MultiSource Virtualization Plugin

Overview

The multiple source query plugin allows SQuirreL users to create a virtual data source that may consist of multiple data sources on different servers and platforms. The user can enter one SQL query to combine and join information from multiple sources. Any database that has a JDBC driver is supported including NoSQL databases such as MongoDB.

Benefits

- The plugin allows SQuirreL to natively support multiple source queries.
- No data source or server changes are required.
- The plugin supports standard SQL including joins, group by, aggregation, LIMIT, and ordering where tables may come from one or more sources.
- The plugin will perform function translation when a user requests a function or SQL feature/syntax that is not supported by a certain source.
- Users can export the virtualization information and use the virtualization driver in other Java programs and reporting software.

Installation

1. Download and install SQuirreL. Add the MultiSource plugin during installation or add to a current version by unzipping `multisource.zip` in the `plugins` folder.

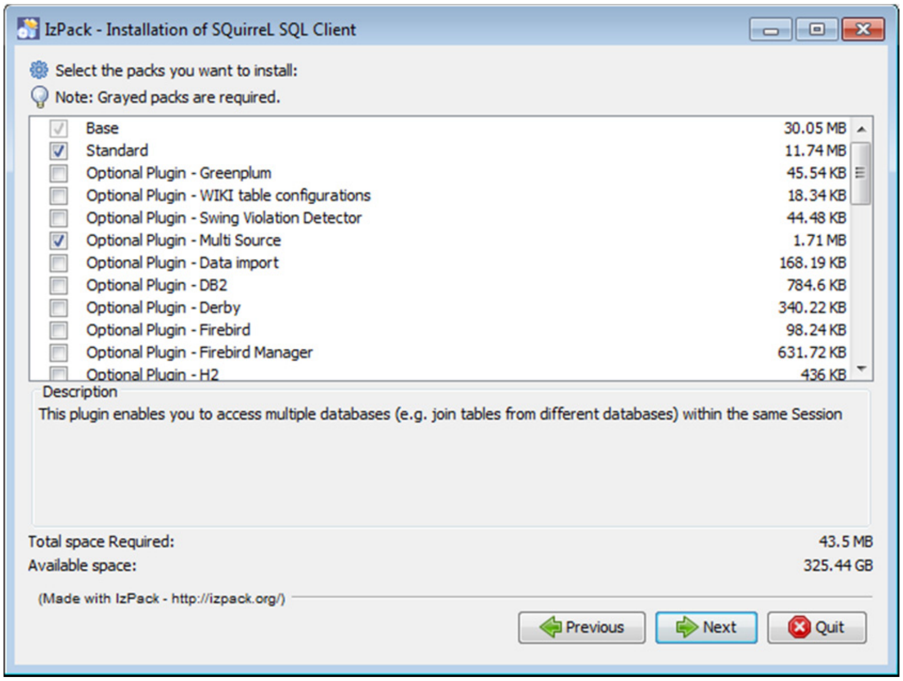

IzPack - Installation of SQuirreL SQL Client

Select the packs you want to install:

Note: Grayed packs are required.

☑	Base	30.05 MB
☑	Standard	11.74 MB
□	Optional Plugin - Greenplum	45.54 KB
□	Optional Plugin - WIKI table configurations	18.34 KB
□	Optional Plugin - Swing Violation Detector	44.48 KB
☑	Optional Plugin - Multi Source	1.71 MB
□	Optional Plugin - Data import	168.19 KB
□	Optional Plugin - DB2	784.6 KB
□	Optional Plugin - Derby	340.22 KB
□	Optional Plugin - Firebird	98.24 KB
□	Optional Plugin - Firebird Manager	631.72 KB
□	Optional Plugin - H2	436 KB

Description

This plugin enables you to access multiple databases (e.g. join tables from different databases) within the same Session

Total space Required: 43.5 MB

Available space: 325.44 GB

(Made with IzPack - http://izpack.org/)

2. The plugin contains `unityjdbc.jar` in folder `squirrel/plugins/multisource`. Copy this jar into the `squirrel/lib` folder. **Make sure to add other database JDBC jars into the `squirrel/lib` folder or JRE classpath. The plugin cannot access drivers in custom classpaths used by SQuirreL.**

UnityJDBC JAR Location

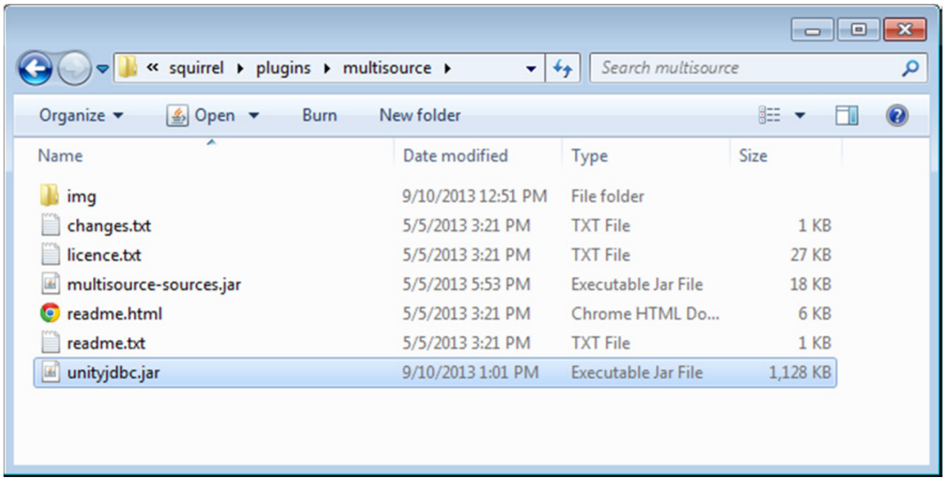

Name	Date modified	Type	Size
img	9/10/2013 12:51 PM	File folder	
changes.txt	5/5/2013 3:21 PM	TXT File	1 KB
licence.txt	5/5/2013 3:21 PM	TXT File	27 KB
multisource-sources.jar	5/5/2013 5:53 PM	Executable Jar File	18 KB
readme.html	5/5/2013 3:21 PM	Chrome HTML Do...	6 KB
readme.txt	5/5/2013 3:21 PM	TXT File	1 KB
unityjdbc.jar	9/10/2013 1:01 PM	Executable Jar File	1,128 KB

Put UnityJDBC JAR in `squirrel/lib` folder

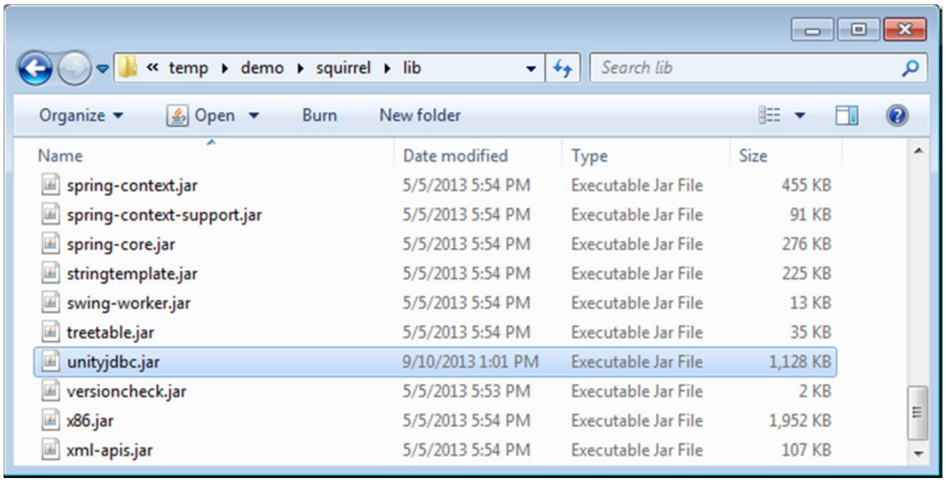

Name	Date modified	Type	Size
spring-context.jar	5/5/2013 5:54 PM	Executable Jar File	455 KB
spring-context-support.jar	5/5/2013 5:54 PM	Executable Jar File	91 KB
spring-core.jar	5/5/2013 5:54 PM	Executable Jar File	276 KB
stringtemplate.jar	5/5/2013 5:54 PM	Executable Jar File	225 KB
swing-worker.jar	5/5/2013 5:54 PM	Executable Jar File	13 KB
treetable.jar	5/5/2013 5:54 PM	Executable Jar File	35 KB
unityjdbc.jar	9/10/2013 1:01 PM	Executable Jar File	1,128 KB
versioncheck.jar	5/5/2013 5:53 PM	Executable Jar File	2 KB
x86.jar	5/5/2013 5:54 PM	Executable Jar File	1,952 KB
xml-apis.jar	5/5/2013 5:54 PM	Executable Jar File	107 KB

3. Start SQuirreL. The multisource plugin should be visible in the plugin list.

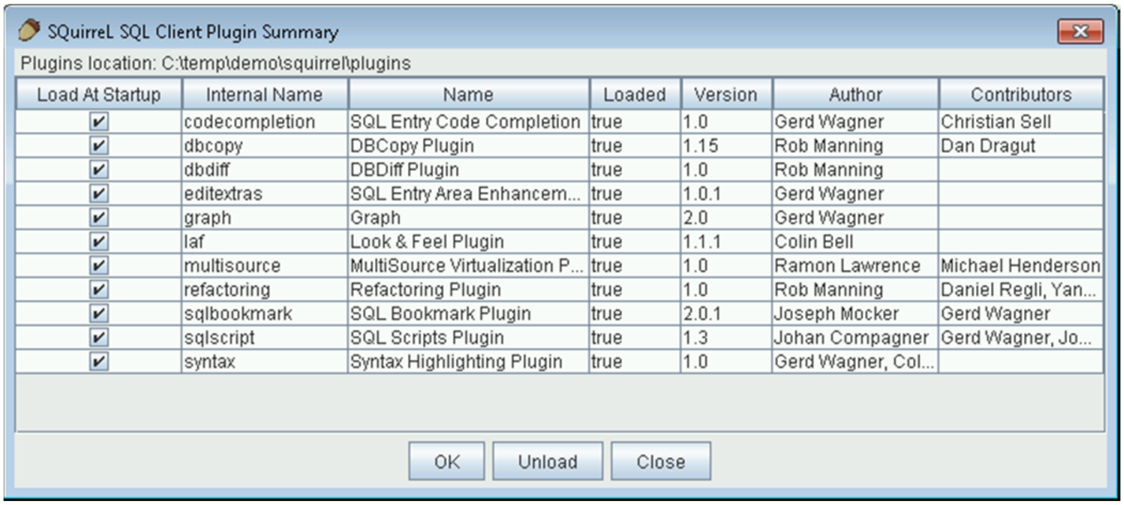

SQuirreL SQL Client Plugin Summary

Plugins location: C:\temp\demo\squirrel\plugins

Load At Startup	Internal Name	Name	Loaded	Version	Author	Contributors
☑	codecompletion	SQL Entry Code Completion	true	1.0	Gerd Wagner	Christian Sell
☑	dbcopy	DBCopy Plugin	true	1.15	Rob Manning	Dan Dragut
☑	dbdiff	DBDiff Plugin	true	1.0	Rob Manning	
☑	editextras	SQL Entry Area Enhancem...	true	1.0.1	Gerd Wagner	
☑	graph	Graph	true	2.0	Gerd Wagner	
☑	laf	Look & Feel Plugin	true	1.1.1	Colin Bell	
☑	multisource	MultiSource Virtualization P...	true	1.0	Ramon Lawrence	Michael Henderson
☑	refactoring	Refactoring Plugin	true	1.0	Rob Manning	Daniel Regli, Yan...
☑	sqlbookmark	SQL Bookmark Plugin	true	2.0.1	Joseph Mocker	Gerd Wagner
☑	sqlscript	SQL Scripts Plugin	true	1.3	Johan Compagner	Gerd Wagner, Jo...
☑	syntax	Syntax Highlighting Plugin	true	1.0	Gerd Wagner, Col...	

How It Works

1. Register data source aliases in SQuirreL as usual. This example has connections to a Microsoft SQL Server database, a MySQL database, an Oracle database, a PostgreSQL database, and a MongoDB database all containing the TPC-H benchmark data. Note that any database with a JDBC driver is supported including those accessible using the JDBC-ODBC bridge. For MongoDB support, the latest MongoDB Java driver should also be in the `squirrel/lib` folder.

Microsoft SQL Server

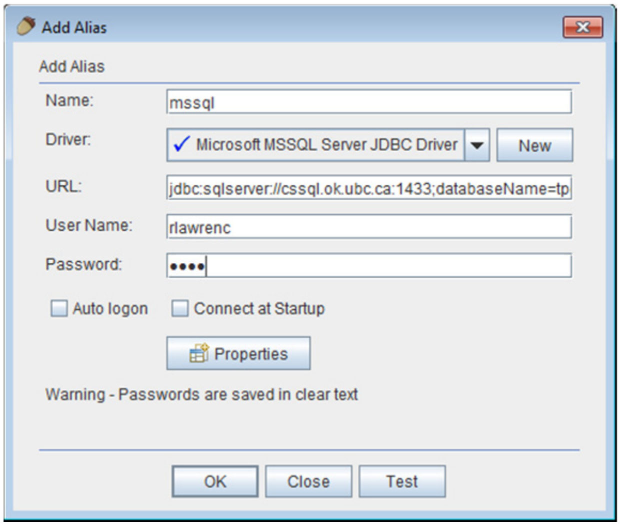

Name: mssql
Driver: ✓ Microsoft MSSQL Server JDBC Driver
URL: jdbc:sqlserver://cssql.ok.ubc.ca:1433;databaseName=tp
User Name: rlawrenc
Password: ••••
□ Auto logon □ Connect at Startup
Warning - Passwords are saved in clear text

MySQL

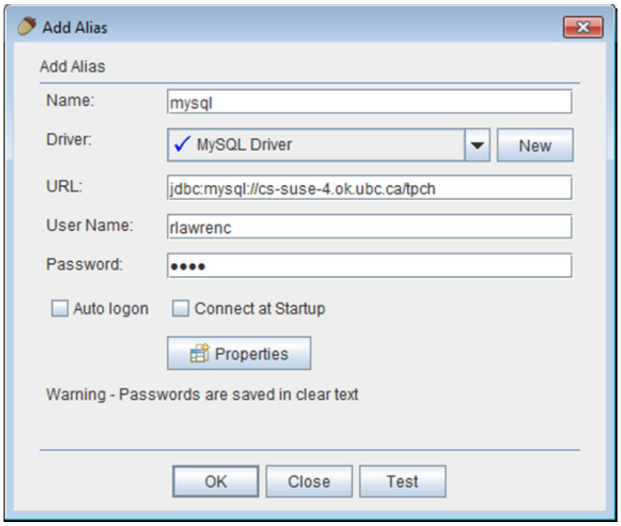

Name: mysql
Driver: ✓ MySQL Driver
URL: jdbc:mysql://cs-suse-4.ok.ubc.ca/tpch
User Name: rlawrenc
Password: ••••
□ Auto logon □ Connect at Startup
Warning - Passwords are saved in clear text

Oracle

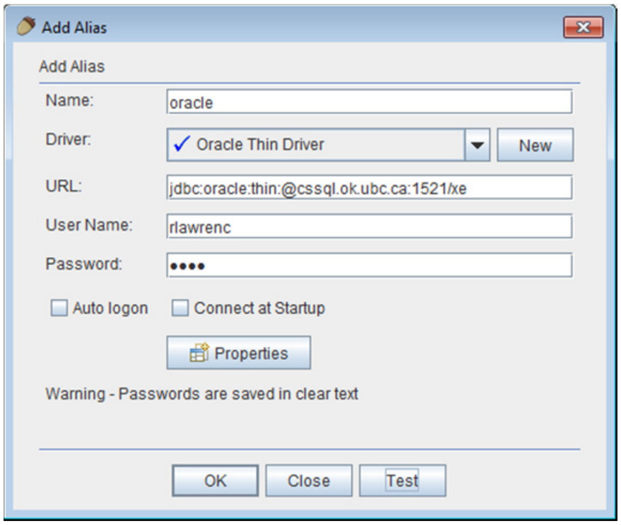

Name: oracle
Driver: ✓ Oracle Thin Driver
URL: jdbc:oracle:thin:@cssql.ok.ubc.ca:1521/xe
User Name: rlawrenc
Password: ••••
□ Auto logon □ Connect at Startup
Warning - Passwords are saved in clear text

Postgres

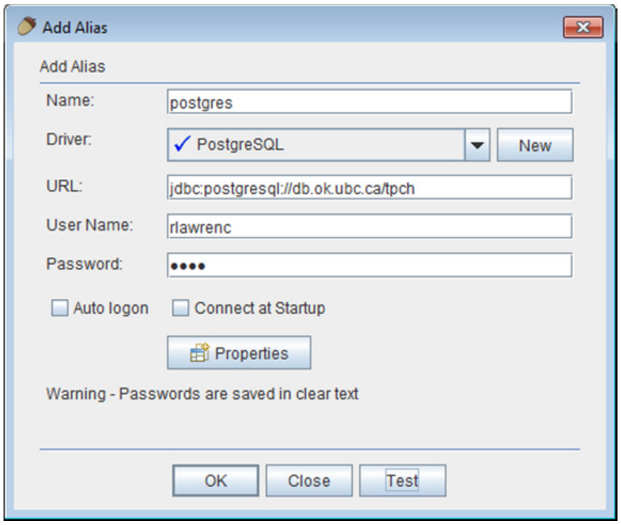

Name: postgres
Driver: ✓ PostgreSQL
URL: jdbc:postgresql://db.ok.ubc.ca/tpch
User Name: rlawrenc
Password: ••••
□ Auto logon □ Connect at Startup
Warning - Passwords are saved in clear text

Register MongoDB Driver

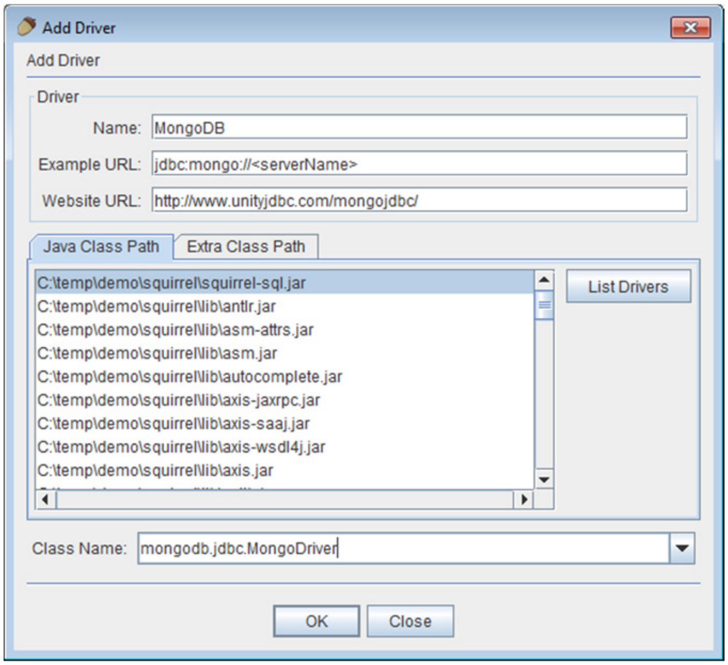

Add Driver

Driver

Name: MongoDB

Example URL: jdbc:mongo://<serverName>

Website URL: http://www.unityjdbc.com/mongojdbc/

```
C:\temp\demo\squirrel\squirrel-sql.jar
C:\temp\demo\squirrel\lib\antlr.jar
C:\temp\demo\squirrel\lib\asm-attrs.jar
C:\temp\demo\squirrel\lib\asm.jar
C:\temp\demo\squirrel\lib\autocomplete.jar
C:\temp\demo\squirrel\lib\axis-jaxrpc.jar
C:\temp\demo\squirrel\lib\axis-saaj.jar
C:\temp\demo\squirrel\lib\axis-wsdl4j.jar
C:\temp\demo\squirrel\lib\axis.jar
```

Class Name: mongodb.jdbc.MongoDriver

MongoDB

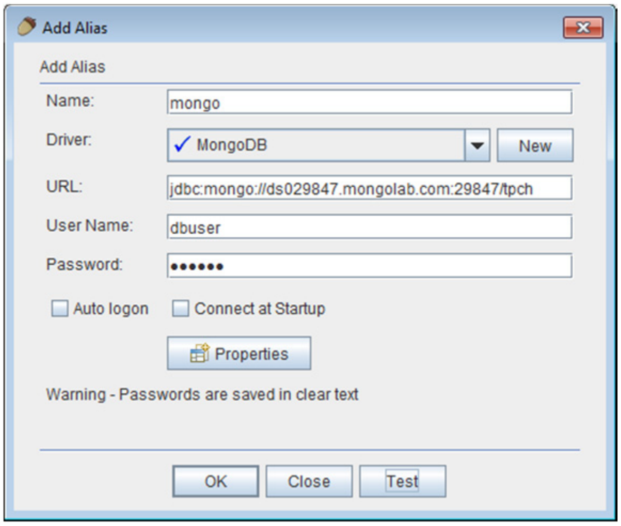

Add Alias

Name: mongo

Driver: ✓ MongoDB

URL: jdbc:mongo://ds029847.mongolab.com:29847/tpch

User Name: dbuser

Password: ••••••

☐ Auto logon ☐ Connect at Startup

Warning - Passwords are saved in clear text

2. Make sure you have registered the UnityJDBC MultiSource Virtualization driver (during installation). Create an alias consisting of virtual sources. The name field can be any name. It does not have to contain "virtual".

Create Virtual Source

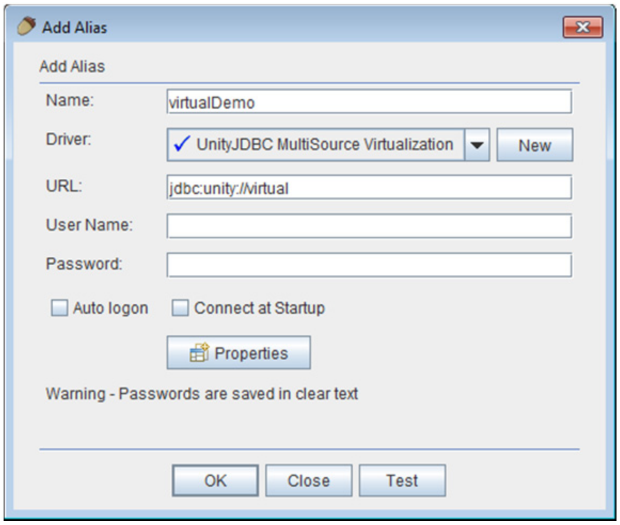

Add Alias

Name: virtualDemo

Driver: ✓ UnityJDBC MultiSource Virtualization

URL: jdbc:unity://virtual

User Name:

Password:

☐ Auto logon ☐ Connect at Startup

Warning - Passwords are saved in clear text

Connect to Virtual Source

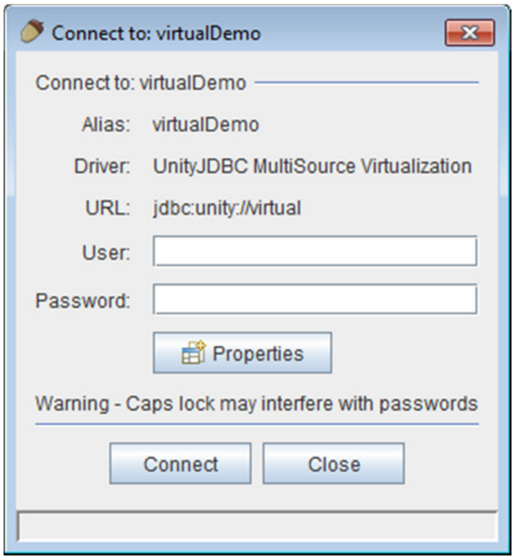

Connect to: virtualDemo

Alias: virtualDemo

Driver: UnityJDBC MultiSource Virtualization

URL: jdbc:unity://virtual

User:

Password:

Warning - Caps lock may interfere with passwords

3. Add existing JDBC connections (SQuirreL aliases) to the virtual source. The example adds all five of the sources created above. On the source name (in this case `virtualDemo`), right click and select `(Virtualization) Add source`.

Before add Microsoft SQL (mssql) source

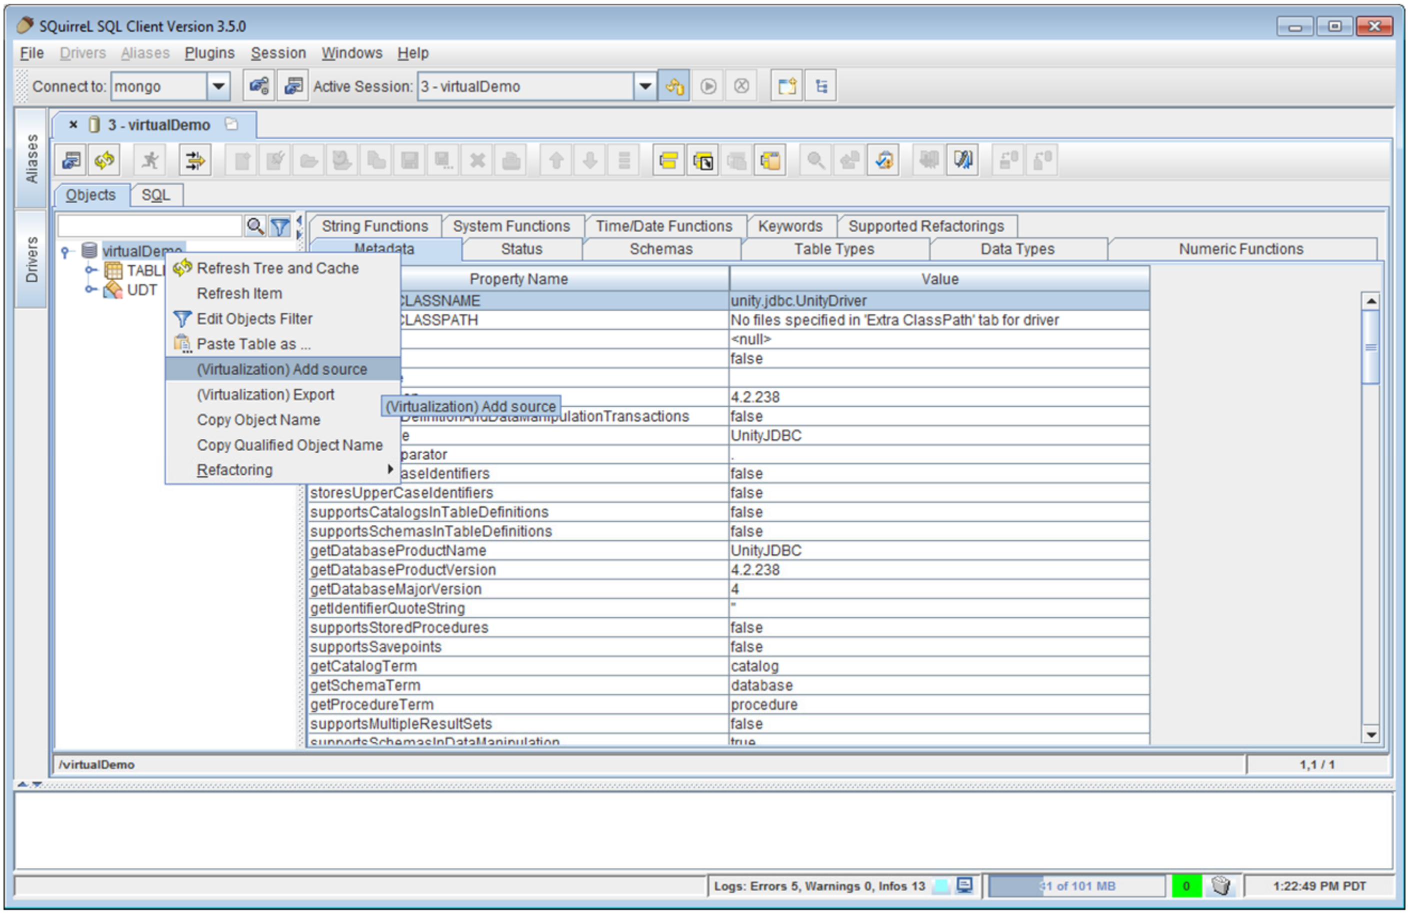

Property Name	Value
CLASSNAME	unity.jdbc.UnityDriver
CLASSPATH	No files specified in 'Extra ClassPath' tab for driver
	<null>
	false
	4.2.238
DefinitionAndDataManipulationTransactions	false
e	UnityJDBC
parator	.
aseIdentifiers	false
storesUpperCaseIdentifiers	false
supportsCatalogsInTableDefinitions	false
supportsSchemasInTableDefinitions	false
getDatabaseProductName	UnityJDBC
getDatabaseProductVersion	4.2.238
getDatabaseMajorVersion	4
getIdentifierQuoteString	"
supportsStoredProcedures	false
supportsSavepoints	false
getCatalogTerm	catalog
getSchemaTerm	database
getProcedureTerm	procedure
supportsMultipleResultSets	false
supportsSchemasInDataManipulation	true

Select the source to add and click the Add button. Progress is shown.

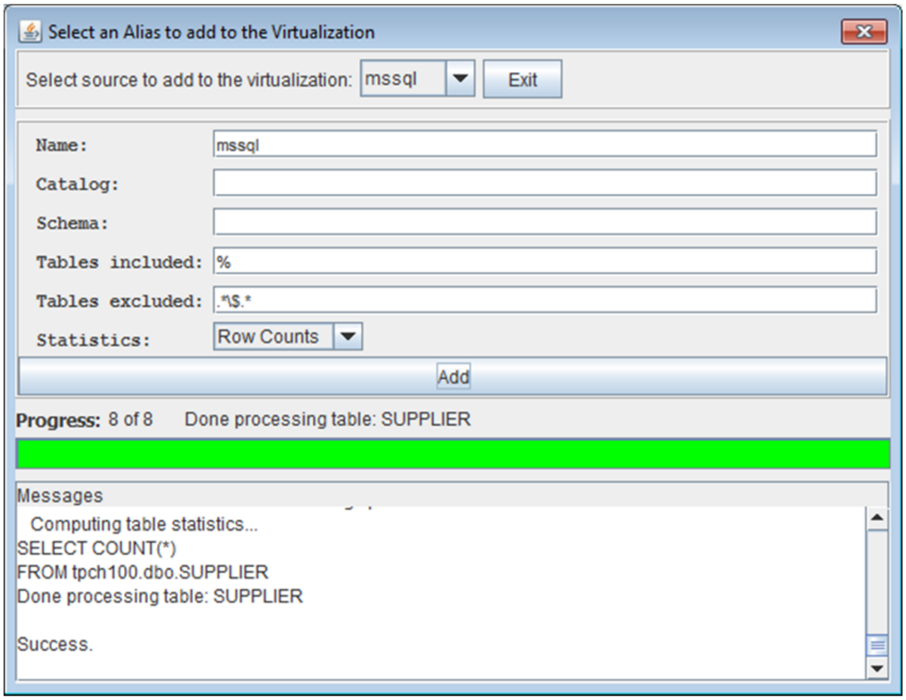

Select source to add to the virtualization: mssql

Name: mssql

Catalog:

Schema:

Tables included: %

Tables excluded: .*\$.*

Statistics: Row Counts

Progress: 8 of 8 Done processing table: SUPPLIER

```
  Computing table statistics...
SELECT COUNT(*)
FROM tpch100.dbo.SUPPLIER
Done processing table: SUPPLIER

Success.
```

After adding the Microsoft source. Tables are visible in the object tree view.

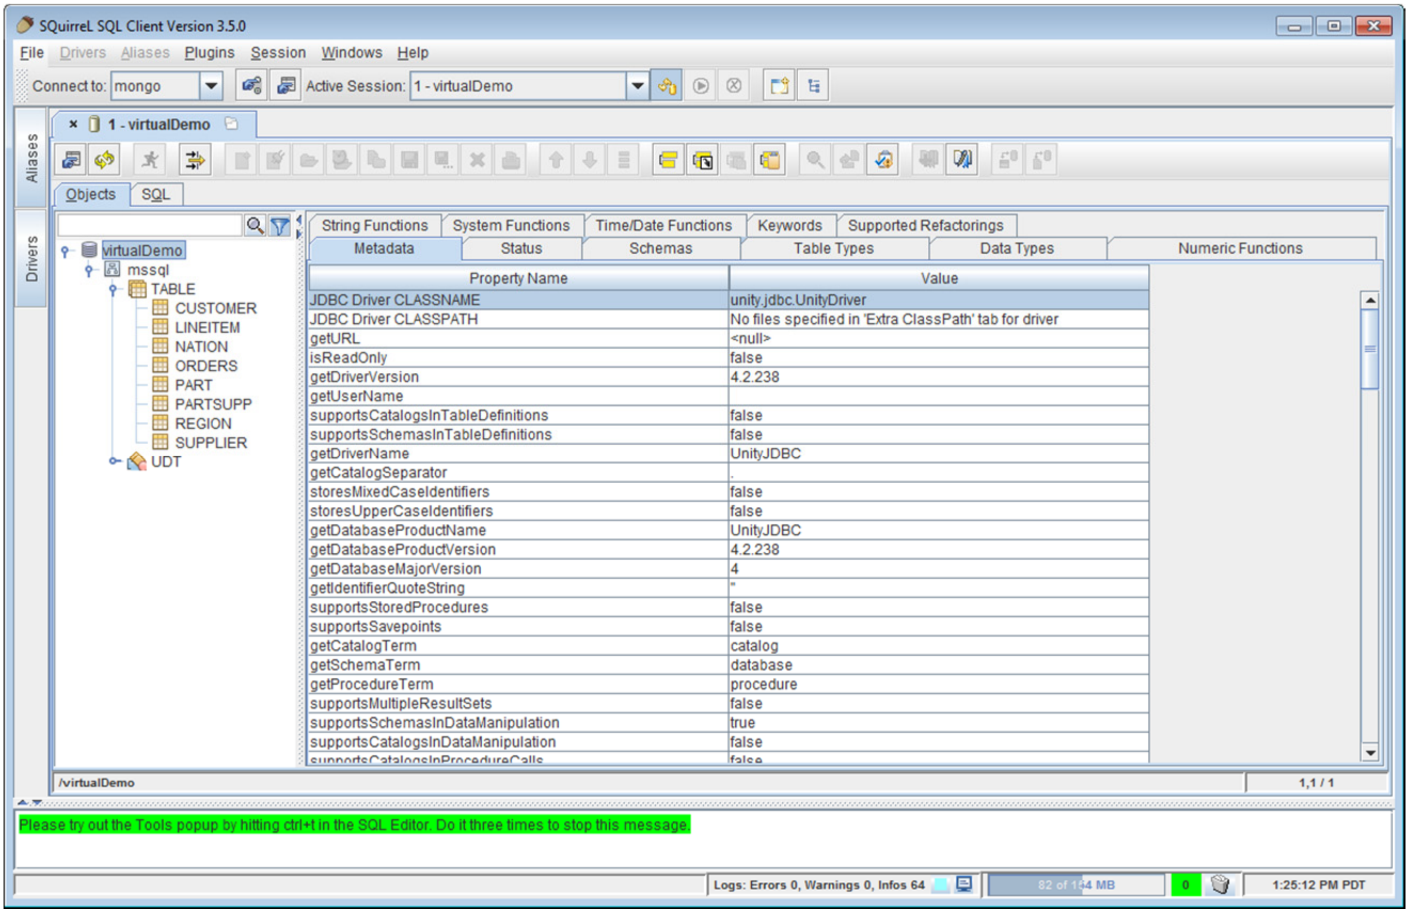

Property Name	Value
JDBC Driver CLASSNAME	unity.jdbc.UnityDriver
JDBC Driver CLASSPATH	No files specified in 'Extra ClassPath' tab for driver
getURL	<null>
isReadOnly	false
getDriverVersion	4.2.238
getUserName	
supportsCatalogsInTableDefinitions	false
supportsSchemasInTableDefinitions	false
getDriverName	UnityJDBC
getCatalogSeparator	.
storesMixedCaseIdentifiers	false
storesUpperCaseIdentifiers	false
getDatabaseProductName	UnityJDBC
getDatabaseProductVersion	4.2.238
getDatabaseMajorVersion	4
getIdentifierQuoteString	"
supportsStoredProcedures	false
supportsSavepoints	false
getCatalogTerm	catalog
getSchemaTerm	database
getProcedureTerm	procedure
supportsMultipleResultSets	false
supportsSchemasInDataManipulation	true
supportsCatalogsInDataManipulation	false
supportsCatalogsInProcedureCalls	false

/virtualDemo

1,1 / 1

Please try out the Tools popup by hitting ctrl+t in the SQL Editor. Do it three times to stop this message.

Logs: Errors 0, Warnings 0, Infos 64

82 of 164 MB

0

1:25:12 PM PDT

4. Users can add as many sources as they wish. You can also rename the source in the virtual view. It does not have to be the same as the alias name used by SQuirreL.
    - When adding Oracle sources, make sure to specify a schema so that system tables and tables from all schemas are not extracted.
    - You can also filter tables added by catalog name, schema name, and table filters.
    - The table filters to include are specified in SQL (JDBC) syntax with a % as a wild-card match. The table exclusion filters are specified as Java regular expressions.

Adding an Oracle Source with a Schema

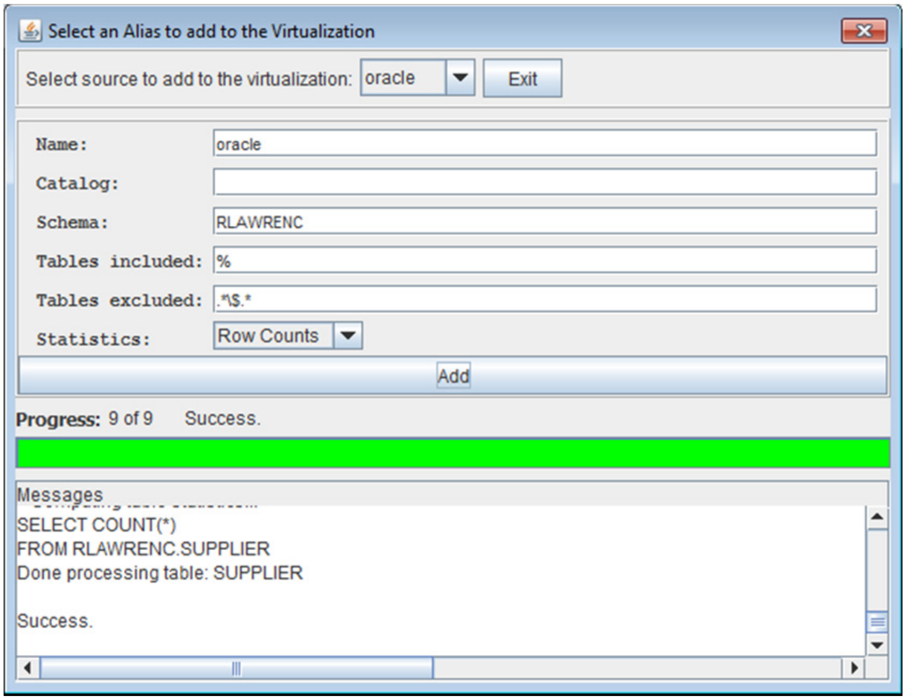

Select source to add to the virtualization: oracle

Name: oracle

Catalog:

Schema: RLAWRENC

Tables included: %

Tables excluded: .*\$.*

Statistics: Row Counts

Progress: 9 of 9 Success.

```
SELECT COUNT(*)
FROM RLAWRENC.SUPPLIER
Done processing table: SUPPLIER

Success.
```

Object Tree View with all Five Sources Added

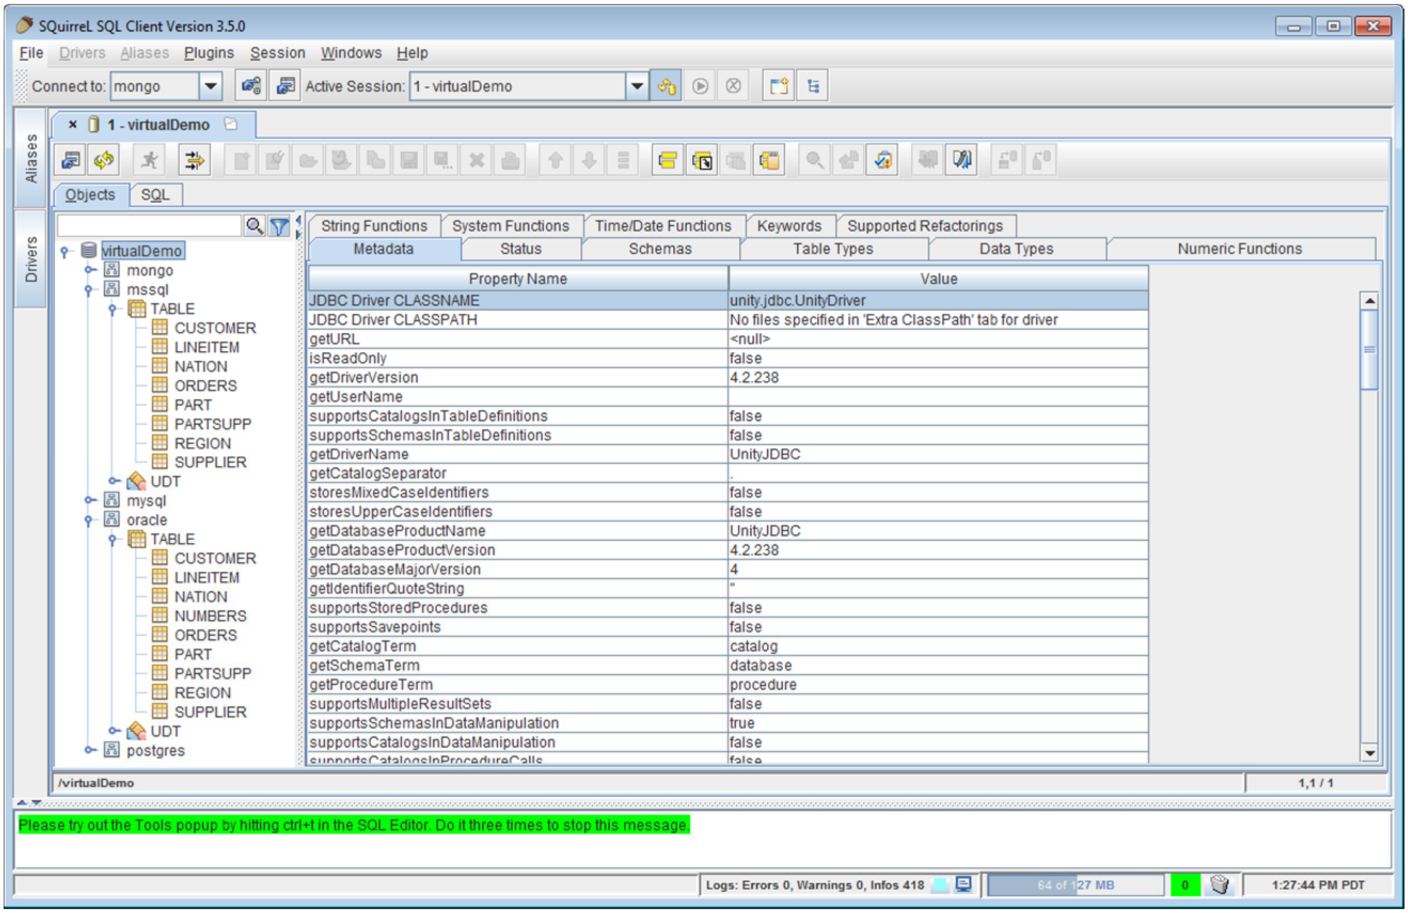

- virtualDemo
  - mongo
  - mssql
    - TABLE
      - CUSTOMER
      - LINEITEM
      - NATION
      - ORDERS
      - PART
      - PARTSUPP
      - REGION
      - SUPPLIER
    - UDT
  - mysql
  - oracle
    - TABLE
      - CUSTOMER
      - LINEITEM
      - NATION
      - NUMBERS
      - ORDERS
      - PART
      - PARTSUPP
      - REGION
      - SUPPLIER
    - UDT
  - postgres

Property Name	Value
JDBC Driver CLASSNAME	unity.jdbc.UnityDriver
JDBC Driver CLASSPATH	No files specified in 'Extra ClassPath' tab for driver
getURL	<null>
isReadOnly	false
getDriverVersion	4.2.238
getUserName	
supportsCatalogsInTableDefinitions	false
supportsSchemasInTableDefinitions	false
getDriverName	UnityJDBC
getCatalogSeparator	.
storesMixedCaseIdentifiers	false
storesUpperCaseIdentifiers	false
getDatabaseProductName	UnityJDBC
getDatabaseProductVersion	4.2.238
getDatabaseMajorVersion	4
getIdentifierQuoteString	"
supportsStoredProcedures	false
supportsSavepoints	false
getCatalogTerm	catalog
getSchemaTerm	database
getProcedureTerm	procedure
supportsMultipleResultSets	false
supportsSchemasInDataManipulation	true
supportsCatalogsInDataManipulation	false
supportsCatalogsInProcedureCalls	false

/virtualDemo 1,1 / 1

Please try out the Tools popup by hitting ctrl+t in the SQL Editor. Do it three times to stop this message.

Logs: Errors 0, Warnings 0, Infos 418 64 of 127 MB 0 1:27:44 PM PDT

5. The user can execute an SQL query that spans multiple sources and get a single result. The virtualization is transparent to the user and SQuirreL.

SQL Query joining tables in MySQL, Microsoft SQL Server, Oracle, PostgreSQL, and MongoDB

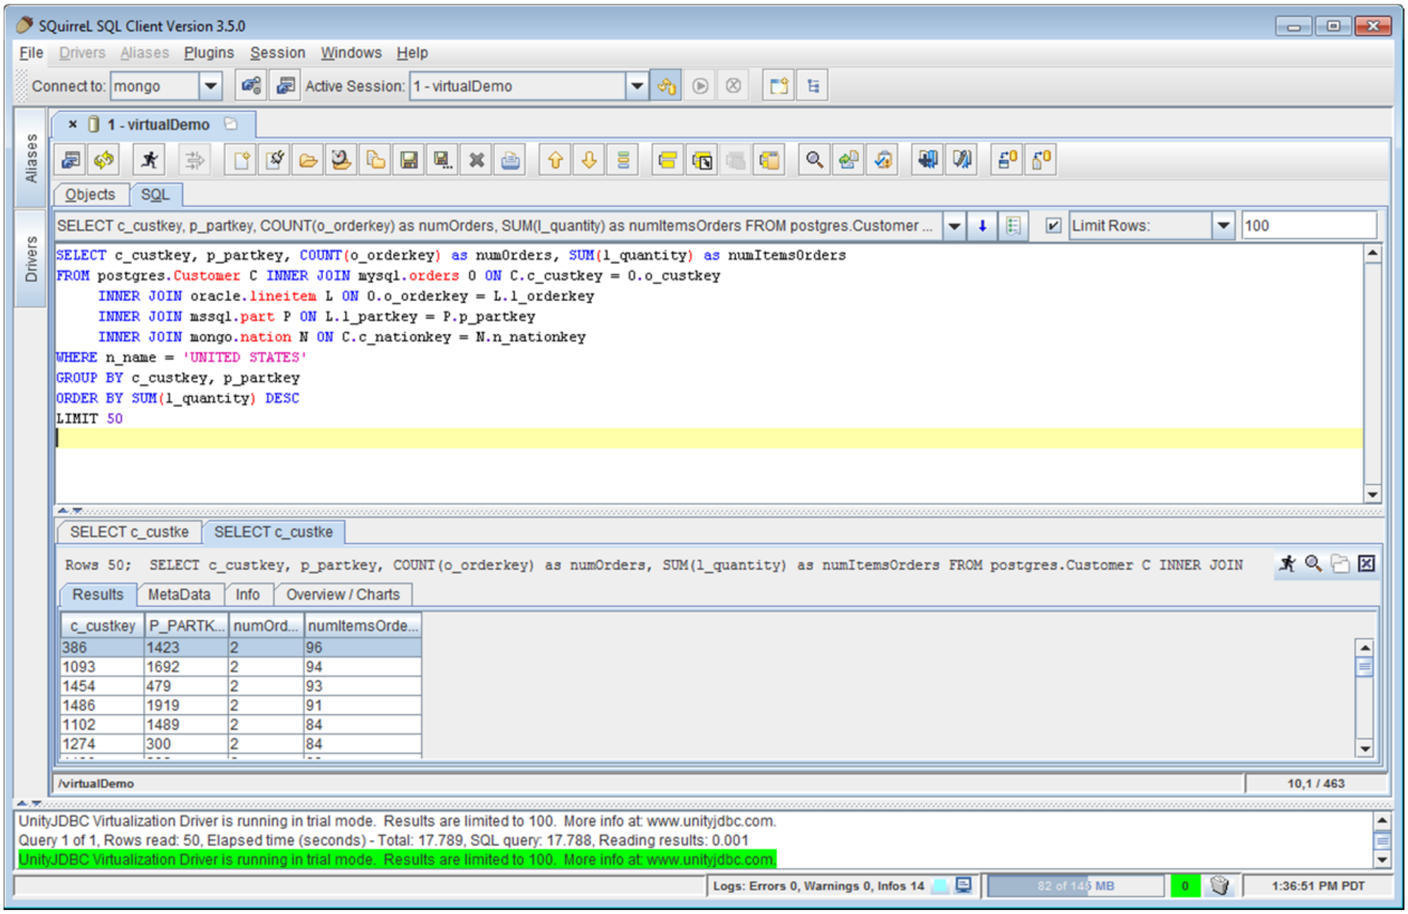

Connect to: mongo

Active Session: 1 - virtualDemo

SELECT c_custkey, p_partkey, COUNT(o_orderkey) as numOrders, SUM(l_quantity) as numItemsOrders FROM postgres.Customer ...

☑ Limit Rows: 100

```
SELECT c_custkey, p_partkey, COUNT(o_orderkey) as numOrders, SUM(l_quantity) as numItemsOrders
FROM postgres.Customer C INNER JOIN mysql.orders O ON C.c_custkey = O.o_custkey
     INNER JOIN oracle.lineitem L ON O.o_orderkey = L.l_orderkey
     INNER JOIN mssql.part P ON L.l_partkey = P.p_partkey
     INNER JOIN mongo.nation N ON C.c_nationkey = N.n_nationkey
WHERE n_name = 'UNITED STATES'
GROUP BY c_custkey, p_partkey
ORDER BY SUM(l_quantity) DESC
LIMIT 50
```

```
Rows 50;  SELECT c_custkey, p_partkey, COUNT(o_orderkey) as numOrders, SUM(l_quantity) as numItemsOrders FROM postgres.Customer C INNER JOIN
```

c_custkey	P_PARTK...	numOrd...	numItemsOrde...
386	1423	2	96
1093	1692	2	94
1454	479	2	93
1486	1919	2	91
1102	1489	2	84
1274	300	2	84

/virtualDemo 10,1 / 463

```
UnityJDBC Virtualization Driver is running in trial mode.  Results are limited to 100.  More info at: www.unityjdbc.com.
Query 1 of 1, Rows read: 50, Elapsed time (seconds) - Total: 17.789, SQL query: 17.788, Reading results: 0.001
UnityJDBC Virtualization Driver is running in trial mode.  Results are limited to 100.  More info at: www.unityjdbc.com.
```

Logs: Errors 0, Warnings 0, Infos 14 | 82 of 145 MB | 0 | 1:36:51 PM PDT

A Translation Example

The plugin uses features of the UnityJDBC driver to perform virtualization and translate functions that are not implemented by certain sources. For example, MSSQL does not support TRIM(), but you can do the same result using RTRIM(LTRIM()). Unity will automatically translate a TRIM() function specified in a MSSQL query to the correct syntax supported by the database.

This translation is supported for common databases and can be freely extended by user-defined functions and translations for each database dialect.

EXPLAIN can be used to understand how a SQL query is translated to queries on individual sources.

Example TRIM() Translation for MSSQL

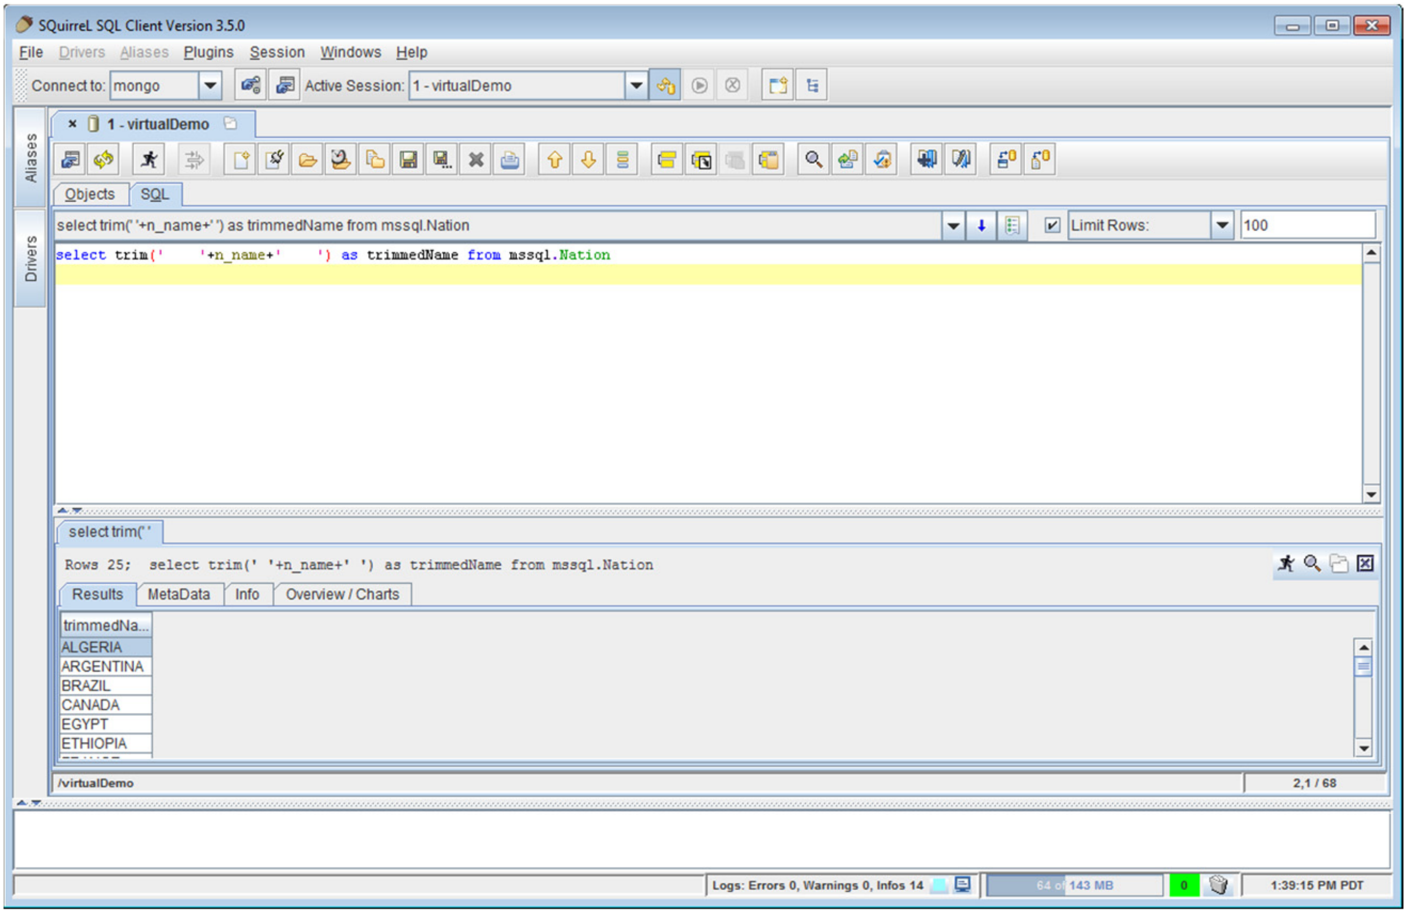

```
select trim('    '+n_name+'    ') as trimmedName from mssql.Nation
```

```
Rows 25;  select trim(' '+n_name+' ') as trimmedName from mssql.Nation
```

trimmedNa...
ALGERIA
ARGENTINA
BRAZIL
CANADA
EGYPT
ETHIOPIA

Encryption

The plugin saves its configuration information in the folder `multisource` under the SQuirreL user folder. If no password is specified when creating an alias, schema and connection information (including passwords) are stored in plain text. To encrypt the configuration files, specify a password during connection (user id is ignored).

Configuration File Location

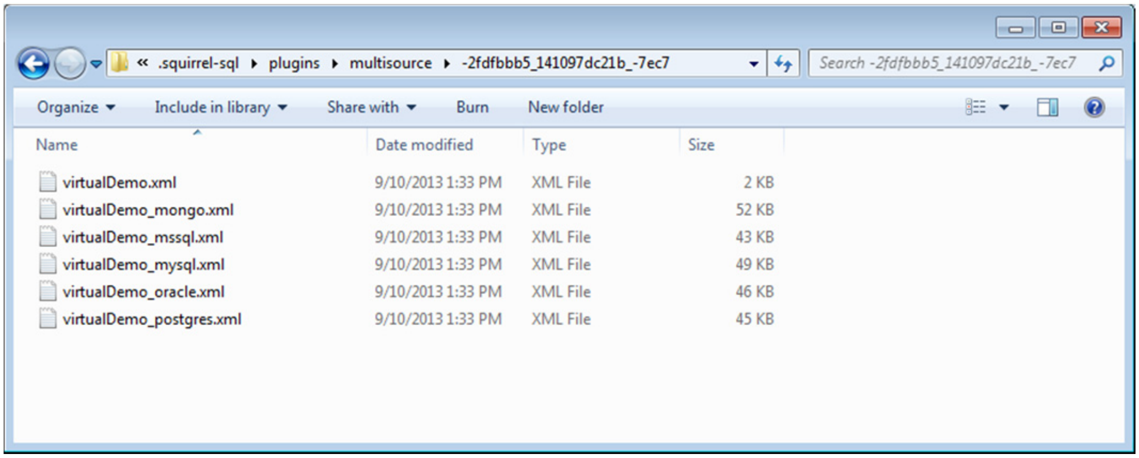

« .squirrel-sql ▸ plugins ▸ multisource ▸ -2fdfbbb5_141097dc21b_-7ec7

Name	Date modified	Type	Size
virtualDemo.xml	9/10/2013 1:33 PM	XML File	2 KB
virtualDemo_mongo.xml	9/10/2013 1:33 PM	XML File	52 KB
virtualDemo_mssql.xml	9/10/2013 1:33 PM	XML File	43 KB
virtualDemo_mysql.xml	9/10/2013 1:33 PM	XML File	49 KB
virtualDemo_oracle.xml	9/10/2013 1:33 PM	XML File	46 KB
virtualDemo_postgres.xml	9/10/2013 1:33 PM	XML File	45 KB

Specify Encryption by Providing a Password

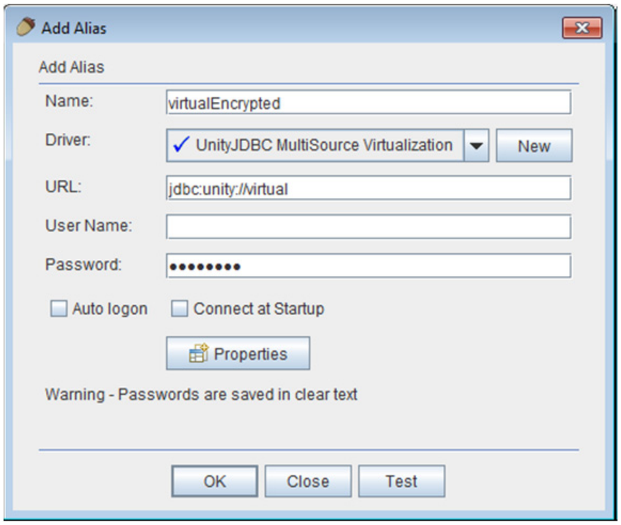

Add Alias

Add Alias

Name: virtualEncrypted

Driver: ✓ UnityJDBC MultiSource Virtualization

URL: jdbc:unity://virtual

User Name:

Password: ••••••••

□ Auto logon □ Connect at Startup

Warning - Passwords are saved in clear text

Virtualization Export for use with Other Programs and Reporting Software

Once a virtualization has been created, all the information necessary is available in the XML files created and stored in the `plugin/multisource` folder in the SQuirreL user directory. These files can be copied and moved to another location and used with any Java/JDBC program or reporting software. These files can be copied directly or exported from within SQuirreL.

Exporting Virtualization Information for Use in Other Programs

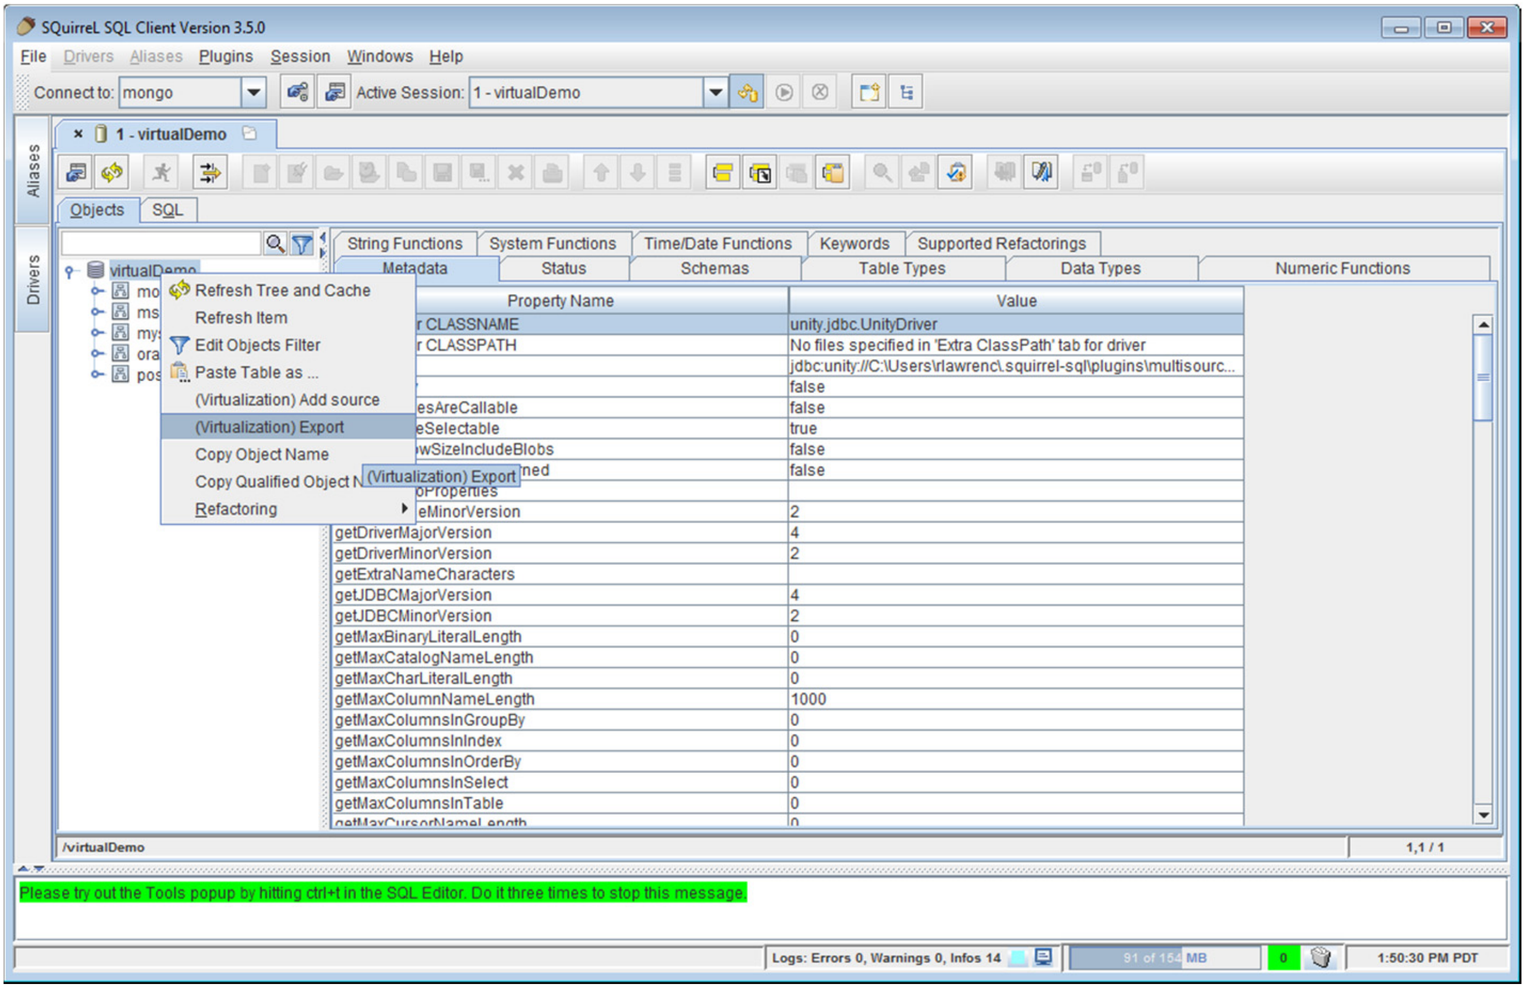

Property Name	Value
r CLASSNAME	unity.jdbc.UnityDriver
r CLASSPATH	No files specified in 'Extra ClassPath' tab for driver
	jdbc:unity://C:\Users\rlawrenc\.squirrel-sql\plugins\multisourc...
	false
esAreCallable	false
eSelectable	true
wSizeIncludeBlobs	false
ned	false
oProperties	
eMinorVersion	2
getDriverMajorVersion	4
getDriverMinorVersion	2
getExtraNameCharacters	
getJDBCMajorVersion	4
getJDBCMinorVersion	2
getMaxBinaryLiteralLength	0
getMaxCatalogNameLength	0
getMaxCharLiteralLength	0
getMaxColumnNameLength	1000
getMaxColumnsInGroupBy	0
getMaxColumnsInIndex	0
getMaxColumnsInOrderBy	0
getMaxColumnsInSelect	0
getMaxColumnsInTable	0
getMaxCursorNameLength	0

/virtualDemo 1,1 / 1

Please try out the Tools popup by hitting ctrl+t in the SQL Editor. Do it three times to stop this message.

Logs: Errors 0, Warnings 0, Infos 14 91 of 154 MB 0 1:50:30 PM PDT

Existing XML configuration files can be used by specifying an absolute or relative path in the JDBC URL. For example, if the user has saved the `sources.xml` file (and the associated schema files for sources in the directory `C:\tmp`, then a JDBC connection URL is: `jdbc:unity://c:\tmp\sources.xml`.

Plugin Limits

The plugin source code, like all of SQuirreL, is released under the GNU Lesser General Public License. The UnityJDBC virtualization driver is released under a commercial license. The UnityJDBC driver included in the plugin is fully functioning with no time limits allowing an unlimited number of sources and queries. The only limitation is the size of the result set is limited to the first 100 rows. (Note there is no limit on the number of rows extracted from each source. `SELECT COUNT(*) FROM` *`table`* with a 1 million row table works as it only returns one result row.) Use `LIMIT 100` to get the first 100 results of a query.

For more information and technical support for the MultiSource plugin contact:

UnityJDBC Support, support@unityjdbc.com, 250-807-9390

UnityJDBC driver information: www.unityjdbc.com