

InfoSphere Information Server DataStage

Configure ODBC for DataStage on UNIX or Linux



© 2012 IBM Corporation

This presentation will discuss how to configure ODBC connectivity for DataStage® on UNIX or Linux.

Objectives

- ODBC drivers
- Configure .odbc.ini file
- Specifying data source name in uvodbc.config file
- Test ODBC driver connectivity
- Troubleshooting

This presentation will discuss the various components involved in configuring ODBC connectivity to a data source in DataStage. It discusses the ODBC drivers used by DataStage and how to configure the .odbc.ini file to create a data source name for a database. It also will discuss how to specify the data source name in the uvodbc.config file for a specific DataStage project and how to test the ODBC driver connectivity and troubleshooting ODBC issues.

Three key components involved to configure ODBC connectivity in DataStage are ODBC drivers, the .odbc.ini file and the uvodbc.config file.

ODBC drivers

- DataStage product ships with Data Direct Drivers
- Installed under `/opt/IBM/InformationServer/Server/branded_odbc`

The first component in configuring an ODBC connection is the ODBC drivers. DataStage ships with several database specific drivers from Data Direct Technologies. The data direct drivers are by default installed under the directory `/opt/IBM/InformationServer/Server/branded_odbc`.

Configuring .odbc.ini file (1 of 2)

- Located by default under \$DSHOME
- Defines all ODBC data sources within DataStage
- Contains pre-configured examples for drivers shipped with DataStage

```
[Oracle Wire Protocol]
Driver=/u2/IS810/IBM/InformationServer/Server/branded_odbc/lib/VMora24.so
Description=DataDirect Oracle Wire Protocol driver
ApplicationUsingThreads=1
ArraySize=60000
CachedCursorLimit=32
CachedDescLimit=0
CatalogIncludesSynonyms=1
CatalogOptions=0
DefaultLongDataBuffLen=1024
DescribeAtPrepare=0
EnableDescribeParam=0
EnableNcharSupport=0
EnableScrollableCursors=1
EnableStaticCursorsForLongData=0
EnableTimestampWithTimeZone=0
HostName=Oracle server
LocalTimeZoneOffset=
LockTimeOut=-1
LogonID=
Password=
PortNumber=Oracle server port
ProcedureRetResults=0
SID=Oracle SID
UseCurrentSchema=1
```

The second component is the .odbc.ini file which defines all the Data Sources within DataStage. The path of the .odbc.ini file used by DataStage is defined by the variable ODBCINI in the dsenv file and by default is located under the DataStage engine directory. The data source within the .odbc.ini file defines the data direct driver to be used and provides information needed by the data source to connect to the database.

The .odbc.ini file contains several pre-configured examples for the various database drivers. A new data source is configured by creating a copy of the pre-configured example and editing the copy according to the database specifications. This example will use the pre-configured Oracle wire protocol entry displayed on this slide.

Configuring .odbc.ini file (2 of 2)

- Create copy of sample DSN
 - Edit attributes
- Add entry to [ODBC Data Sources] section
- ODBC Driver documentation
 - InformationServer/Server/branded_odbc/docs

```
[ODBC Data Sources]
OracleTest=Oracle Test Database

[OracleTest]
Driver=/n2/YS910/IBM/InformationServer/Server/branded_odbc/lib/VMora24.so
Description=DataDirect Oracle Wire Protocol driver
ApplicationUsingThreads=1
ArraySize=60000
CachedCursorLimit=32
CachedDescLimit=0
CatalogIncludesSynonyms=1
CatalogOptions=0
DefaultLongDataBuffLen=1024
DescribeAtPrepare=0
EnableDescribeParam=0
EnableNcharSupport=0
EnableScrollableCursors=1
EnableStaticCursorsForLongData=0
EnableTimestampWithTimeZone=0
HostName=134.168.53.16
LocalTimeZoneOffset=
LockTimeOut=-1
PortNumber=1521
ProcedureRetResults=0
SID=DNSH1
UseCurrentSchema=1
```

5 Configure ODBC for DataStage on UNIX or Linux © 2012 IBM Corporation

The example displayed on this slide shows the details of the data source named OracleTest that was created after copying and editing the pre-configured Oracle Wire protocol displayed in the previous slide to connect to the Oracle database DNSH1. It is also necessary to make an entry at the top of the .odbc.ini file under the section ODBC Data Sources with the DSN Name and a description. Some components of DataStage look for the entry in the [ODBC Data Sources] section.

Additional documents that provide more information on the data source attributes for various databases are located under InformationServer/Server/branded_odbc/docs.

Specifying data source names in uvodbc.config file

- Located in \$DSHOME and each project
 - uvodbc.config in project is accessed first
 - Uses uvodbc.config in DSEngine if uvodbc.config does not exist in project
- Syntax for uvodbc.config

```
<DSN_NAME>
DBMSTYPE = ODBC
```

 - Where DSN_NAME is data source name specified in .odbc.ini file
 - Case must match .odbc.ini file
- Space on each side of equal sign required

```
[ODBC DATA SOURCES]
<localuv>
DBMSTYPE = UNIVERSE
network = TCP/IP
service = uvserver
host = 127.0.0.1

<SAMPLEDB>
DBMSTYPE = ODBC

<OracleTest>
DBMSTYPE = ODBC
```

6

Configure ODBC for DataStage on UNIX or Linux

© 2012 IBM Corporation

The third and final component is the uvodbc.config file. This file specifies the ODBC data source name (DSN) for each database. The uvodbc.config file exists at two locations. A common file exists at DSEngine directory and project specific file exists under each DataStage projects directory. The copy in DSEngine will only be accessed if the uvodbc.config file does not exist in the project directory.

Add the data source name to the uvodbc.config file in your project directory using the syntax displayed on this slide. The Data Source Name must match the Data Source Name in the .odbc.ini file exactly, including the case. Also, be sure that every equal sign in this file has a space on both sides.

Testing ODBC driver connectivity

- Connection can be tested using DS_CONNECT command

1. Setup the environment by sourcing the dsenv file

```
./dsenv
```

2. Start the DataStage server Shell

```
./bin/dssh
```

3. Log into the DataStage Project

```
LOGTO project_name
```

4. Test the connection

```
DS_CONNECT dsn
```

5. Enter the user name and password when prompted

Once the DSN has been configured, the next step is to test the ODBC connectivity to the data source. This can be done using the UV command DS_CONNECT. Before testing the connectivity, change directories to DSEngine directory and setup the environment by sourcing the dsenv file. Start a DataStage server shell by running the bin/dssh command from the DSEngine directory. Log to the specific DataStage project using the LOGTO command followed by the DataStage project name. Test the connection to the data source by executing the command DS_CONNECT followed by DSN name. When prompted, enter the user name and password to connect to the specific database.

When testing is complete, enter .Q to quit the connection.

Troubleshooting - dctestlib

- dctestlib can be used to test data direct ODBC driver
- Help diagnose configuration problems
- Attempts to load specified ODBC driver and provides error information if load fails
- Located under InformationServer/Server/branded_odbc/bin
- Source dsenv file before running dctestlib
- Syntax to run dctestlib

```
dctestlib <libname_with_fullpath>
```

If your ODBC connection is not working properly, the first thing to check is if the ODBC driver library will load properly. dctestlib can be used to test the library. Located under the InformationServer/Server/branded_odbc/bin directory, dctestlib attempts to load the specified driver and provides error information if the load fails. Source the dsenv file before running dctestlib to ensure that your test environment matches the DataStage environment. If the library fails to load, check your environment variables to be sure they are set correctly and that you have the proper database clients loaded if you are using non-wired ODBC drivers.

Troubleshooting - ODBC tracing

- Set Trace=1 under [ODBC] DSN
- Set TraceFile to an absolute path
- All ODBC calls for all processes go into same trace file
 - Trace with no one else using odbc
- Set Trace=0 as soon as test is complete
- For More information on ODBC tracing, see IBM Education Assistant module “Problem Determination – ODBC Tracing”

http://publib.boulder.ibm.com/infocenter/ieduasst/inv1r0/index.jsp?topic=/com.ibm.iea.datastage/datastage/8.1/ProblemDetermination/DS_odbcTracing/player.html

```
[ODBC]
UNICODE=UTF-8
IANAAppCodePage=4
InstallDir=/u2/IS810/IBM/InformationServer/Server/branded_odbc
Trace=1
TraceDll=/u2/IS810/IBM/InformationServer/Server/branded_odbc/lib/odbc trac.so
TraceFile=/tmp/odbc_trace.out
UseCursorLib=0
```

If the library loads but there are still connection problems, the next step is to do an ODBC trace. The .odbc.ini file also contains a data source entry [ODBC] at the end of the file. ODBC tracing can be turned ON by setting Trace=1. You should change the TraceFile attribute to the absolute path to your trace file to make it easy to find. Be sure the directory where the trace file is going to be written to has plenty of room as those files can become large. All the trace records for all processes go to one file so it is best to trace with no one else using any other ODBC data sources. Be sure to turn the Trace attribute back to 0 as soon as your trace is completed. For additional information on ODBC tracing, see the IBM Education Assistant module “Problem Determination – ODBC Tracing” using the URL displayed on this slide.

Troubleshooting - QEWS

- DataDirect ODBC drivers should not be used outside of DataStage
- Creates QEWS entries and places drivers in demo mode
- QEWS entries can also result in Invalid license file warnings
- For additional information, see technote links
<https://www.ibm.com/support/docview.wss?uid=swg21454584>
<https://www.ibm.com/support/docview.wss?uid=swg21406961>

It is important to note that the DataDirect ODBC drivers shipped with Information Server product should not be used outside of DataStage. If the drivers are used outside of DataStage, it shows a warning message and creates a QEWS entry at the top of the data source entry in the .odbc.ini file. If the QEWS entries exist, the drivers will be in demo mode and can potentially expire. The QEWS entries can also result in Invalid license file warnings when running DataStage jobs. Remove the QEWS entries and ensure the drivers are not used outside of DataStage.

Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, DataStage, and InfoSphere are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. Other company, product, or service names may be trademarks or service marks of others.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2012. All rights reserved.