. The environment variable is different for each driver manager, as shown in [Using Environment Variables to Specify the Location of the Configuration Files](#).
2. If there is no variable set, or if no file is found in that location, the driver manager looks in the user's home directory, `~/`.

### Note:

In this directory the file must have a preceding dot, for example:  
`~/odbc.ini`

3. The driver manager looks in the `/etc` directory, for example `/etc/odbc.ini`.

### Note:

There is no dot in front of the `odbc.ini` file in the `/etc` directory.

## Locating the odbcinst.ini File

The driver manager looks for this file in the following locations, in the order specified:

1. If the corresponding environment variable is set, the driver manager looks in the directory specified by this variable. The environment variable is different for each driver manager, as shown in [Using Environment Variables to Specify the Location of the Configuration Files](#).
2. If there is no variable set, or if no file is found in that location, the driver manager looks in the user's home directory.

**Note:**

In this directory the file must have a preceding dot, for example:  
`~/.odbcinst.ini`

3. The driver manager looks in the `/etc` directory, for example `/etc/odbcinst.ini`

**Note:**

There is no dot in front of the `odbcinst.ini` file in the `/etc` directory.

### Locating the `.simba.QuickStart.ini` configuration file

The Simba SDK searches for the `.simba.QuickStart.ini` file in the following locations, in the specified order:

1. The path, including the file name, specified using the **SIMBAINI** environment variable.

**Note:**

The name of this environment variable can be rebranded.

2. The connector directory, as a non-hidden `.ini` file.
3. The directory that the client application is launched from.
4. In `$HOME`, as a hidden `.ini` file
5. In `/etc/` as a non-hidden `.ini` file

### Using Environment Variables to Specify the Location of the Configuration Files

The location of configuration files is determined by environment variables, as shown below:

Environment Variable	File
ODBCINI	Specifies the full path (including the file name) of the <code>odbc.ini</code> file for the iODBC and unixODBC driver managers.
ODBCINSTINI	Specifies the full path (including the file name) of the <code>odbcinst.ini</code> file for the iODBC driver manager.
ODBCSYSINI	Specifies the directory containing the <code>odbcinst.ini</code> file for the unixODBC driver manager.
SIMBAINI	Specifies the full path (including the file name) of the connector configuration file. By default, this file is called <code>simba.drivervname.ini</code> .  <b>Note:</b> the name of this environment variable can be rebranded.

For example, these environment variables could be set as shown below:

**Example:**

```
export ODBCINI=/usr/local/odbc/myodbc.ini
export ODBCINSTINI=/usr/local/odbc/myodbcinst.ini
export ODBCSYSINI=/usr/local/odbc
```

If the environment variables are not set, the driver manager assumes that the configuration files exist in the user's home directory using the default file names `.odbc.ini` and `.odbcinst.ini`.

## [Configure the Connector and Data Source](#)

### Data Retrieval

In the Data Store Interface (DSI), the following methods perform the actual task of retrieving data from your data store:

- Each `MetadataSource` implementation of `GetMetadata()`
- `TabbedUnicodeFileReader::GetData()`

Both methods provide a way to uniquely identify a column within the current row. For `MetadataSource`, the Simba SQL Engine passes in a unique column tag (see



DSIOutputMetadataColumnTag). For QSTable, the Simba SQL Engine will pass in the column index .

In addition, both methods accept the following three parameters:

- `in_data`

The SQLData into which you must copy the value of your cell. This class is a wrapper around a buffer managed by the Simba SQL Engine. To access the buffer, you call its `GetBuffer()` method. The data you copy into the buffer must be formatted as a SQL Type (see <https://msdn.microsoft.com/en-us/library/ms710150%28VS.85%29.aspx> for a list of data types and definitions). Therefore, if your data is not stored as SQL Types, you will need to write code to convert from your native format.

The type of this parameter is governed by the metadata for the column that is returned by the class. Thus, if you set the SQL Type of column 1 in `QSTable::InitializeColumns()` to `SQL_INTEGER`, then when `TabbedUnicodeFileReader::GetData()` is called for column 1, you will be passed a SQLData that wraps an int data type. For `MetadataSource`, the type is associated with the column tag (see `DSIOutputMetadataColumnTag.h`).

#### Example:

```
If SqlData was of type SQL_INTEGER:
simba_int32 value = 5;
//This is one way
memcpy(in_data->GetBuffer(), &value, sizeof(simba_int32));
// This is another way; both work equally well
*reinterpret_cast<simba_int32*>(in_data->GetBuffer()) = 5;
```

When working with variable length data, for example character or binary data, you must call `SetLength()` before calling `GetBuffer()`. Not doing so may result in a heap violation. See `QSTypeUtilities.h` for an example on how to handle character or binary data.

- `in_offset`

Character, wide character and binary data types can be retrieved in parts. This value specifies where, in the current column, the value should be copied from. The value is usually 0.

- `in_maxSize`

The maximum size (in bytes) that can be copied into the `in_data` parameter.

For character or binary data, copying data that is greater than this size can result in a data truncation warning or a heap violation.

## SqlData Types

`SqlData` objects represent the SQL types and encapsulate the data in a buffer. To get the underlying SQL type that a `SqlData` object represents, use `GetMetadata()` - `>GetSqlType()`. This retrieves the associated `SQL_*` type.

For information on how SQL types map to C++ types, see [SQL Data Types in Developing Connectors for Data Stores Without SQL](#)

## Data Conversion in Practice

In the QuickStart example, when `GetData()` is called the values are read from the tabbed Unicode file (in `TabbedUnicodeFileReader::GetData`), converted to `simba_wstrings` (in `QSTable::ReadWholeColumnAsString`) and then converted to the requested SQL data type (in `QSTable::ConvertData`). This works well because the data source is a text file and a good cross-platform example.

For your data source, if you're already getting data of the correct type—integers, for example—then ideally you should skip the conversion to strings so you can achieve better performance. Be aware of which data types map to which SQL Types, as well as how to represent them in the expected format. Then you can set the buffer in an appropriate manner.

## NULL Values

To represent a null value, directly set the `SqlData` object as null:

```
in_data->SetNull(true);
```

## Server Configuration

Your custom ODBC connector can be recompiled as a server and deployed in a client-server configuration. The connection settings for the connector are normally retrieved directly from the ODBC DSN. However, when the connector is a server, the settings cannot be retrieved directly because the DSN refers to the client instead of a specific connector. Also, to enforce security, clients do not have control over server-specific settings.

For information about making a connection to a connector that is compiled and built as a server, see the [SimbaClient/Server Developer Guide](#).

## How to Add Schema Support

Some applications require support for schemas (note that Microsoft Excel does not require schema support). If the data store supports schemas, the connector can provide access to them. To support schemas, a connector must handle `DSI_SCHEMAONLY_METADATA`.

### To add schema support to a connector:

1. By default, `QSqlConnection::SetConnectionPropertyValues()` disables schema support via `DSIPropertyUtilities::SetSchemaSupport()`. Change this value to `true` to enable schema support.
2. Make the following additional changes:
  - a. `QSMetadataHelper::GetNextTable()`  
In the QuickStart sample connector, a blank schema is returned because schema support is not enabled by default. To enable schema support, return the schema in the `Identifier` to allow Simba SDK to open the correct table.
  - b. `QSDatabase::OpenTable()`  
Modify this method to verify the given schema and return the correct table for the given catalog, schema and table name.
  - c. `QSTable::GetSchemaName()`  
Modify this method to return the schema to which the table belongs

## Install the Evaluation License

You can use Simba SDK for 30 days after installing the evaluation license. The evaluation license is emailed to the person who registered the product.

Typically, you use Simba SDK to create your custom ODBC connector, then use a test ODBC-enabled application to retrieve data using the connector. you are running.

### Install the Evaluation License on Unix, Linux, and macOS

#### To license the Simba SDK:

- Save the license file under the `$HOME` directory, either as a hidden or a non-hidden file. For example:

```
/home/<user_id>/SimbaEngineSDK.lic for a non-hidden file
```

```
/home/<user_id>/.SimbaEngineSDK.lic for a hidden file
```

## Troubleshooting

This section contains solutions to common problems.

### Specified Driver Could Not Be Loaded

On Linux, Unix, and macOS platforms, the error `Specified Driver Could Not Be Loaded` may be returned when you try to connect to a data source. Here are some of the reasons the driver manager cannot load the specified connector:

- You are using a 64-bit driver manager, but have 32-bit paths defined in the `odbc.ini` or `odbsinst.ini` files
- You define the path to the debug version of the connector, but you built the non-debug version

### Unable to Connect to the Connector

In the `odbc.ini` and `odbcinst.ini` files, be sure there are no characters such as spaces or tabs at the end of a value. This can cause failure to connect.

### Is My Driver Manager 32-or-64 bit?

You can use the `file` command.

#### Example:

This version of `iodbctest` (and therefore the iODBC driver manager) is 64-bit:

```
/usr/local/bin> file iodbctest
iodbctest: ELF 64-bit LSB executable, x86-64, version .....
```

*Testing and Troubleshooting* in the guide [Developing Connectors for Data Stores Without SQL](#)

## Contact Us

For more information or help using this product, please contact our Technical Support staff. We welcome your questions, comments, and feature requests.

**Note:**

To help us assist you, prior to contacting Technical Support please prepare a detailed summary of the Simba SDK version and development platform that you are using.

You can contact Technical Support via the Magnitude Support Community at [www.magnitude.com](http://www.magnitude.com).

You can also follow us on Twitter @SimbaTech and @Mag\_SW.

## Third-Party Trademarks

Simba, the Simba logo, Simba SDK, 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.

Kerberos is a trademark of the Massachusetts Institute of Technology (MIT).

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

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

Microsoft SQL Server, SQL Server, Microsoft, MSDN, Windows, Windows Azure, 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.

Solaris is a registered trademark of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

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

Ubuntu is a trademark or registered trademark of Canonical Ltd. or its subsidiaries in Canada, United States and/or other countries.

All other trademarks are trademarks of their respective owners.