
DbVisualizer 6.0

Users Guide



Copyright Onseven Software AB 1998-2007

<http://www.dbvis.com>

Table of Contents

Getting Started and General Overview.....	8
Introduction	8
Installing.....	8
Installation structure.....	8
Java Properties.....	8
Install license key for DbVisualizer Personal.....	9
Uninstalling the license key.....	9
Useful Resources.....	10
Starting DbVisualizer.....	10
Command line arguments.....	10
The Main Window and Common Components.....	11
Standard Components in the User Interface.....	12
Grid, Graph and Chart.....	12
Context Sensitive Components.....	13
Tooltips.....	13
Grids.....	14
Problem resolution.....	18
How to satisfy the DbVisualizer support team.....	19
Load JDBC Driver and Get Connected.....	20
Introduction	20
What is a JDBC Driver?.....	20
Get the JDBC driver file(s).....	21
Connection Wizard.....	21
Driver Manager.....	26
JDBC Driver Finder.....	26
Loading and Configuring Drivers Manually.....	28
Setup a JDBC driver.....	29
JDBC drivers that requires several JAR or ZIP files.....	31
The JDBC-ODBC bridge.....	31
Loading JNDI Initial Contexts.....	32
Errors (why are some paths red?).....	33
Several versions of the same driver.....	33
Setup a database connection.....	33
Setup using JDBC driver.....	33
Setup using JNDI lookup.....	35
Connection Properties.....	36
Database Profile.....	38
Driver Properties.....	39
Always ask for userid and/or password.....	40
Using variables in the Connection details.....	41
Connect to the Database.....	42
Connections Overview.....	44
Database Objects Explorer.....	46
Introduction	46
Create a Database Connection.....	47
Database Connection object.....	47
Alias.....	48
Default database and schema.....	48
Remove and copy database connection objects.....	48
Database Connection detailed information.....	48
Search.....	49
Organizing Database Connections in Folders.....	49
Connections overview.....	50
Database Objects Tree.....	51
Standard Actions.....	51
Object Actions.....	52
Common Object Actions.....	53
Objects Tree Filtering.....	55
Show Table Row Count.....	56
Database Profiles.....	56

Database Specific Support.....	56
Generic profile.....	57
Catalog/Database object.....	57
Schema object.....	58
Table Type object.....	59
Table object.....	60
Procedure object.....	61
Object Views.....	62
Grid.....	62
Form.....	63
Source.....	64
Table Row Count.....	65
Table Data.....	66
Right-click menu.....	66
Where Filter.....	68
Quick Filter.....	71
Monitor row count.....	73
Editing.....	73
DDL Viewer.....	74
References.....	75
Navigator.....	76
Procedure Editor.....	77
SQL Commander.....	78
Introduction.....	78
Physical Database Connections and Transactions.....	79
Editor.....	79
Database Connection, Catalog and Schema.....	80
Limiting Result Set size (Max Rows/Chars).....	81
Load from and save to file.....	82
Load Recent.....	83
Editor Preferences.....	84
Multiple editors.....	84
Permissions.....	85
Charsets and Fonts.....	85
Key Bindings.....	86
Client side Comments.....	86
Auto Completion.....	87
SQL Formatter.....	90
SQL History.....	91
SQL Bookmarks.....	91
Execution.....	92
SQL->Execute.....	92
SQL->Execute Current.....	92
SQL->Execute Buffer.....	93
SQL->Execute Explain Plan (Oracle, SQL Server and DB2).....	93
Auto Commit, Commit and Rollback.....	96
SQL Scripts.....	97
Anonymous SQL blocks.....	97
Stored Procedures.....	98
Client Side Commands.....	98
@run - run SQL script from file.....	98
@cd <directory> - change directory.....	98
@<file> - run SQL script from file.....	98
@export - export result sets to file.....	99
@exit [nocheck] - Exit DbVisualizer.....	101
@echo - Echo text.....	101
@window iconify - Iconify the main window.....	101
@window restore - Raise the main window.....	101
@desc table - Describe the columns in table.....	101
@spool log - Save log to file.....	101
@stop on error - Stop execution if any error occurs.....	102
@stop on warning - Stop execution if any warning occurs.....	102
Parameterized queries (variables).....	102
Variable Syntax.....	103

Output View.....	104
Log.....	104
Log controls.....	105
Auto clear log.....	105
Result Set.....	105
Result set menu.....	108
Editing.....	108
Multiple result sets produced by a single SQL statement.....	109
Text.....	111
Chart.....	111
DBMS Output (Oracle).....	112
Query Builder.....	113
Introduction.....	113
Current Limitations.....	114
Creating a Query.....	114
Adding Tables.....	115
Joining Tables.....	116
Manually Joining Tables.....	116
Joining Tables Automatically.....	116
Join Properties.....	117
Remove Tables and Joins.....	118
Query Details.....	118
Columns.....	119
Conditions.....	120
Grouping.....	121
Sorting.....	121
SQL Preview.....	122
Testing the Query.....	122
Loading a Query from the SQL Editor.....	122
Properties controlling Query Builder.....	123
Express joins as JOIN clause or WHERE condition.....	123
Table and Column Name qualifiers.....	123
Delimited Identifiers.....	123
Drag style and Diagram Size.....	124
Monitor and Charts.....	125
Introduction.....	125
Monitor an SQL statement	125
Monitor table row count.....	127
Monitor table row count difference.....	128
Monitor window.....	130
Charts.....	131
Chart Controls.....	132
Data.....	132
Layout.....	132
Chart View.....	134
Zooming.....	134
Rotating.....	134
Export.....	135
Create and Alter Table.....	137
Introduction.....	137
Create Table.....	137
Columns tab.....	138
Primary Key tab.....	140
Foreign Keys tab.....	141
Unique Constraints tab (database-specific).....	141
Check Constraints tab (database-specific).....	142
Indexes tab (MySQL only).....	143
SQL Preview.....	143
Execute.....	144
Alter Table.....	145
Create Index.....	146
Columns.....	147
SQL Preview.....	147
Execute.....	147

Edit Table Data.....	148
Introduction.....	148
Features that support editing.....	148
Update and Delete must match one table row.....	148
Edit Multiple Rows.....	149
Data Type checking.....	149
New Line and Carriage Return.....	149
Grid Editor.....	150
Insert row.....	150
Update row.....	151
Delete row(s).....	151
Duplicate row(s).....	151
Copy/Paste.....	151
Paste data from Excel and OpenOffice.....	151
Insert pre-defined values (Set Selected Cells).....	152
Undo Edit(s).....	153
Key Column(s) Chooser.....	153
Preview Changes.....	154
Saving Changes.....	154
Transaction Control.....	154
Permissions.....	155
Errors.....	155
Form Editor/Viewer.....	155
Cell Editor/Viewer.....	157
Binary/BLOB.....	158
Image Viewers.....	159
XML Viewer.....	160
Serialized Java Objects Viewer.....	161
Hex Viewer.....	162
Large text data/CLOB.....	163
Import from File.....	163
Export to File.....	163
Table Data Navigation.....	164
Introduction.....	164
Data Navigation.....	165
Adding Context Information to the Graph.....	167
Arranging the Graph.....	168
Exporting and Printing the Graph.....	170
Procedure Editor.....	171
Introduction.....	171
Create Procedure.....	171
Edit and Compile.....	173
Running in SQL Commander.....	176
SQL Bookmarks.....	177
Introduction.....	177
What's a bookmark in DbVisualizer?.....	177
The Bookmarks Main Menu.....	177
Bookmark Editor.....	178
Bookmark list.....	178
New and History root folders.....	180
SQL Editor.....	180
Monitor information.....	180
The Note field.....	180
Executing an SQL bookmark or folder of SQL bookmarks.....	180
Tool Properties.....	182
Customizing DbVisualizer.....	182
The user preferences (XML) file.....	182
General Settings.....	182
Appearance.....	183
Fonts.....	185
Key Bindings.....	185
Database Connection.....	188
Driver Manager.....	188
Permissions.....	188

SQL Commander Permissions.....	188
Inline Editor Permissions.....	189
Time Zone.....	189
Data Formats.....	190
Date, Time and Timestamp formats.....	190
Number formats.....	190
Table Data.....	191
Variables.....	191
Transaction.....	191
Bookmarks.....	191
Monitor.....	191
Form Viewer.....	192
Grid.....	192
Copy.....	192
Colors.....	192
Binary/BLOB and CLOB Data.....	192
SQL Editor.....	193
Statement Delimiters.....	193
SQL Formatting.....	193
Auto Completion.....	194
Comments.....	194
Debug.....	194
Database Settings.....	195
Authentication.....	196
Delimited Identifiers.....	196
Qualifiers.....	196
Physical Connection.....	196
Transaction.....	197
SQL Statements.....	197
Connection Hooks.....	198
Objects Tree.....	198
SQL Editor.....	201
Query Builder.....	201
Database Specific settings.....	201
Data Types (Oracle).....	201
Data Types (DB2 and JavaDB/Derby).....	201
Explain Plan (Oracle, SQL Server and DB2).....	202
Explain Plan (Oracle).....	202
Explain Plan (DB2).....	202
System Tables (Oracle).....	202
Export, Import and Print.....	203
Introduction.....	203
Export Grid data.....	203
Settings page.....	204
Output Format.....	205
Encoding.....	205
Data Format.....	205
Quote Text Data.....	206
Options.....	206
Settings.....	207
Data page.....	208
Generate Test Data.....	209
Preview.....	211
Output Destination.....	211
Export Text data.....	213
Export Graph data.....	213
Export Chart data.....	214
Import Table Data.....	215
Source File.....	216
Settings.....	216
Data Formats.....	219
Import Destination.....	221
Import process.....	222
Print.....	223

Printer Setup.....	223
Grid, Chart and Plain Text.....	223
Graph.....	224
Print Preview.....	225
Plug-in Framework.....	226
Introduction.....	226
What features in DbVisualizer relies on the database profile?.....	226
How does DbVisualizer know what database profile to use?.....	228
XML structure.....	228
XML skeleton.....	229
<DatabaseProfile>.....	230
<Commands> - The SQLs used to interact with the database.....	230
<Command>.....	231
Result set.....	231
<Input> - Setting command input.....	232
<Output> - Redefine command output.....	233
<ObjectsTreeDef> - Definition of the Database Objects Tree.....	233
<GroupNode> - Static objects used for grouping.....	234
<DataNode> - Dynamic objects created via SQL.....	235
<Command>.....	236
<SetVar>.....	237
<ObjectsViewDef> - Definition of the Object Views.....	238
<ObjectView>.....	239
<DataView>.....	239
Viewers.....	240
<Command>.....	250
<Message>.....	250
Extending ObjectView.....	251
<ObjectsActionDef> - Definition of user actions.....	252
Variables.....	253
<ActionGroup>.....	255
<Action>.....	255
<Input>.....	256
<SetVar>.....	261
<Confirm>.....	261
<Result>.....	262
<Command>.....	262
Conditional processing.....	262
When are conditional expressions processed?.....	262
Conditional processing when database connection is established.....	263
Conditional processing during command execution.....	263
Current limitations.....	264
DB2 for Windows/Linux - Database Specific Support.....	265
Introduction.....	265
Database Objects Tree.....	265
Explain Plan.....	267
Informix IDS - Database Specific Support.....	268
Introduction.....	268
Database Objects Tree.....	268
JavaDB/Derby - Database Specific Support.....	271
Introduction.....	271
Database Objects Tree.....	271
Mimer SQL - Database Specific Support.....	273
Introduction.....	273
Database Objects Tree.....	273
MySQL - Database Specific Support.....	275
Introduction.....	275
Database Objects Tree.....	275
Oracle - Database Specific Support.....	277
Introduction.....	277
Database Objects Tree.....	277
Application Object Support (PL/SQL).....	279
DBMS_OUTPUT Support.....	279
Explain Plan.....	279

Data Type support.....	279
PostgreSQL - Database Specific Support.....	280
Introduction.....	280
Database Objects Tree.....	280
SQL Server - Database Specific Support.....	282
Introduction.....	282
Database Objects Tree.....	282
Explain Plan.....	284
Sybase ASE - Database Specific Support.....	285
Introduction.....	285
Database Objects Tree.....	285
set showplan.....	287

Getting Started and General Overview

Introduction

DbVisualizer is a feature rich, intuitive multi-database tool for developers and database administrators, providing a single powerful interface across a wide variety of operating systems. With its easy-to-use and clean interface, DbVisualizer has proven to be one of the most cost effective database tools available, yet to mention that it runs on all major operating systems and supports all major RDBMS that are available. Users only need to learn and master one application. DbVisualizer integrates transparently with the operating system being used.

This document gives a overview, installation tips and general information about the product.

Note: All documents in the Users Guide are primarily focusing on the DbVisualizer Personal edition. Some of the described features are not available in the Free edition.

The screenshots throughout the users guide are produced on Windows XP using the Windows Look and Feel, but DbVisualizer lets you choose among other Look and Feels as well.

Installing

Installing DbVisualizer is no different then installing other modern products. The standard installation procedure is performed using a graphical application, and you just need to click through the questions that are displayed. Follow the [instructions](#) at the DbVisualizer web site if you need information on how to start the installation procedure specifically for your platform.

Installation structure

The installer and launcher for DbVisualizer is based on the install4j™ product (<http://www.install4j.com>). The structure of the installation directory (referred as DBVIS-HOME throughout the users guide) contains the following. (The exact content may differ between platforms):

```
.install4j/  
doc/  
lib/  
resources/  
wrapper/  
dbvis.vmoptions  
dbvis.exe  
README.txt  
uninstall.exe
```

The **dbvis.exe** file is used to start DbVisualizer. The remaining files and directories are only of interest if you need to do nonstandard customization. For information on how to increase the memory for the Java process that runs DbVisualizer, and also on how to modify the Java version being used, please read the on-line [FAQ](#) for the latest information.

Java Properties

DbVisualizer utilizes a few Java properties that you can use to modify characteristics of the application. These

DbVisualizer specific properties are available in the DBVIS-HOME/resources/dbvis-custom.prefs file.

Note: You rarely need to modify these properties, as the default values are sufficient for most usage.

The following are the properties handled by DbVisualizer:

Property	Description
<code>dbvis.driver.ignore.dir=lib:resources:.install4j</code>	Specify directories from DBVIS-HOME that should not be listed in the Driver Manager "System Classpath" list. Directories are separated with ":". Accepted values: one or several directory names starting from DBVIS-HOME.
<code>dbvis.grid.encode=false</code>	Specifies if encoding of data in result set grids will be performed or not. If set to true then make sure the <code>dbvis.grid.fromEncode</code> and/or <code>dbvis.grid.toEncode</code> are also set.
<code>dbvis.grid.fromEncode=ISO8859_1</code>	Encoding used when translating text data that is fetched from the database
<code>dbvis.grid.toEncode=GBK</code>	Encoding used when translating data that will appear in the result set grid
<code>dbvis.usegetObject=false</code>	Specifies if the generic <code>ResultSet.getObject()</code> method in JDBC will be used in favor of the data type specific get methods or not. Default is false.
<code>dbvis.savedatacolumns=false</code>	Column layout changes such as reordering and/or visibility are saved for all grids in the Objects Views *except* for the "Data" grid. This property can be used to also include the layout in the "Data" grid. Note: This will result in DbVisualizer saving the layout for each table that is displayed in the Data grid = huge XML file...
<code>dbvis.disabledataedit=false</code>	Specifies if table data editing should be completely disabled, i.e. the form and inline editors. Note: This only has an effect when used with a licensed edition.
<code>dbvis.showactionresult=false</code>	This defines whether the result for all actions should be displayed or only failures (default).

Note: These properties may change in future versions of DbVisualizer. Some are also experimental and may be removed or instead introduced in the DbVisualizer GUI.

Install license key for DbVisualizer Personal

If you have a license key file for DbVisualizer Personal, then start DbVisualizer and open the Help->License Key window. Enter the name of the license file in the License Key File field, or launch the file chooser by pressing the "..." button to the right of the license file field. Once the file is loaded, press the Install License button.

Uninstalling the license key

If you ever need to uninstall the license key, you can do so by removing (or renaming) the following file:

Operating System	File Name
Windows	C:\Documents and Settings\ <user>\.dbvis\dbvis.license</user>
UNIX/Linux	/home/<user>/.dbvis/dbvis.license

Useful Resources

Resources related to DbVisualizer that are useful:

1. The home of [DbVisualizer](#)
2. The [FAQ](#) which is regularly updated with frequently asked questions and known problems
3. The [User Guide](#)
4. The [Databases and JDBC Drivers](#) online page. This page gives information about supported databases and JDBC drivers
5. The Ming [forums](#)
6. The online [problem report](#) form. This is the recommended channel for product support and general questions

Starting DbVisualizer

How to start DbVisualizer depends on the operating system you are using.

- Windows
Locate the DbVisualizer submenu in the Start menu. Select the DbVisualizer entry in that menu
- Linux/Unix
Open a shell and change directory to the DbVisualizer installation directory. Execute the `dbvis` program
- Mac OS X
Double click on the DbVisualizer application or DbVisualizer.app application bundle.

Command line arguments

DbVisualizer supports a range of command line arguments. These are listed in the Help->About menu choice, under the Command Line tab, in DbVisualizer.

```
Usage: dbvis [-help] [-up <path>] [-sqlfile <path>]
           [-windowtitle <title>]
           [connect options] [remote options]
```

General Options:

<code>-help</code>	Display this help
<code>-up <path></code>	Use an alternate user preferences file
<code>-sqlfile <path></code>	Load file into the SQL Commander editor
<code>-encoding <encoding></code>	Optional encoding for loaded file
<code>-windowtitle <title></code>	Additional window title
<code>-execute</code>	Will execute SQL file automatically
<code>-invisible</code>	No windows will be displayed

Driver Connect Options:

<code>-driver</code>	Setup and connect using the following Driver options:
<code>-alias <name></code>	Database alias
<code>-drivername <name></code>	Driver name
<code>-path <path></code>	Path to driver class
<code>-class <class></code>	JDBC Driver class

```
-url <url>           Connection URL
-userid <user>       Userid to connect as
-password <pw>       Connect password
```

JNDI Connect Options:

```
-jndi Setup and connect using the following JNDI options:
-alias <name>        Database alias
-drivername <name>   Driver name
-path <path>         Path to initial context class
-class <class>       Initial context class
-url <url>           Provider URL
-lookup <name>       Lookup name
-userid <user>       Userid to connect as
-password <pw>       Connect password
```

Remote Options:

```
-attachremote Attach to remote DbVisualizer instance
-enableremote Enable remote attachment
-host <host>      Remote host name (default: localhost)
-port <port>      Remote port (default: 8787)
```

The Main Window and Common Components

As you can see in the in the screenshot below, the DbVisualizer interface has a Database Objects Tree to the left and two tabs to the right.

Database Objects Tree

This tree keeps (at the top level) all the Database Connection objects (or folder objects, used to organize Database Connections). Use this tree to navigate and explore the database. Clicking on an object will change the view in the Object View tab to show details about the selected object.

Object View

This tab shows detailed information about the object represented by the selected tree node. The content of the Object View tab depends on the type of the selected object.

SQL Commander

The SQL Commander lets you execute any SQL statements and scripts.

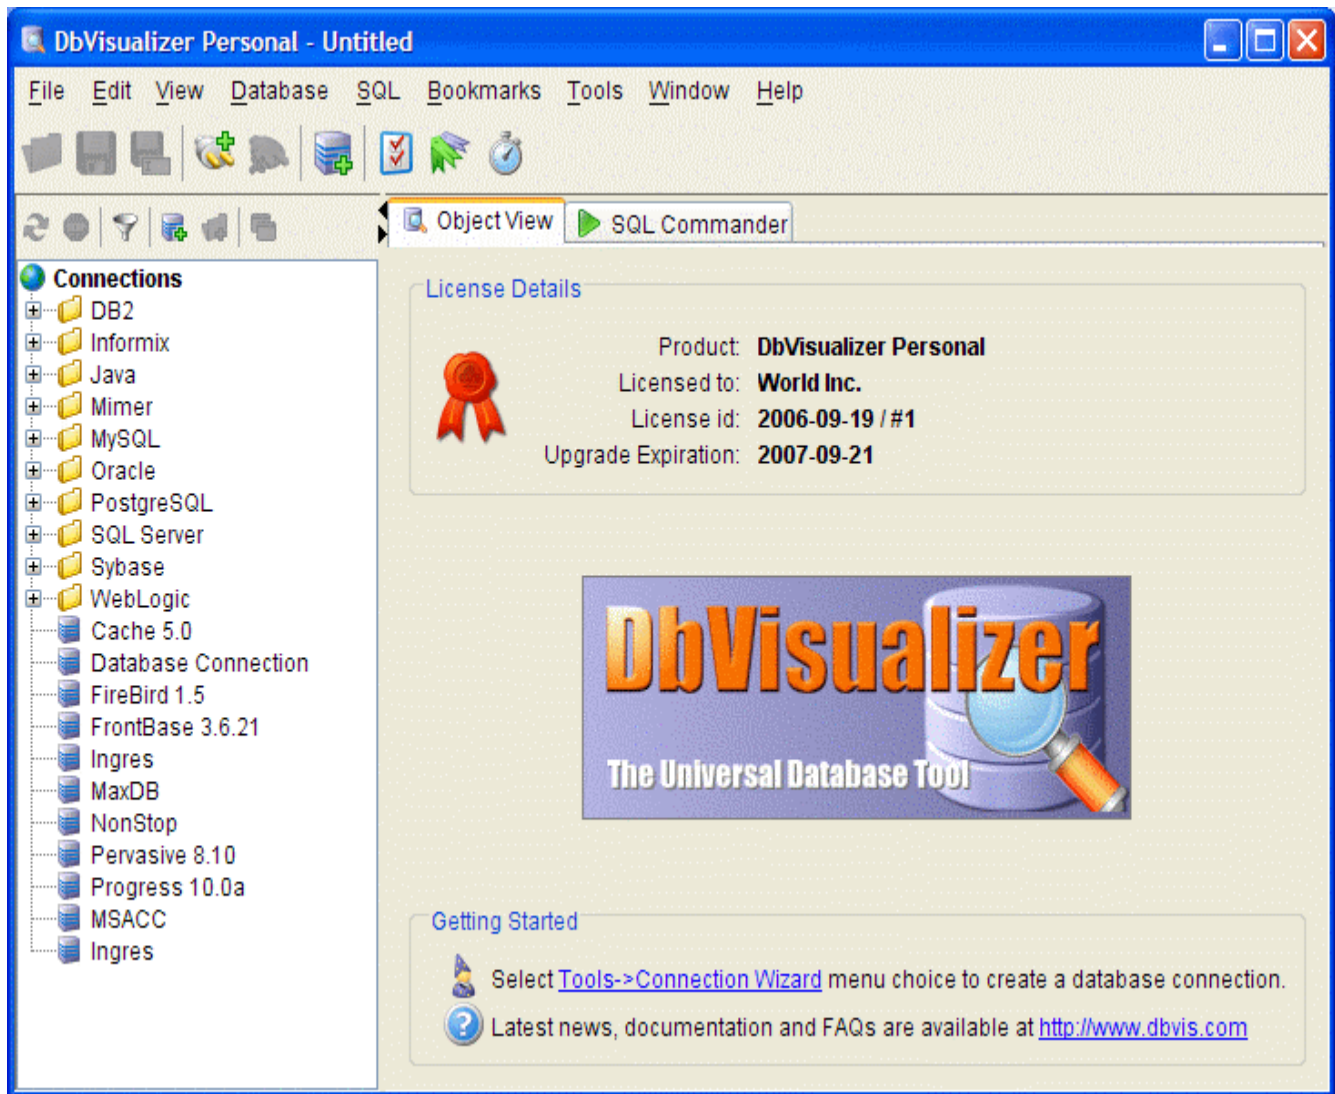


Figure: The DbVisualizer main window

Standard Components in the User Interface

The following sections introduce generic features and components that you find in many parts of DbVisualizer.

Grid, Graph and Chart

Grid, graph and chart are three terms that are often used in the application and in the documentation. The following screenshots show what they represent.

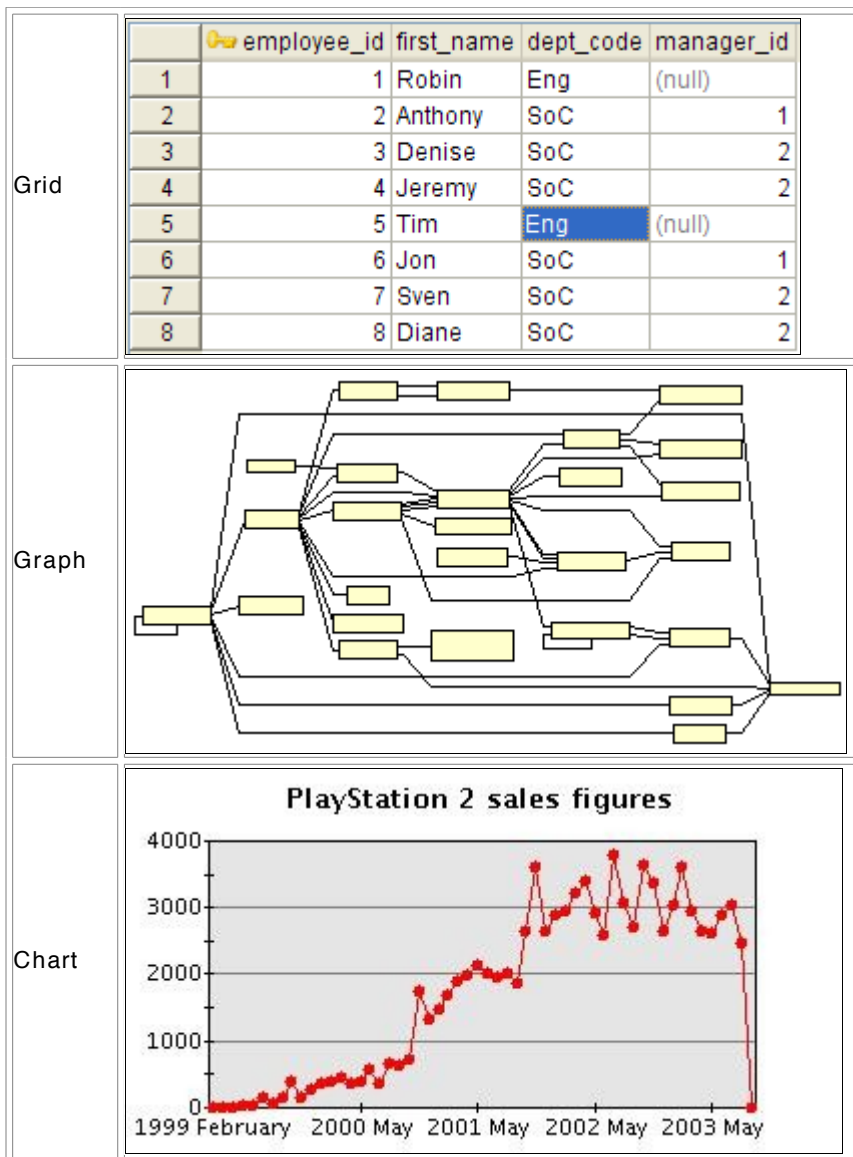


Figure: The grid, graph and chart terms

Note: The documentation uses the term `grid` for the user interface component that represents tabular data, rather than `table`, to avoid confusion with a database table.

Context Sensitive Components

All components in the user interface (e.g., buttons and menu items) are context sensitive. They are enabled only if they can be used in the current scope.

Tooltips

Tooltips are used to provide more details about a component. They are also used to express status information. An example is the grid column header tooltip that shows information about the column. To see a tooltip, let the mouse hover over an area of the user interface, e.g., a button or grid header. If there is a tooltip for the area, it

will pop up in about a second.

```
employee_id INTEGER
Not NULL, Key Column
JDBC: INTEGER (type: 4), Java: Long
```

Figure: Tooltip example

Grids

Grids are used heavily in DbVisualizer and require a brief introduction.

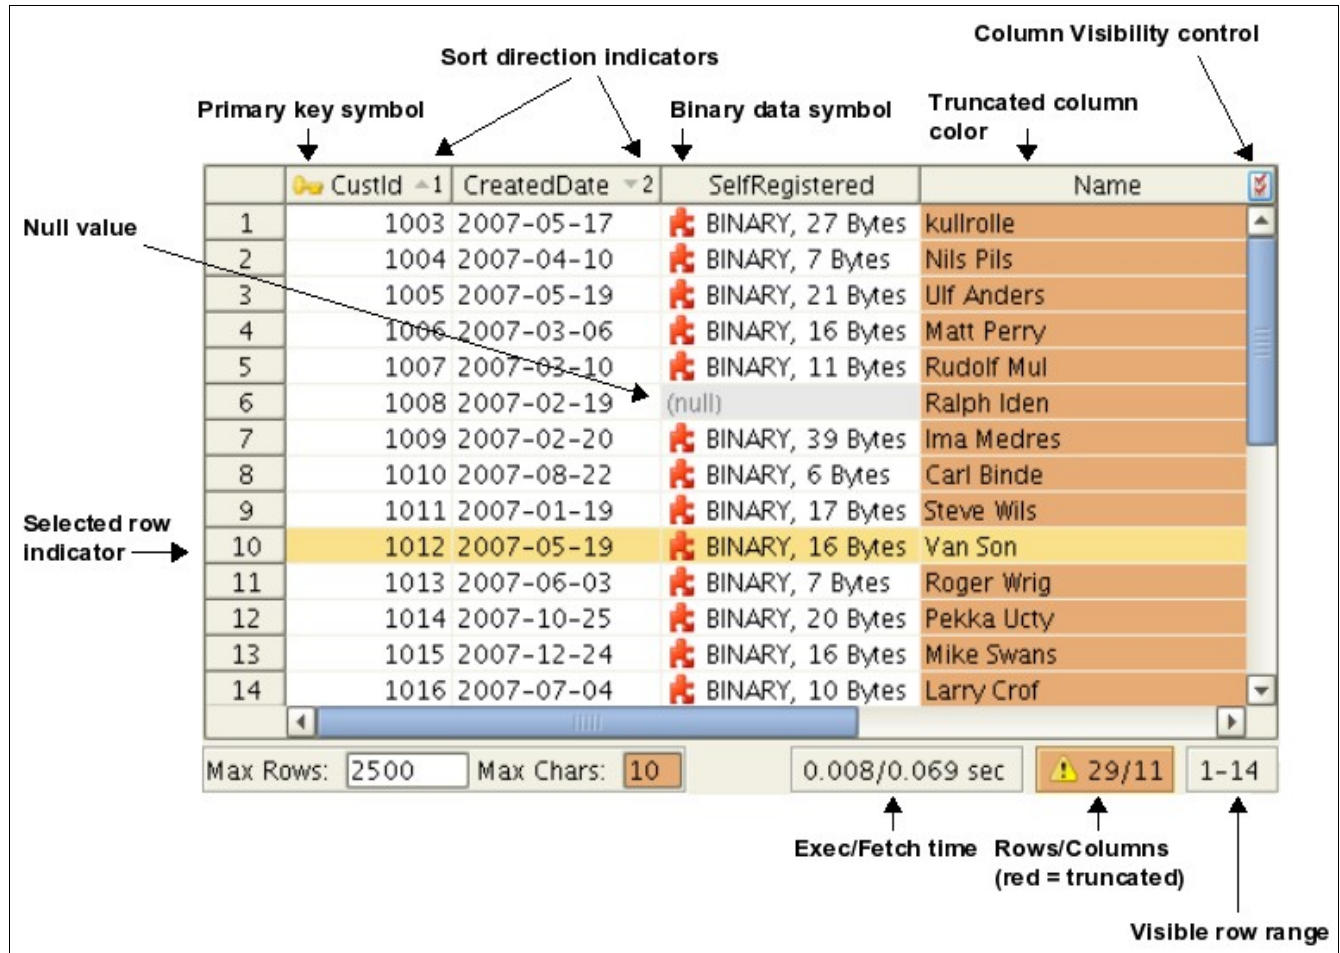


Figure: Grid overview

The screenshot shows the grid and controls that are available for the grid in the Data tab, but the differences are minor compared to grids used in other places.

Sorting

You can sort the grid based on the values in one or more columns. When you click on a column header, the grid is sorted in ascending order on the values in that column, indicated by an up-arrow in the column header. If you click the same column header again, the grid is sorted in descending order, indicated by a down-arrow in the

column header. If you click a third time, the data is shown in the order it was received from the database and the sort indicator disappears.

To sort on more than one column, Ctrl-click (keep the Ctrl key pressed when clicking) on additional columns. The grid is then sorted on the values in the first column you clicked on (indicated with a 1 next to the arrow), and then all rows with the same value in the first column are sorted on the values in the second sort column (indicated with a 2 next to the arrow), and so on.

Right-click menu

The generic right-click grid menu contains the following operations:

















 Select All	Ctrl-A
 Select Row	Ctrl+Shift-J
 Copy Selection	Ctrl-C
 Copy Selection (With Column Header)	Ctrl-H
 Export...	Ctrl+Alt-E
 Export Selection...	
 Print...	
 Print Selection...	
 Print Preview...	
 Save Selected Cell...	Ctrl+Shift-S
 Reload	Ctrl-R
 Find...	Ctrl-F
 Browse Row in Window...	
 Browse Cell in Window...	
 Describe Data...	
 Calculate Selection...	Ctrl+Shift-C

Figure: Grid right click menu

Menu Choice	Description
Select All	Select all cells (i.e., all rows and columns) in the grid
Select Row	Select all cells in the row
Copy Selection	Copy all selected cells onto the system clipboard
Copy Selection (With Column Header)	Copy all selected cells including column header onto the system clipboard
Export	Copy the export dialog
Export Selection	Export the selection using the standard export feature
Print	Open the print dialog for printing the complete grid
Print Selection	Open the print dialog for printing just the selected rows/columns
Print Preview	Open the print preview dialog
Save Selected Cell	Save the value of the selected cell to a file, selected with a file chooser dialog
Reload	Reload the grid with data from the database

Find	Open the find dialog
Browse Row in Window	Display all data for the selected row in a separate window. Note: for a read/write grid, this entry is named Edit Row in Window.
Browse Cell in Window	Display the cell value in a separate window. This is especially useful for BLOB/CLOB data. Note: for a read/write grid, this entry is named Edit Cell in Window.
Describe Data	Show detailed information about the columns in the grid
Calculate Selection	Displays some metrics about the current selection. This is especially useful for numeric fields. Read more in Calculate Selection below.

The menu may contain additional entries based on the current scope, e.g., entries for editing cell values for a read/write grid.

Calculate Selection

The Calculate Selection feature performs basic calculations on the current selection. It is primarily used for selections of cells holding numbers. The following is an example of what it shows.

IT_PROG	4800	(null)
IT_PROG	4800	(null)
IT_PROG	4200	(null)
FI_MGR	12000	(null)
FI_ACCOUNT	9000	(null)
FI_ACCOUNT	8200	(null)
FI_ACCOUNT	7700	(null)
FI_ACCOUNT	7800	(null)
Number of Cells:	6	Sum: 43,000
Valid Numbers:	6	Min: 4,200
Null Values:	0	Avg: 7,166.67
Bytes:	25	Max: 12,000
PU_CLERK	2800	(null)
PU_CLERK	2600	(null)
PU_CLERK	2500	(null)

Figure: The calculate selection popup

Property	Description
Number of Cells	Shows the number of cells in the selection.
Valid Numbers	Lists the number of valid numbers in the selection.
Null Values	Shows the total of null values in the selection.
Bytes	Shows the total number of bytes in the selection after that the data has been translated to text
Sum	Shows the total summary of the selection
Min	Shows the minimum number in the selection
Avg	Shows the average value of the selection by doing sum / number of valid numbers
Max	Shows the maximum number in the selection

Click the red cross icon, or anywhere in the popup, to close it.

Column Visibility

The Column Visibility dialog controls which columns you want to appear in a grid. You open the column visibility dialog by clicking the button above the vertical scrollbar in the grid.

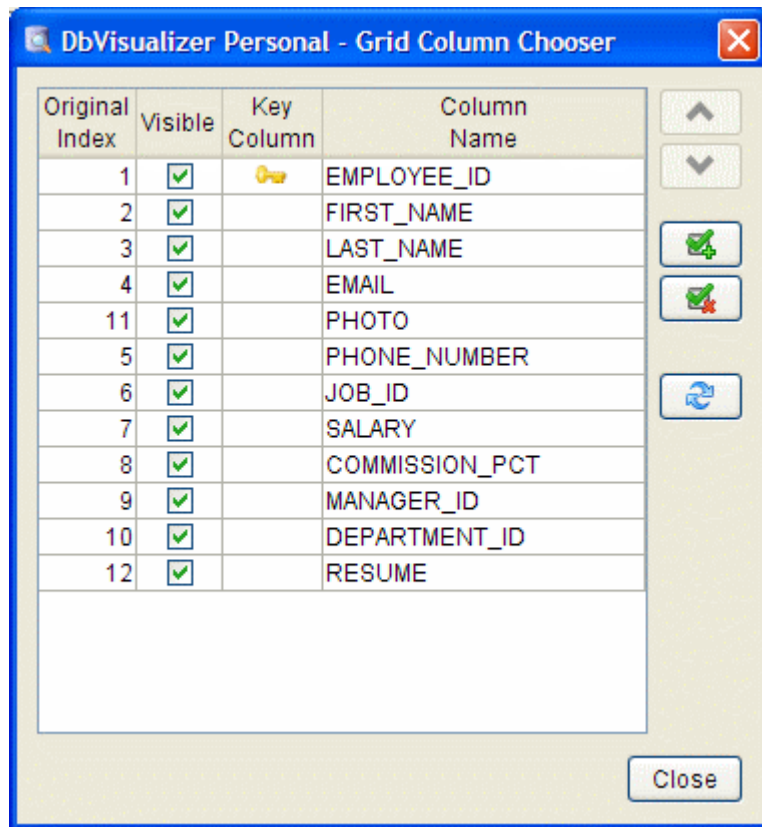


Figure: The Column Visibility dialog

The Column Visibility dialog shows all columns that are available in the grid.

The checkmark in front of a column name indicates that the column is visible in the grid, while an unchecked box indicates that it is excluded from the grid. Click the checkmark to change the visibility of a column. You can change the visibility for all columns at once using the two visibility buttons in the dialog.

The order of the columns can also be adjusted in this dialog. Just select a row and use the Up and Down buttons to move it up (left in grid) or down (right in grid).

If you want to revert your changes, you can click on the Default Layout button to reset the grid, i.e., making all column visible and put them in their default locations.

Note 1: Modifications of column visibility, size and order are saved between invocations of DbVisualizer for all grids in the various Object View tabs except for the Data tab.

Note 2: If you modify the column visibility in the Data tab, the changes persists throughout the session. For instance, if you remove the Name column in the Data tab for the table EMPLOYEE, the Name column remains excluded when you reload the table or come back to the Data tab for that table later in the same session. You must manually make it visible again to bring it back. The changes are, however, reset when you restart the application.

Auto Resize

The column header right-click menu contains a number of options for automatic resizing of column widths.

All Columns (consider Header)
All Columns (ignore Header)
Current Column (consider Header)
Current Column (ignore Header)
Default Column Widths

Figure: Auto Resize menu

Menu Choice	Description
All Columns (consider Header)	Resize all columns to fit the widest cell value, or the column header if it is wider than any cell value
All Columns (ignore Header)	Resize all columns to fit the widest cell value
Current Column (consider Header)	Resize the current column to fit the widest cell value, or the column header if it is wider than any cell value
Current Column (ignore Header)	Resize the current column to fit the widest cell value
Default Column Widths	Set all column widths to their default width

Quick Filter

All areas that hold a grid in DbVisualizer also provide a Quick Filter field.

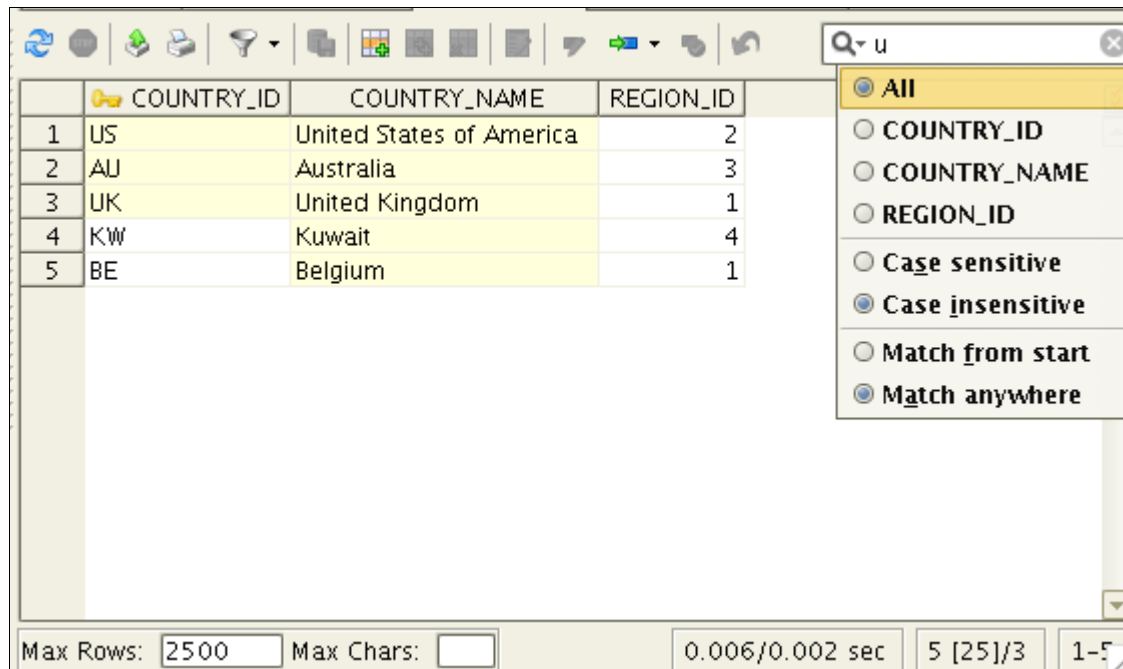


Figure: Grid with Quick Filter

When you type in the Quick Filter field, DbVisualizer matches the value with cell values in the grid and filters out all rows that do not have a match in at least one cell. The Quick Filter pull-down menu (click on the down arrow

next to the magnifying glass) lets you choose if the filter should match cells in all columns or just one selected column, case or case insensitive matching, and where in the cell the value must match.

Problem resolution

Even though we make our very best to ensure the quality of DbVisualizer, you may run into problems of different kinds. The runtime environment for DbVisualizer is rather complicated when it comes to tracking the source of a problem, since it's not only DbVisualizer that may cause the problem but also the JDBC driver, or even the database engine.

There are a few things that you can try before reporting a problem, depending on the nature of the problem:

1. Make sure you are using the latest version of [Java](#) available for your platform (Java 5 or later).
2. Make sure you are using a the [version](#) of the JDBC driver that we've tested DbVisualizer with, or a later, production quality version.
3. Read the DbVisualizer [FAQ](#).
4. Check the online [Forums](#).
5. Read the DbVisualizer [Users Guide](#).

If you cannot find a solution to resolve the problem, please do the following and email us the debug output:

- Problem during installation or when starting DbVisualizer
[Debugging the installer](#)
- Error or problem while using DbVisualizer
[Debugging DbVisualizer](#)

Use the DbVisualizer [problem report form](#) or email support@dbvis.com. We appreciate detailed reports as well as screenshots when possible.

How to satisfy the DbVisualizer support team

Quite often we get incomplete problem reports and need to follow up for additional information. If you encounter a problem, please follow these steps to include the details we need to help you:

1. Select the Connection tab
2. In the Connection Message area, select the right click menu
3. In the menu, select Copy
4. This copies the system details to the clipboard. Then paste the details into an email or in the problem report form.
5. In addition, we really appreciate it if you provide us with screenshots. An image says more then ... you know.

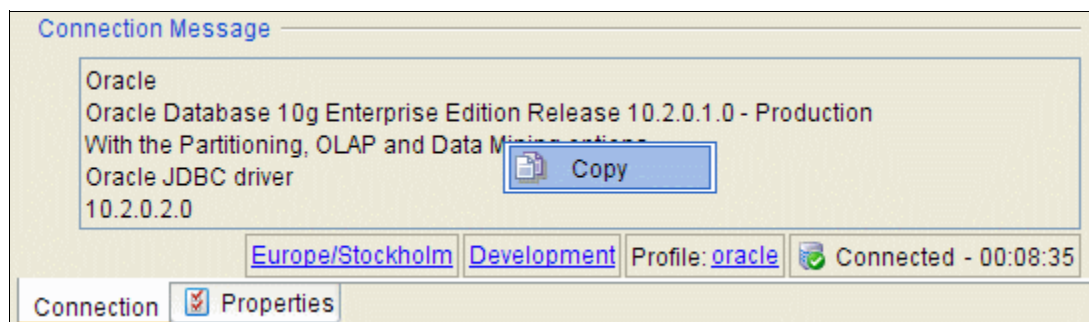


Figure: The connection message right click menu

Load JDBC Driver and Get Connected

Introduction

This document describes the way JDBC drivers are managed in DbVisualizer and all aspects about getting connected to your database(s).

Fast track: If you are impatient, please go ahead and read the [Connection Wizard](#) section. It is the recommended way to create database connections in DbVisualizer.

What is a JDBC Driver?

DbVisualizer is, as you know, a generic tool for administration and exploration of databases. DbVisualizer is in fact quite simple, since it does not deal directly with how to communicate with each database type. The hard job is done by a JDBC driver, which is a set of Java classes that are either organized in a directory structure or collected into a JAR or ZIP file. The magic of these JDBC drivers is that they all match the JDBC specification and the standardized Java interfaces. This is what DbVisualizer relies on. A JDBC driver implements all details for how to communicate with a specific database and database version, and there are a range of drivers from the database vendors themselves and 3:rd party authors. To establish a connection with a database, DbVisualizer loads the driver and then get connected to the database through the driver.

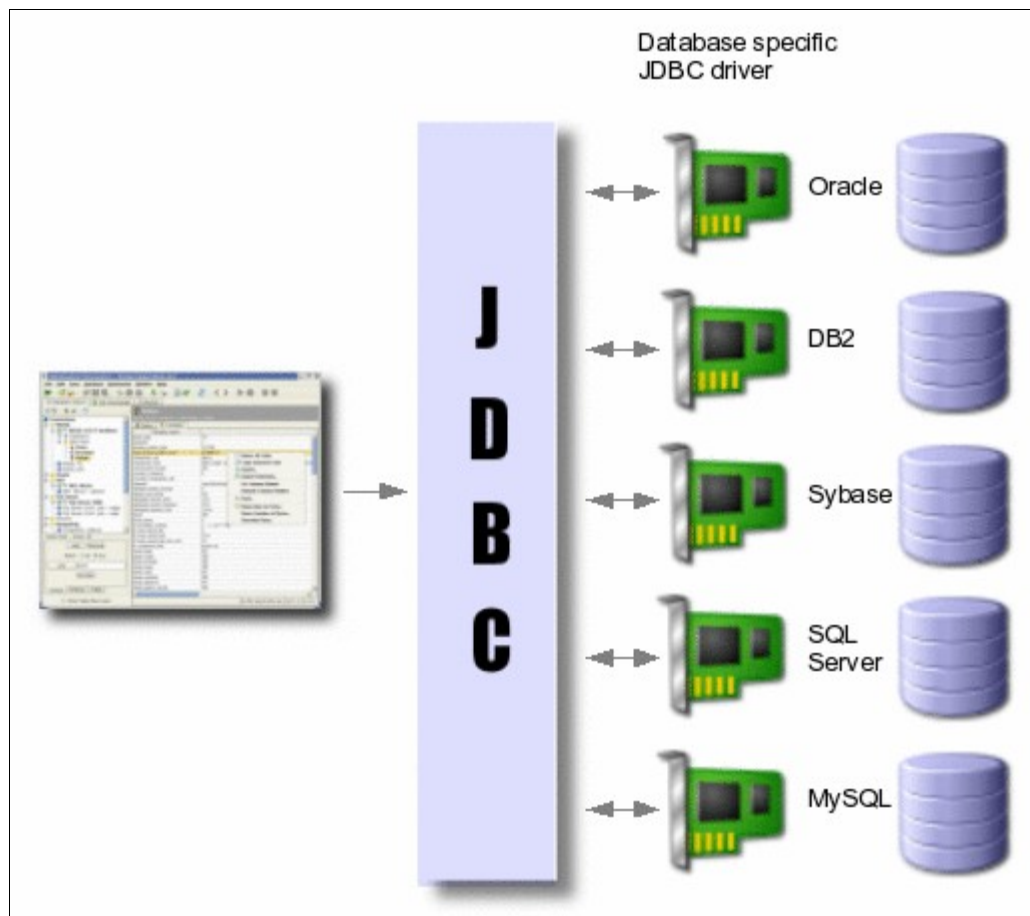


Figure: The runtime environment with the JDBC interface, JDBC driver and sample databases

It is also possible to obtain a database connection using the Java Naming and Directory Interface (JNDI). This technique is widely used in enterprise infrastructures, such as application server systems. It does not replace JDBC drivers but rather adds an alternative way to get a handle to an already established database connection. To enable database "lookup's" using JNDI, an Initial Context implementation must be loaded into the Driver Manager. This context is then used to lookup a database connection. The following sections describe the steps for getting connected using a JDBC Driver, and also how to use JNDI to obtain a database connection.

A complete JDBC driver typically consists of a number of Java classes, located in a JAR, ZIP or a folder, that need to be loaded into the DbVisualizer driver manager. DbVisualizer automatically recognizes the classes that are used to initiate the connection with the database and presents them in the Driver Class list. You must select the correct class in this list to make sure DbVisualizer successfully can initiate the connection. Consult the driver documentation for information of which class to select, or if the number of classes found are low, figure it out by trying each of them. More about this in the following sections.

Get the JDBC driver file(s)

DbVisualizer comes bundled with all commonly used JDBC drivers that have a license that allows for distribution with a third party product. Currently, drivers for DB2, JavaDB/Derby, Mimer, MySQL, and PostgreSQL, as well as the jTDS driver for SQL Server and Sybase, are included with DbVisualizer. If you only need to connect to databases of these types, you can skip the rest of this section and jump straight to the [Connection Wizard](#) section, because by default, DbVisualizer configures all these drivers automatically the first time you start DbVisualizer.

If you need to connect to a database that is not supported by a bundled JDBC driver, you must get a JDBC driver that works with your database type and version. The following online web page contains an up-to-date listing of the database/driver combinations we have tested:

[Databases and JDBC Drivers](#)

Information about almost all drivers that are available is maintained by Sun Microsystems on this page:

[JDBC Data Access API - Drivers](#)

Download the driver to an appropriate directory. Make sure to read the installation instructions provided with the driver. Some drivers are delivered in ZIP or JAR format but need to be unpacked to make the driver files visible to the Driver Manager. The [Databases and JDBC Drivers](#) web page describes where you can download each driver and also what additional steps may be needed to install and load the driver in DbVisualizer.

(Drivers are categorized into 4 types. We're not going to explain the differences here, just give you the hint that the "type 4," aka "thin," drivers are the easiest to maintain, since they are pure Java drivers and do not depend on any external DLL's or dynamic libraries. Even though DbVisualizer works with any type of driver, we recommend that you get a type 4 driver if there is one for your database).

When you have downloaded the JDBC driver into a local folder (and unpacked it, if needed), you can go ahead and create a connect with the Connection Wizard, as described in the next section. You will then asked to load the driver files when the wizard needs them. Alternatively, you can move (or copy) the JDBC driver files to the DBVIS_HOME/jdbc folder, where they will be picked up and loaded automatically by the JDBC Driver Finder the next time you start DbVisualizer. You can read more about this option in the [JDBC Driver Finder](#) section.

Connection Wizard

The Connection Wizard greatly simplifies the steps needed to load the JDBC driver and create a new database connection. You just enter information about the driver file(s) and the connection data on a few wizard pages, and the wizard handles all the details. Once the new database connection has been created, it appear in the database objects tree.

Note: The wizard cannot be used to define database connections via JNDI data sources.

The first wizard screen look like this.



Figure: Connection Wizard - Page 1

In the connection alias field, enter the name of the new database connection. This is the name that will be used for the connection in DbVisualizer, e.g., in the object tree.

Press Next to go to the next page.

On this page, select the driver you are going to use from the list. A red icon in front of the driver name indicates that the driver is not yet ready to use, while a green icon indicates that it has been properly configured (simply press Next to continue).

If the driver you select is not yet configured, a Load Driver File(s) button is displayed. When you click the Load Driver File(s) button, a file chooser is opened. You should select the JAR or ZIP file(s) that contain the driver implementation.



Figure: Connection Wizard - Page 2

In the file chooser, locate the files containing the JDBC driver files. (Select multiple files by pressing the SHIFT key while clicking).




Figure: Connection Wizard - Page 3

Once the driver has been properly loaded, a green icon appears in front of the driver name. Press Next to continue to the last page.

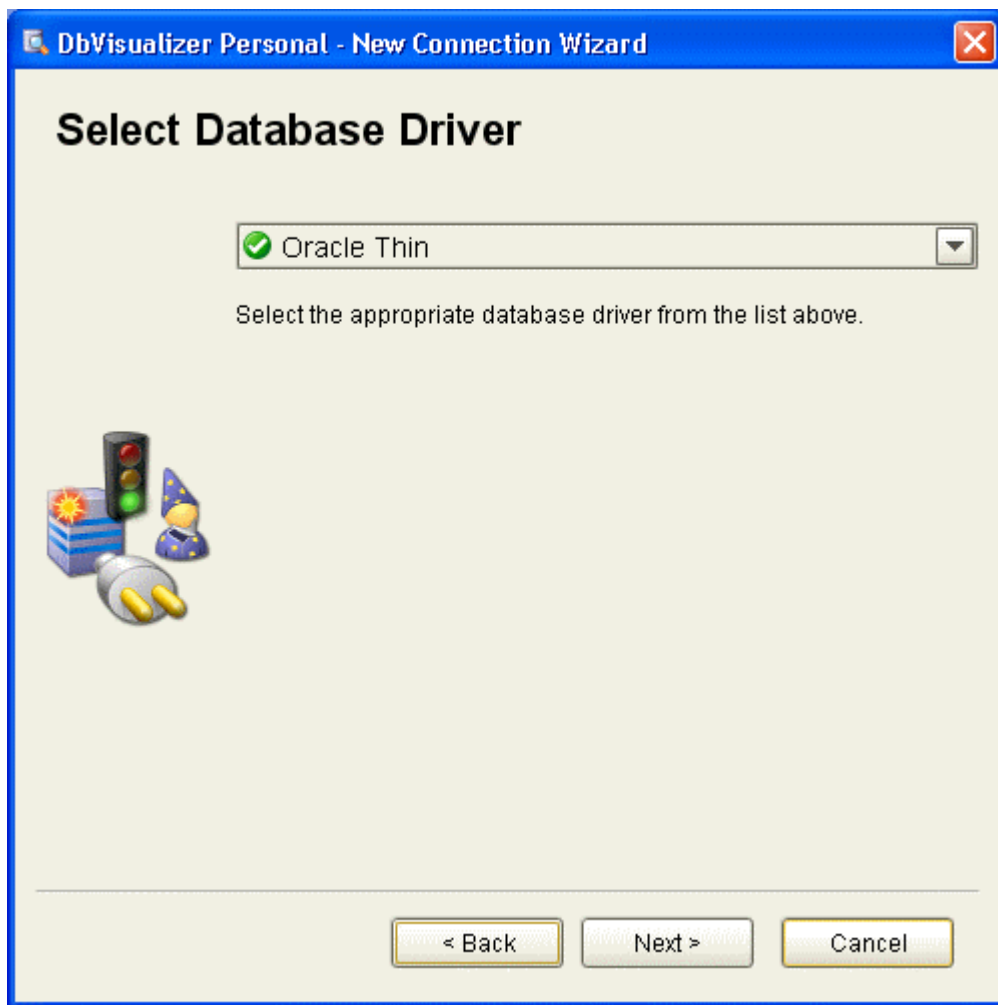


Figure: Connection Wizard - Page 4

On the last wizard pane, enter details for the new database connection. The information that must be provided varies depending on the database type. Please consult the database documentation if you are unsure about how to find the requested information.

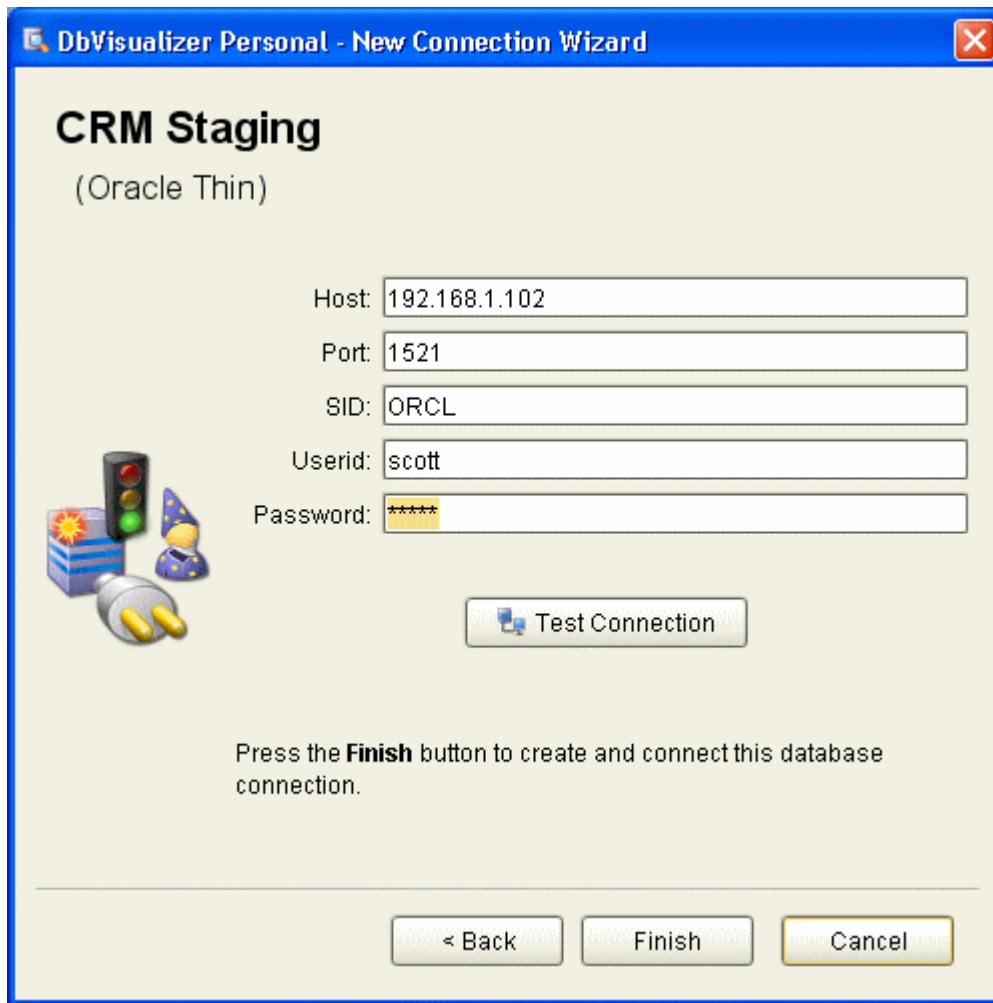


Figure: Connection Wizard - Page 5

Press Test Connection to check that the connection can be established. If the test passes, press Finish to create the new database connection and connect to the database.

We recommend that you skip the rest of this document, unless you:

- want to learn how the driver manager in DbVisualizer works
- need to have several versions of the same JDBC driver loaded simultaneously
- need to establish a connection via the JNDI interfaces (Java Naming and Directory Interface)
- need to add a Driver that do not exist in the wizard list of drivers

Driver Manager

The Driver Manager in DbVisualizer is used to define the drivers that will be used to communicate with the databases. You can manually locate the JDBC driver files and configure the driver, or you can use the JDBC Driver Finder to do most of the work for you, either on demand or automatically.

JDBC Driver Finder

The JDBC Driver Finder is a very powerful part of the Driver Manager that automates most of the driver management work. Given the folders where JDBC drivers are located, it loads and configures new drivers (if any) every time you start DbVisualizer. You can configure the JDBC Driver Finder in Tools Properties, in the General->Driver Manager category.

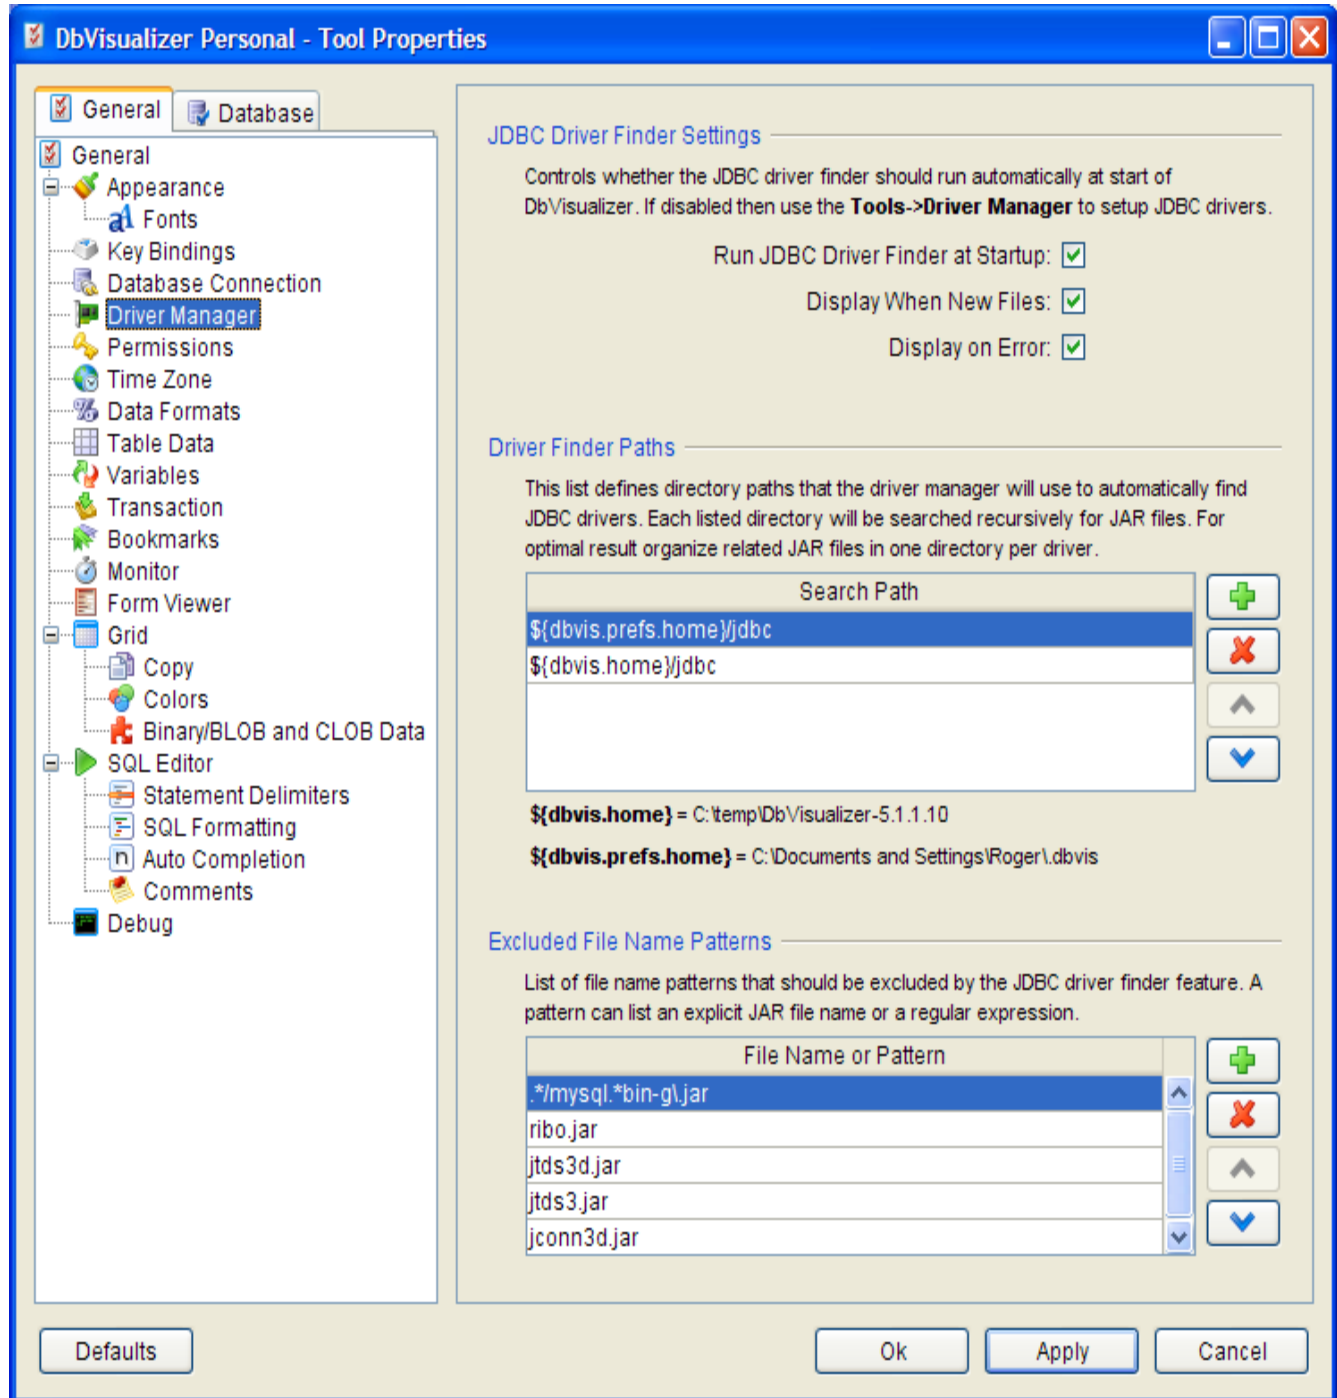


Figure: JDBC Driver Finder properties

Use the following properties to specify the finder behavior:

Load JDBC Driver and Get Connected

Property	Description
Run JDBC Driver Finder at Startup	If enabled, the finder will run automatically every time you start DbVisualizer. If it finds files, it will automatically load and configure them.
Display When New Files	If enabled, the finder window pops-up if it finds any new files when you start DbVisualizer. Otherwise, the finder runs invisibly in the background.
Display on Error	If enabled, the finder window pops-up if it encounters any errors loading and configuring drivers. Otherwise, it is silent about errors and you have to launch the Driver Manager to check for drivers not loaded successfully. Enabling this property is only meaningful if you have displayed the Display When New Files property.

You can also specify the folders the JDBC Driver Finder will search. By default, it will search folders named `jdbc` in the DbVisualizer installation directory (`${dbvis.home}`) and the DbVisualizer preferences folder (`${dbvis.prefs.home}`). These folder paths are shown under the list of Driver Finder Paths.

Finally, you can specify regular expression patterns for filenames that the finder should ignore. This can be useful if you need to store other files besides driver files in the designated folders.

If you let the JDBC Driver Finder load all drivers for you, all you need to do to install a new driver is to put the driver files in one of the folders specified for the finder in Tool Properties and then restart DbVisualizer.

Loading and Configuring Drivers Manually

You can also load and configure JDBC drivers manually using the Driver Manager. If you use JNDI to provide access to the database, you must use this option, since the JDBC Driver Finder does not handle JNDI. Start the Driver Manager dialog using the `Tools->Driver Manager` menu choice.

The left part of the driver manager dialog contains a list of driver names with a symbol indicating whether the driver has been configured or not. The right part displays the driver configuration for the selected driver in terms of the following:

- **Name**
A driver name in the scope of DbVisualizer is a logical name for either a JDBC driver or an Initial Context in JNDI. This is the name shown in the Connection tab setup when selecting which driver to use for a database connection
- **URL Format**
The URL format specifies the pattern for the JDBC URL or a JNDI Lookup name. Its purpose is to assist the user in the Connection tab when entering the URL or lookup name
- **Driver Class**
Defines the main class for the JDBC driver, used for connecting to the database.
- **Web Site**
Link to the DbVisualizer web site, where you can get up-to-date information about how to download the driver.
- **Driver File Paths**
Defines all paths to search for JDBC drivers or Initial Contexts when connecting to the database. The Driver File Paths area is composed of two tabs: the paths in the User Specified tab are used for dynamically loaded JDBC drivers or Initial Context classes, and the System Classpath tab lists all paths that are part of the Java system classpath.

Note: The System Classpath tab is only of interest for the [JDBC-ODBC](#) driver.

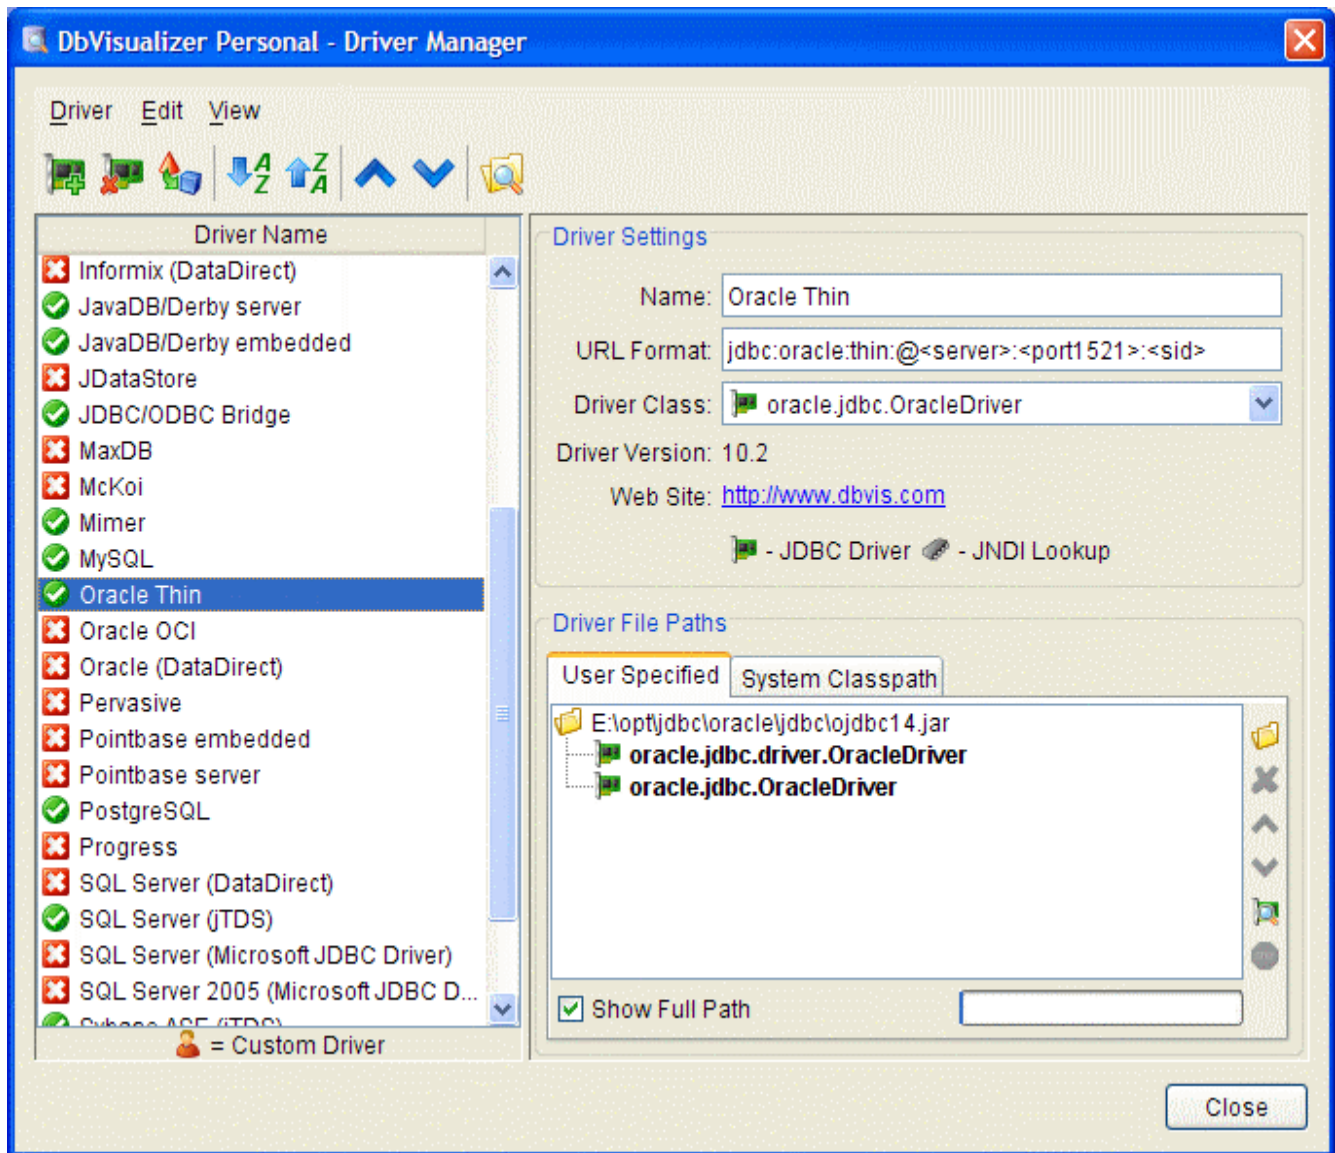


Figure: Driver Manager dialog

Initially, the driver list contains a collection of default drivers. They are not fully configured, as the paths to search for the classes need to be identified. You can edit the list, i.e., create, copy, remove and rename drivers. A driver is ready to use once a driver class has been identified, which is indicated with a green check icon in the list. Drivers that are not ready for use are indicated with a red cross icon.

Note: Only ready (configured) drivers appear in the Connection tab driver list.

The figure shows seven drivers that are ready: DB2 UDB, Informix, JDBC/ODBC Bridge, Mimer, MySQL, Oracle Thin and PostgreSQL.

Setup a JDBC driver

The recommended way to setup a driver is to pick a matching driver name from the list and then simply load the JAR, ZIP or directory that keeps the driver class(es). For instances, if you are going to load the JDBC driver for Oracle, select the Oracle driver in the list. You can also create a new driver or copy an existing one.

Note: Check the following online web page with the most current information about the tested databases and drivers:

Databases and JDBC Drivers

- It lists which databases and drivers we have tested
- Download links to JDBC drivers
- Information of which files to load in the driver manager for each JDBC driver
- Information of which Driver Class to choose

When you have selected the driver to configure, you need to load the driver files. Click the Load button to the right of the User Specified paths tree to show the file chooser and load the driver JAR, ZIP or individual files.




Figure: File Chooser dialog

A JDBC Driver implementation typically consists of several Java classes. If they are packaged in a JAR or a ZIP file, you don't have to worry about the details; just select and load the JAR or ZIP file. For instance, in the example above, use the ojdbc14.jar file.

If the driver classes are not packaged, it is important to select and load the root folder for the JDBC Driver. Java classes are typically organized using a package name structure. Example:

```
oracle.jdbc.driver.OracleDriver
```

Each package part in the name above (separated by ".") is represented by a folder in the file system. The root folder for the driver is the folder named by the first part, i.e., the oracle directory in this example. The class files are stored in the oracle/jdbc/driver sub folder. When the driver classes are located in a folder structure like this,

you must select and load the root folder, so that the Driver Manager gets the complete package structure.

When a connection is established in the Connection tab, DbVisualizer searches the selected drivers path tree's in the following order:

1. User Specified
2. System Classpath

The paths are searched from the top of the tree, i.e., if there are several identical classes in, for example, the dynamic tree, the topmost class will be used. Loading several paths containing different versions of the same driver in one driver definition is not recommended, even though it works (if you do this, you must move the driver you are going to use to the top of the tree). The preferred method for handling multiple versions of a driver is to create several driver definitions.

When you load files in the User Specified paths list, DbVisualizer analyzes each file to find the classes that represent main driver classes. Each such class is listed under the path where it was found in the User Specified paths lists, and it is also added to the Driver Class list in the Driver Settings area above. If there is more than one class in the list, make sure you select the correct Driver Class from the list. Consult the driver documentation (or the [Databases and JDBC Drivers](#) page) for information about which class to select.

JDBC drivers that requires several JAR or ZIP files

Some drivers depend on several ZIP or JAR files, or directories. An example is if you want XML support for an Oracle database. In addition to the standard JAR file for the driver, you then also need to load two additional JAR files. These are not JDBC driver files but adds functionality the driver needs to fully support XML.

Simply select all JARs at once and press Open in the file chooser dialog. The Driver Manager will then automatically analyze each of the loaded files and present any JDBC driver classes or JNDI initial context classes it finds.

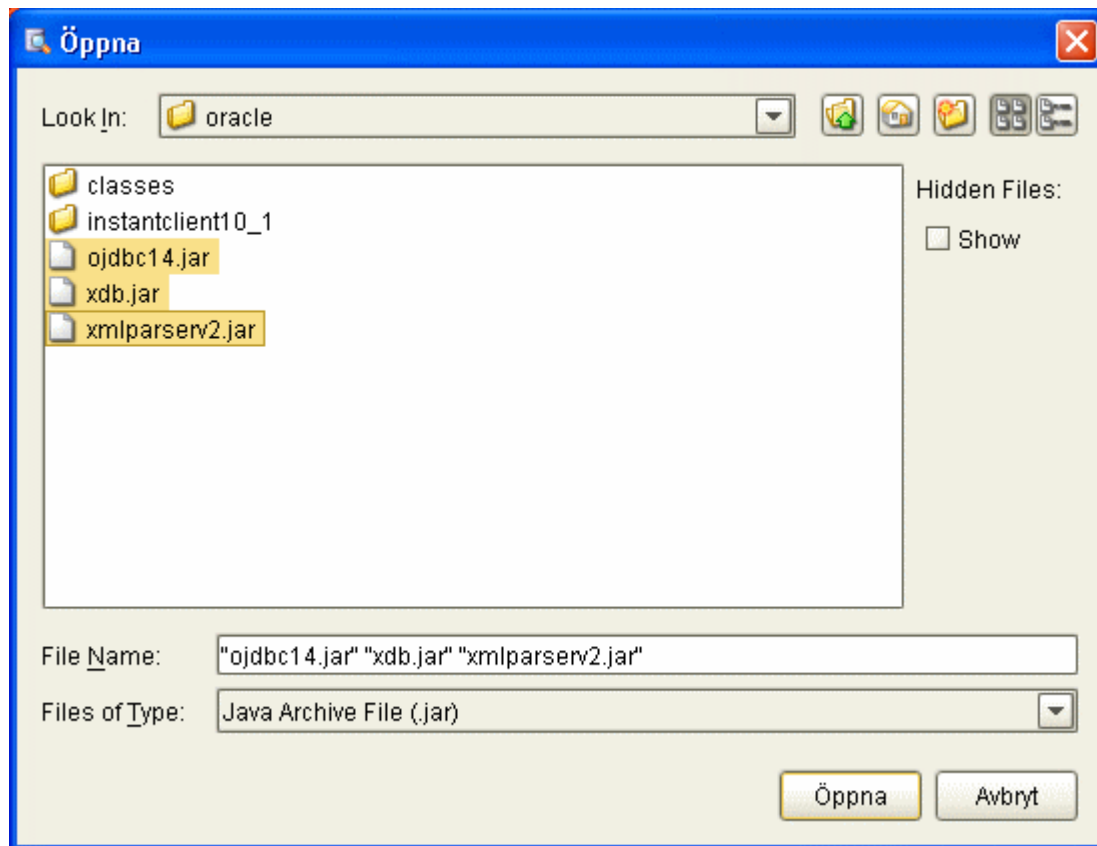


Figure: File Chooser dialog

The JDBC-ODBC bridge

The JDBC-ODBC driver is bundled with most Java installations, but not all (e.g., it is not included with Java for Mac OS X). The `JdbcOdbcDriver` class is included in a JAR file that is commonly named `rt.jar`, stored somewhere in the Java directory structure. DbVisualizer automatically identifies this JAR file in the System Classpath tree. To locate the `JdbcOdbcDriver`, simply press the Find Drivers button to the right of the System Classpath tree. When it is found, make sure the `sun.jdbc.odbc.JdbcOdbcDriver` is selected as the Driver Class in the Driver Settings area.

Loading JNDI Initial Contexts

Initial Context classes are needed to get a handle to a database connection that is registered with a JNDI lookup service. In DbVisualizer, these context classes are similar to JDBC driver classes in that an Initial Context implementation for a specific environment is required.

Note: Remember that the appropriate JDBC driver classes must be loaded into the Driver Manager even if the database connection is obtained using JNDI.

To load Initial Context classes into the Driver Manager, simply follow the steps outlined for loading JDBC drivers. The difference is that you will instead load paths containing Initial Context classes instead of JDBC drivers. When you load a path, DbVisualizer locates all Initial Context classes in the path and lists them in the User Specified paths list.

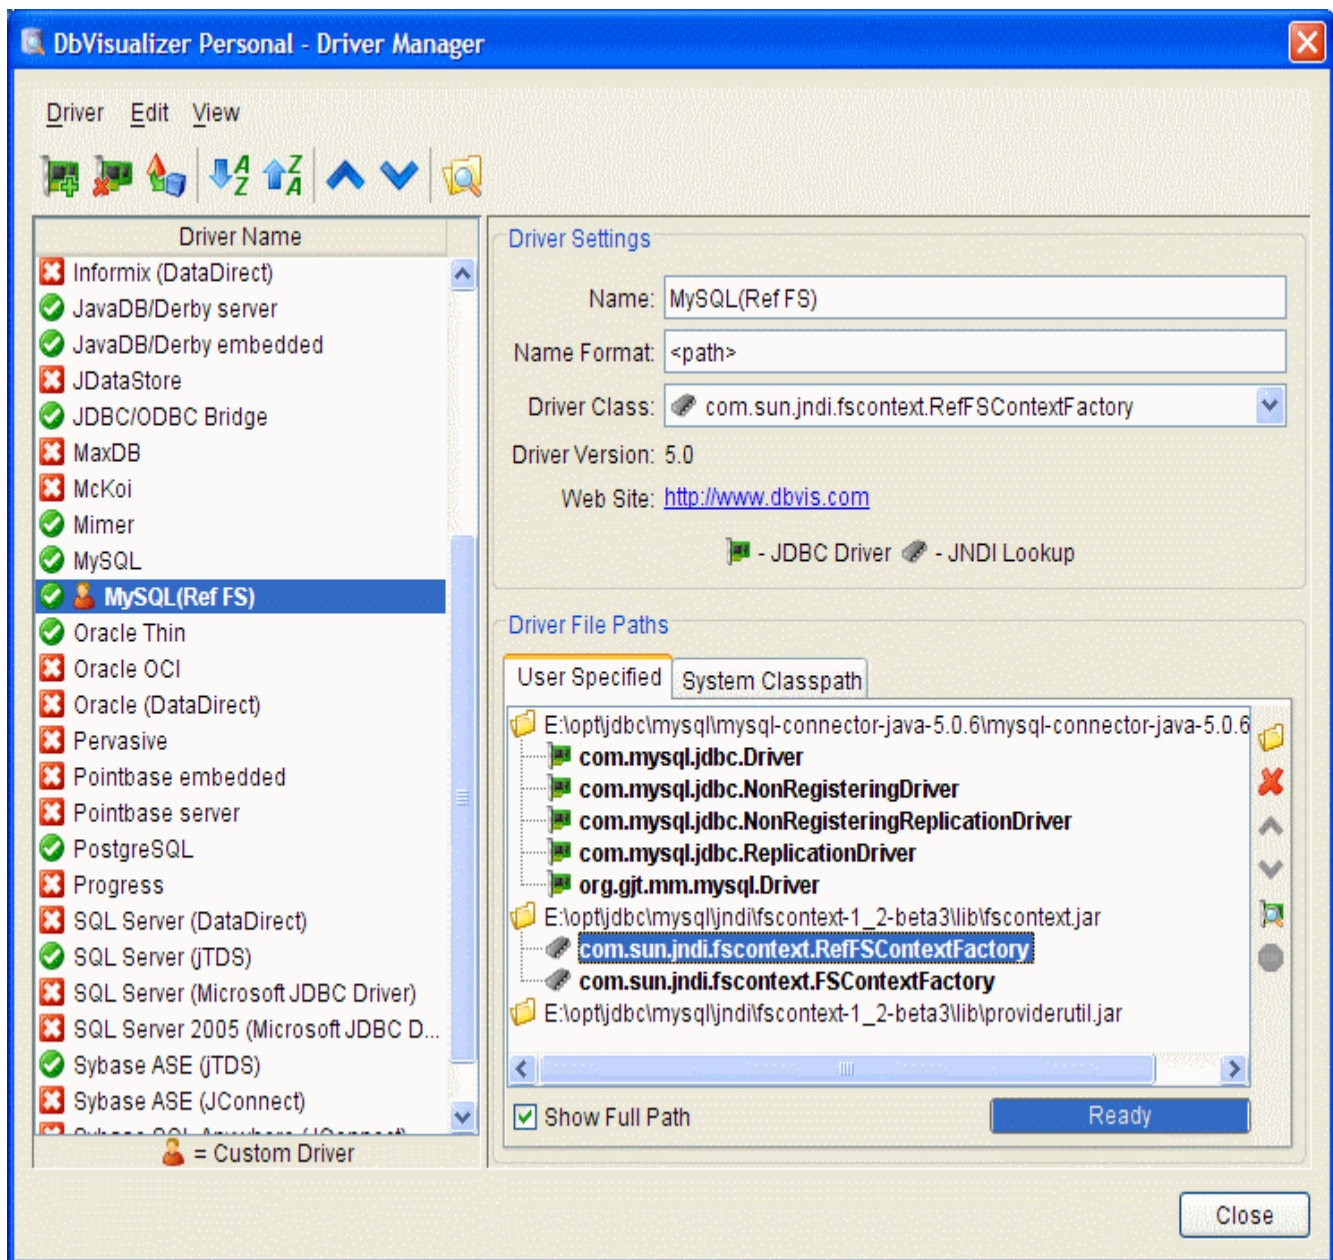


Figure: Driver Manager List with Initial Context classes

Visually, the difference between the identified JDBC drivers and Initial Context classes is the icon in the tree.

The figure shows the JAR files required to first obtain the JNDI handle, and then also the JDBC driver for the database. Check with the application server vendor or similar for more information of what files you need to load to get connected via JNDI.

Errors (why are some paths red?)

A path in red color indicates that the path is invalid. This may happen if the path has been removed or moved after it was loaded into the driver manager. Simply remove the erroneous path and locate the correct one.

Several versions of the same driver

The Driver Manager supports loading and using several versions of the same driver concurrently. We recommend that you create a unique driver definition per version of the driver and name the driver definitions properly, e.g., Oracle 9.2.0.1, Oracle 10.2.1.0.1, etc.

Setup a database connection

This section explains how to setup a Database Connection in the Connection tab.

Setup using JDBC driver

A Database Connection in DbVisualizer is the root of all communication with a specific database. It requires at a minimum that a driver is selected and that a Database URL is specified. A new Database Connection is created using the Database->Add Database Connection menu choice in the main window:

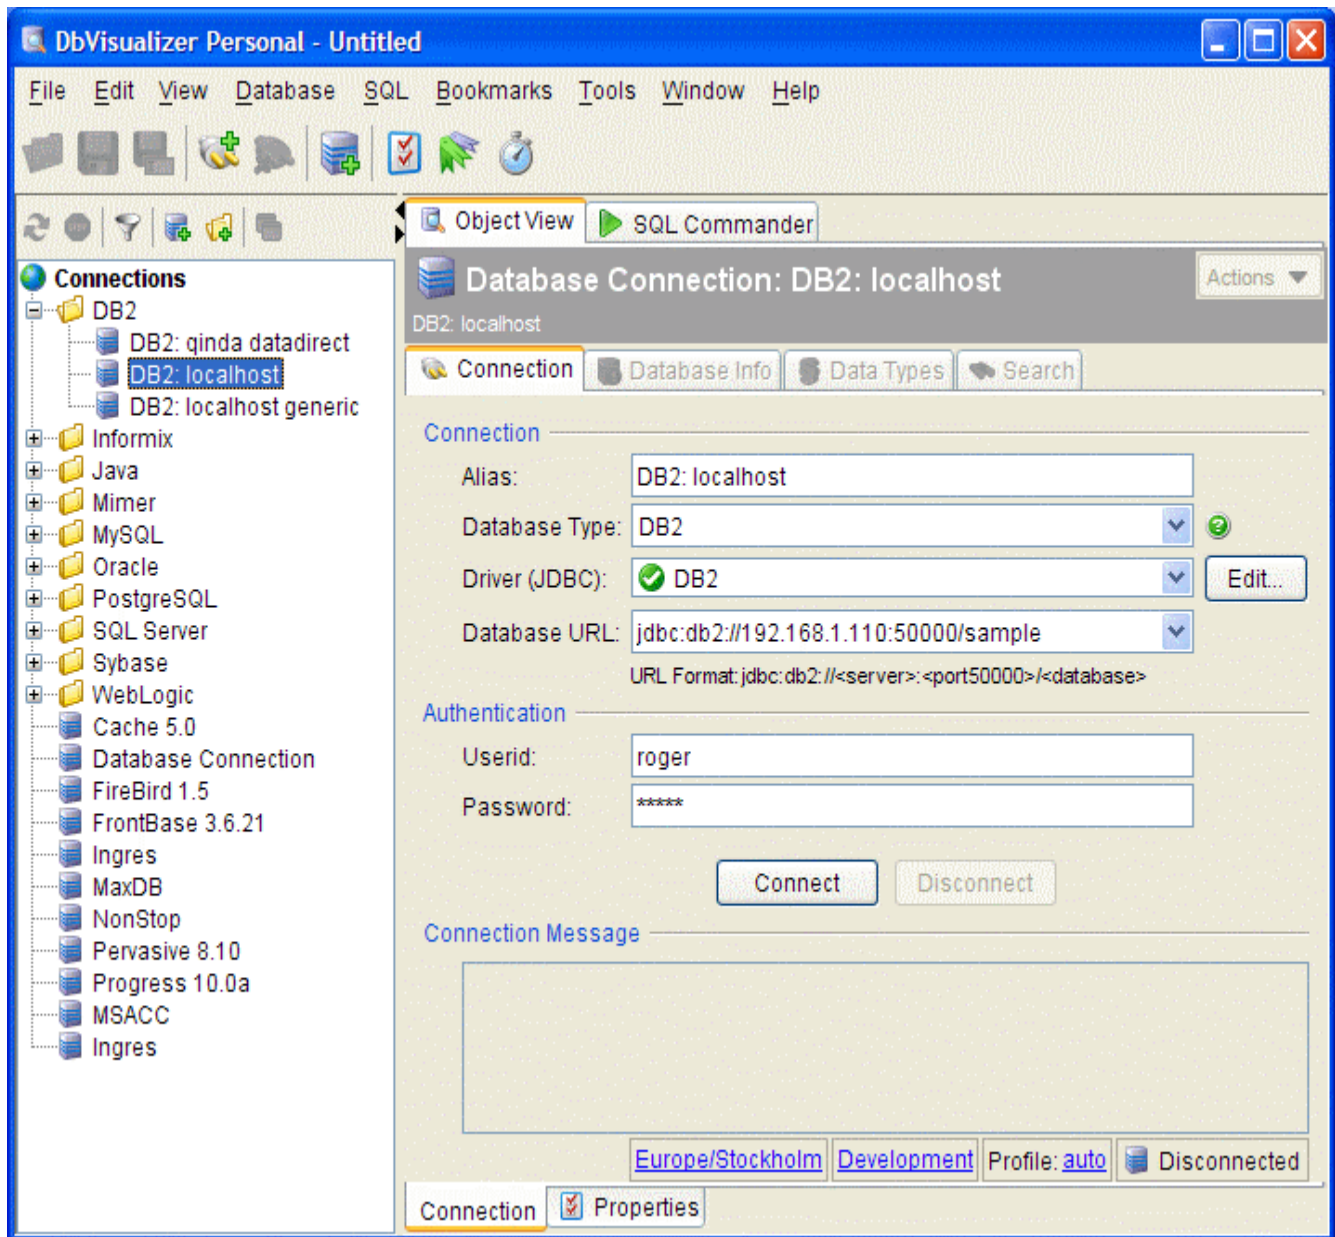


Figure: New Database Connection using JDBC driver

The Connection tab is the only tab that is enabled if you are not connected to the database. Database connection objects appear throughout the application and are by default listed by their URL. A URL can be, and often is, quite complex and long. You can use the Database Alias to set a more readable name for the database connection. The Driver list shows all defined drivers that have been defined properly in the Driver Manager. Just open the list and select the appropriate driver. The URL Format lists the format that the driver supports.

Tip: Put the mouse pointer on the URL Format and click with the mouse to copy the format template into the URL field.

Terms between < and > characters placeholders that need to be replaced with appropriate values, e.g.:

```
jdbc:oracle:thin:@proddb:1521:bookstore
jdbc:sybase:Tds:localhost:2638
jdbc:db2://localhost/crm
```

```
jdbc:microsoft:sqlserver://localhost;DatabaseName=customers
```

Userid and Password are optional but most databases require that they are specified.

Some drivers accept additional proprietary parameters described in the [Connection Properties](#) section.

Setup using JNDI lookup

The information needed to obtain a database connection using JNDI lookup is similar to what is needed for connecting using a JDBC driver.

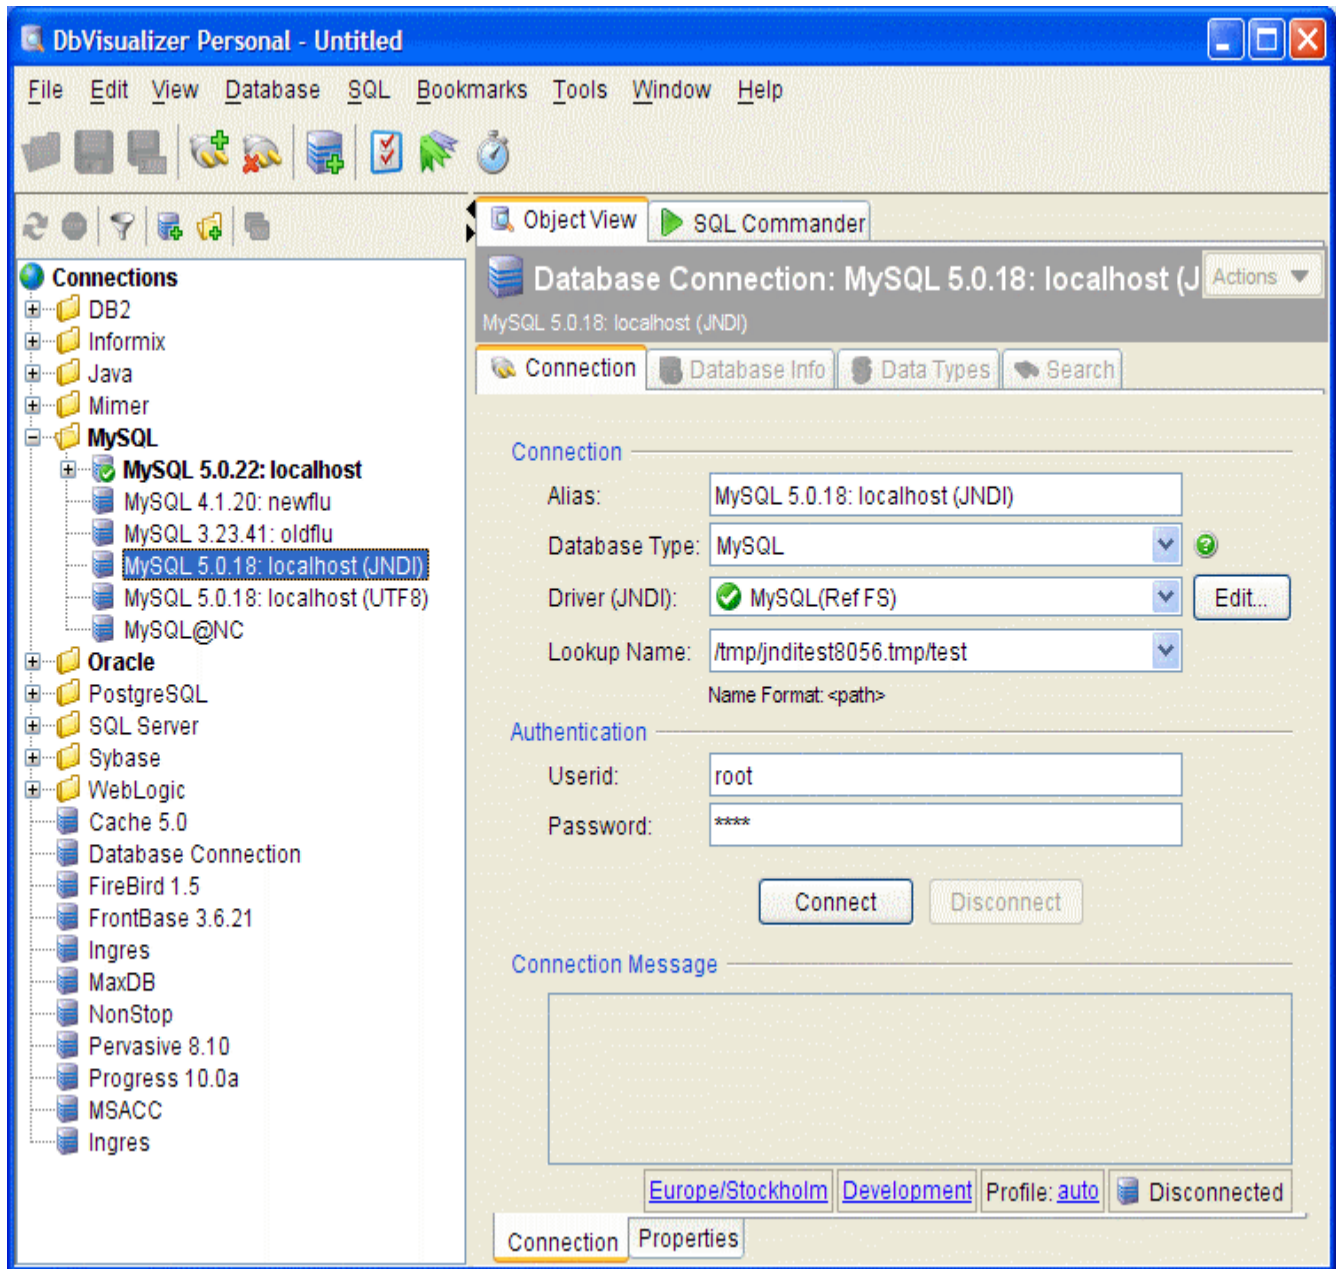


Figure: New Database Connection using JNDI lookup

The figure above shows parameters to connect with a lookup service via the MySQL RefFS driver. The /tmp/jnditest4975.tmp/test lookup name specifies a logical name for the database connection. This example is in its simplest form, since userid and password are not specified, nor where the database connection is finally fetched from. Any errors during the process of getting a handle to the database connection appears in the Connection Message area.

Connection Properties

In addition to the standard connection parameters (URL, Driver, Userid, Password, etc.), there are also a collection of connection properties. Which properties are available depends on the database type. Some database types have more properties than others. Which edition of DbVisualizer you use also affects which connection properties are available.

All supported database types (Oracle, Informix, Mimer, DB2, MySQL, etc.) are listed in the Database tab in the Tool Properties window. For each database type, there are a number of properties that are applied to any database connection of that type. This means, for instance, that a database connection defined as being a PostgreSQL database type will use the PostgreSQL properties defined in Tool Properties. The Connection Properties can then be used to override some settings specifically for one database connection. The advantage with this inheritance model is that property changes that apply to all connections can be made in one place, instead of having to apply a common setting for every database connection of a specific database type.

The following summarize the organization of the properties:

- Tool Properties (Database)
These apply to all database connections of the specific database type.
- Connection Properties
These apply for a specific database connection only.

- "Okay, so there are two places to change the value of a property. Which shall I use?"

This depends on whether the change should be applied to all database connections for a specific database type or just a single one. If the majority of your database connections should use the new property, it is recommended to set it in Tool Properties. Any overridden properties in the Connection Properties tab are indicated with an icon in the Properties tab label.

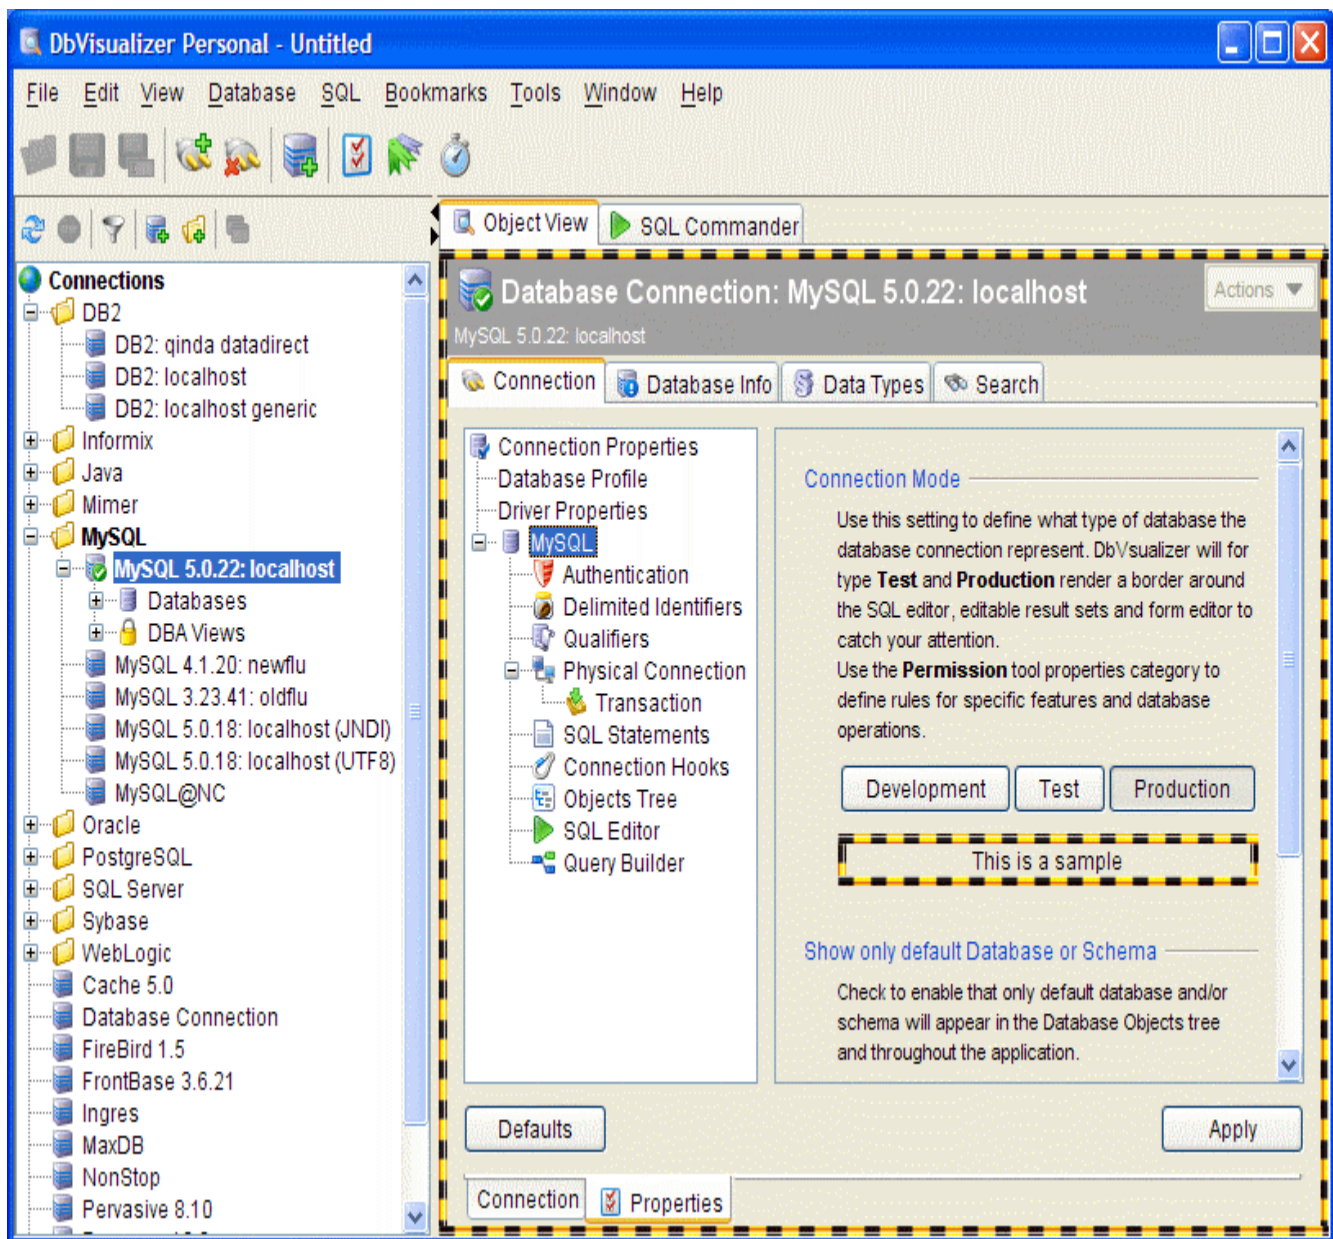


Figure: Connection Properties

The Connection Properties tab is organized in the same way as the Tool Properties window. The difference is that the list only includes the categories that are applicable for the selected database connection. Briefly, the categories are:

- Database Profile
- Driver Properties
- MySQL (The current Database Type)
 - Authentication
 - Delimited Identifiers
 - Qualifiers
 - Physical Connection
 - Transaction
 - SQL Statements
 - Connection Hooks

- Objects Tree
- SQL Editor
- Query Builder

The Database Profile and Driver Properties categories are only available in the Connection Properties tab and not in Tool Properties. The next section explains the Database Profile and Driver Properties categories, while the other categories are described in the [Tool Properties](#) document.

Additional categories may appear in the connection properties depending on the type of database. An example is the category for Explain Plan for Oracle, DB2 and SQL Server.

Database Profile

Please read in the [Database Objects Explorer](#) document for detailed information about database profiles.

The Database Profile category is used to select whether a profile should be automatically detected and loaded by DbVisualizer, or if a specific one should be used for the database connection. The default strategy is to Auto Detect a database profile.

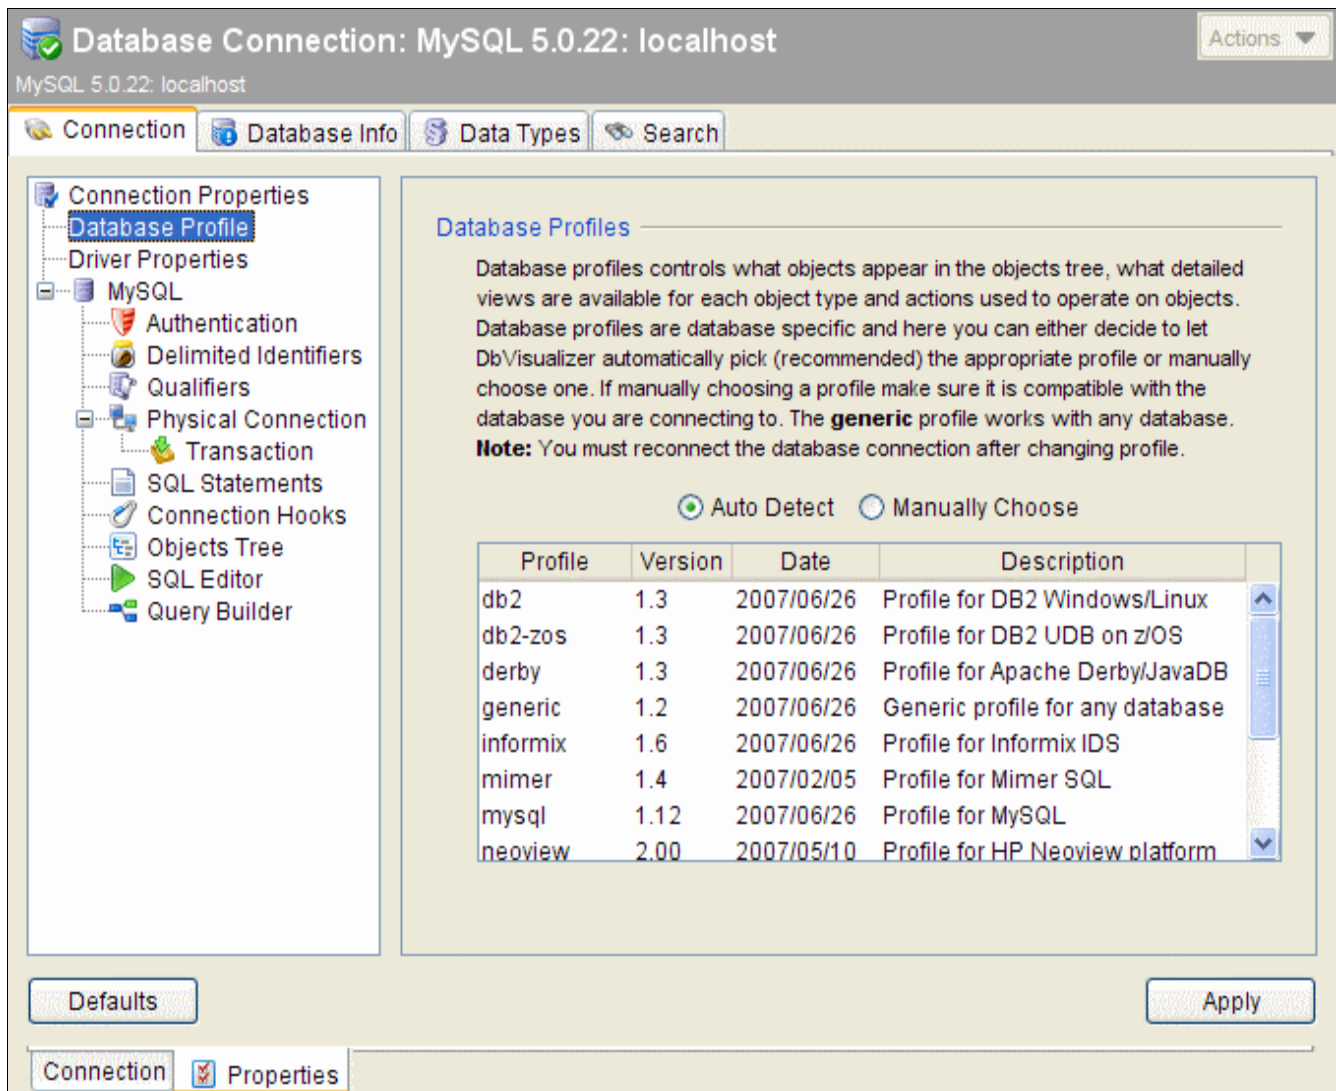


Figure: Database Profile category for a database connection

Note: The way DbVisualizer auto detects a profile is based on the setting of Database Type in the connection details.

If you manually choose a database profile, this choice will be saved between invocations of DbVisualizer.

Driver Properties

The Driver Properties category is used to fine tune a driver or Initial Context before the database connection is established.

Driver Properties for JDBC Driver

Some JDBC drivers support driver specific properties that are not covered in the JDBC specification.

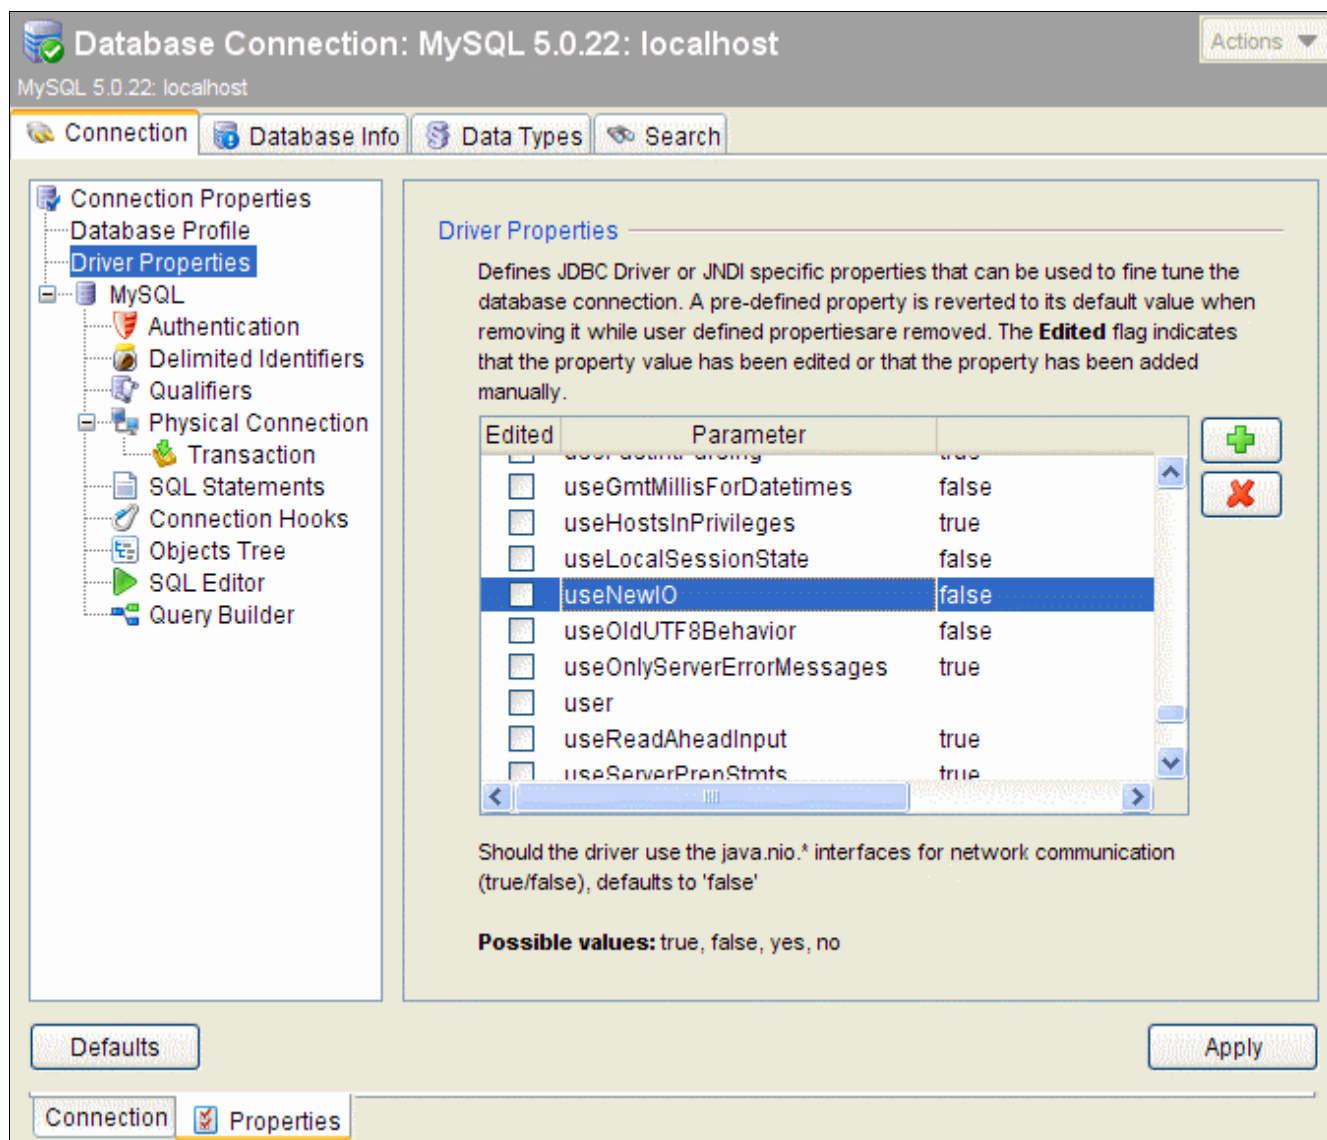


Figure: Driver Properties for JDBC Driver

The list of parameters, their default values and parameter descriptions are determined by the JDBC driver used for the connection. Not all drivers supports additional driver properties. To change a value, just modify it in the list. The first column in the list indicates whether the property has been modified or not, and so, whether DbVisualizer will pass that parameter and value onto the driver at connect time. New parameters can be added using the buttons at the bottom of the dialog. Be aware that additional parameters do not necessarily mean that the driver will do anything with them.

Driver Properties for JNDI Lookup

The Driver Properties category for a JNDI Lookup connection always contains the same parameters.

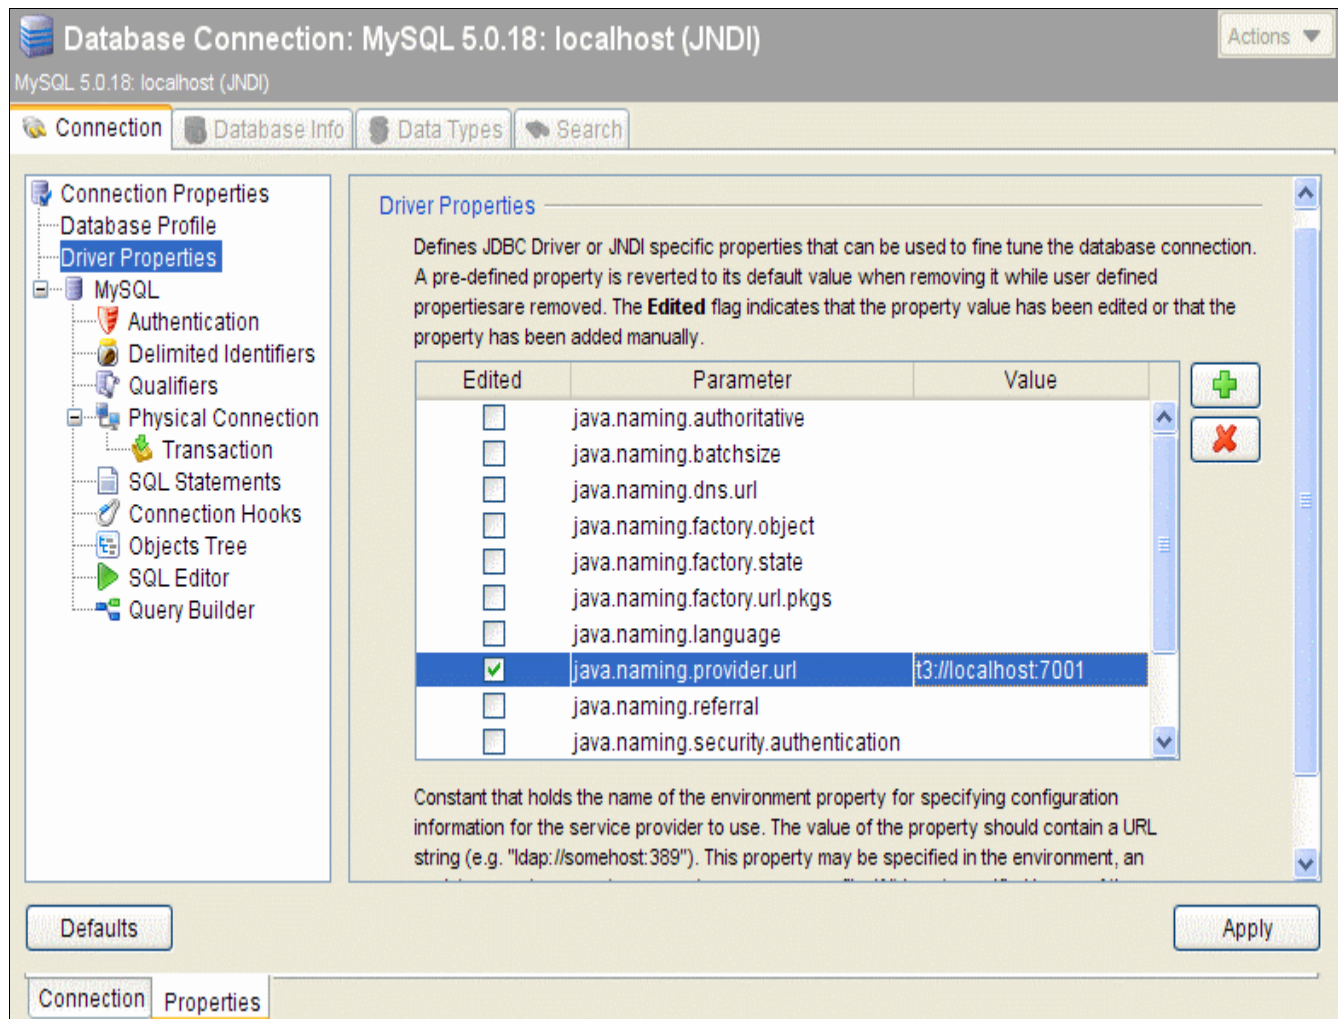


Figure: Driver Properties for JNDI lookup

The list of options for JNDI lookup is determined by the constants in the javax.naming.Context class. To change a value, just modify the value of the parameter. The first column in the list indicates whether the property has been modified or not, and so, whether DbVisualizer will pass that parameter and value onto the driver at connect time. New parameters can be added using the buttons at the bottom of the dialog. Be aware that additional parameters do not necessarily mean that the InitialContext class will do anything with them.

Always ask for userid and/or password

Userid and password information is generally information that should be handled with great care. By default, DbVisualizer saves both userid and password (encrypted) for each database connection. Userid is always saved while password saving can be disabled in the connection properties.

The Require Userid and Require Password connection properties can be enabled to tell DbVisualizer to automatically prompt for userid and/or password when a connection is to be established. Enabling either one or both of these while leaving the Userid and Password fields blank for a database connection ensures that DbVisualizer will not keep this vital information between sessions. The following dialog is displayed if requiring both userid and password.



Figure: Dialog asking for Userid and Password as a result of having Require Userid and Password settings enabled

Using variables in the Connection details

Variables can be used in any of the fields in the Connection tab. This can be useful alternative to having a lot of similar database connection objects. Several variables can be in a single field, and default values can be set for each variable. The following figure shows an example with variables, i.e., variable named delimited by dollar characters, \$\$. . . \$\$.

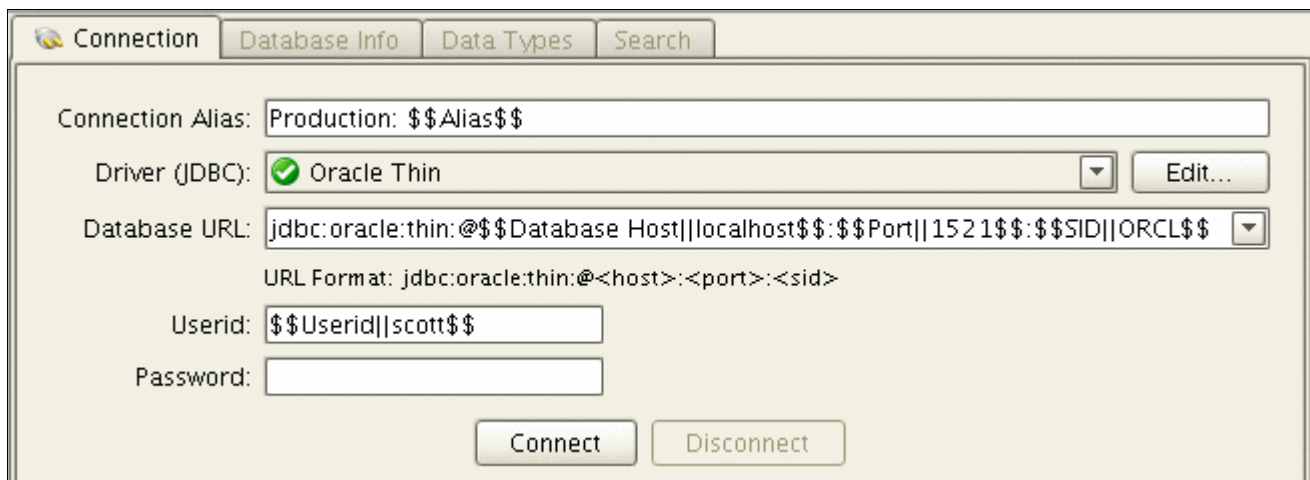


Figure: Connection tab with variables

The following variables appear in the figure:

- `$$Alias$$`
- `$$Database Host||dbhost2|||choices=[dbhost1,dbhost2,dbhost3] $$`
- `$$Port||1521$$`
- `$$SID||ORCL$$`
- `$$Userid||scott$$`

All of these variables define a default value after the "||" delimiter, except for the `$$Alias$$` variable, which have no default value. The default values appear in the connect dialog when you ask for a connection to be established. The `$$Database Host$$` variable includes the `choices` option, with a comma separated list of choices that should appear in a drop-down list. The drop-down list is editable, so the user is not locked into the choices from the list.

The following figure shows the connect dialog based on the connection definition shown above.

Note: Using variables in conjunction with the `Require Userid` and/or `Require Password` settings also works.

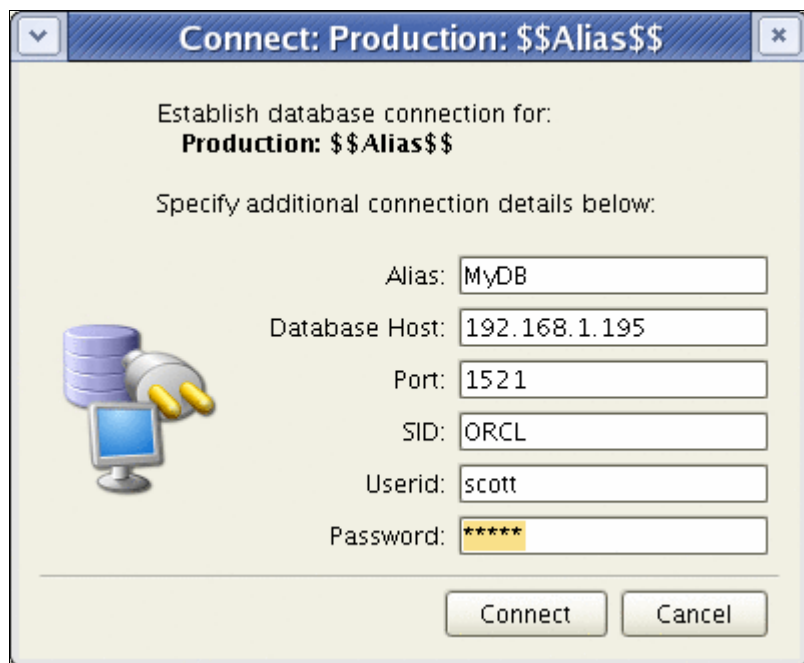


Figure: Connection tab with variables

Enter the appropriate information in the fields and then press the `Connect` button to establish the connection. When the connection is established, DbVisualizer automatically substitutes the variables in the `Connection` tab with the values entered in the connect dialog. At disconnect from the database, they revert back to the original variable definitions.

Connect to the Database

Press `Connect` when all information has been specified. DbVisualizer passes all information you entered on to the selected driver, and when the connection is established, the following appears.

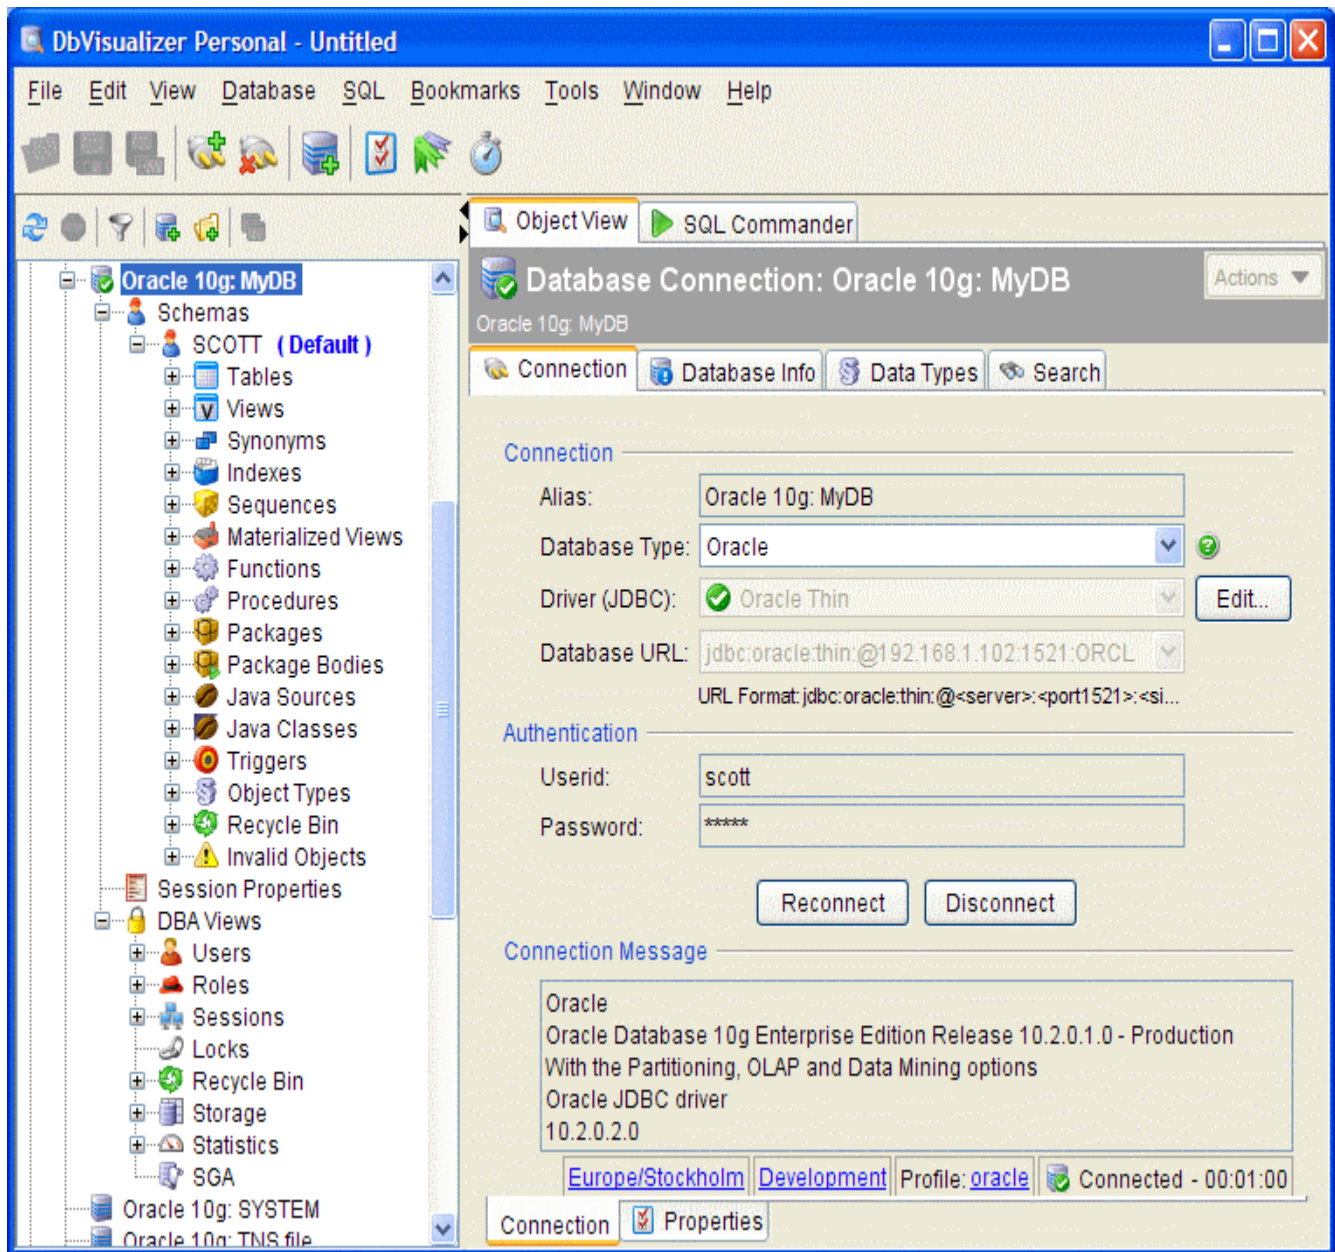


Figure: A freshly initiated database connection using JDBC driver

The Connection Message now lists the name and version of the database as well as the name and version of the JDBC driver. The database connection node in the tree indicates that it is connected. The connection properties cannot be edited while a database connection is established. The Alias can be edited by selecting the database connection node in the tree and then clicking on the name.

The figure above also shows that the database connection node in the tree has been expanded to show its child objects.

If the connection is unsuccessful, it is indicated by an error icon in the tree. The error message as reported by the database or the driver appears in the Connection Message area. Use this information to track down the actual problem. Since these conditions are specific for the combination of driver and database, you should check the driver and database documentation to find out more. Below are a few common problem situations:

Error Message	Explanation
<p>No suitable driver. There is no driver that can handle a connection for the specified URL. The most common reason is that the driver is not loaded in the Driver Manager. Also make sure the URL is correct spelled.</p>	<p>The JDBC support in Java determines what driver to load based on the database URL. If the URL is malformed then there might be no driver that is able to handle the database connection based on that URL. This error is produced when this situation occurs or when the driver is not loaded in the driver manager. The recommendation is to check the JDBC driver documentation for the correct syntax.</p>
<pre>java.sql.SQLException: Io exception: Invalid number format for port number Io exception: Invalid number format for port number</pre>	<p>The URL templates that are available in the Database URL list contains the "<" and ">" place holders. These are there to indicate that the value between them must be replaced with an appropriate value. The "<" and ">" characters must then be removed.</p> <p>This example error message is produced by the Oracle driver when using the following URL:</p> <pre>jdbc:oracle:thin:@<qinda>:<1521>:<fuji></pre> <p>Simply remove the "<" and ">" characters and try again.</p>

Connections Overview

The Connections overview is displayed by selecting the Connections object in the Database Objects Tree. This overview displays all database connections in a list and is handy to get a quick overview of all connections. In addition to the URL, driver, etc. there are a few symbols describing the state of each connection. Double clicking on a connection changes the display to show that specific connection.

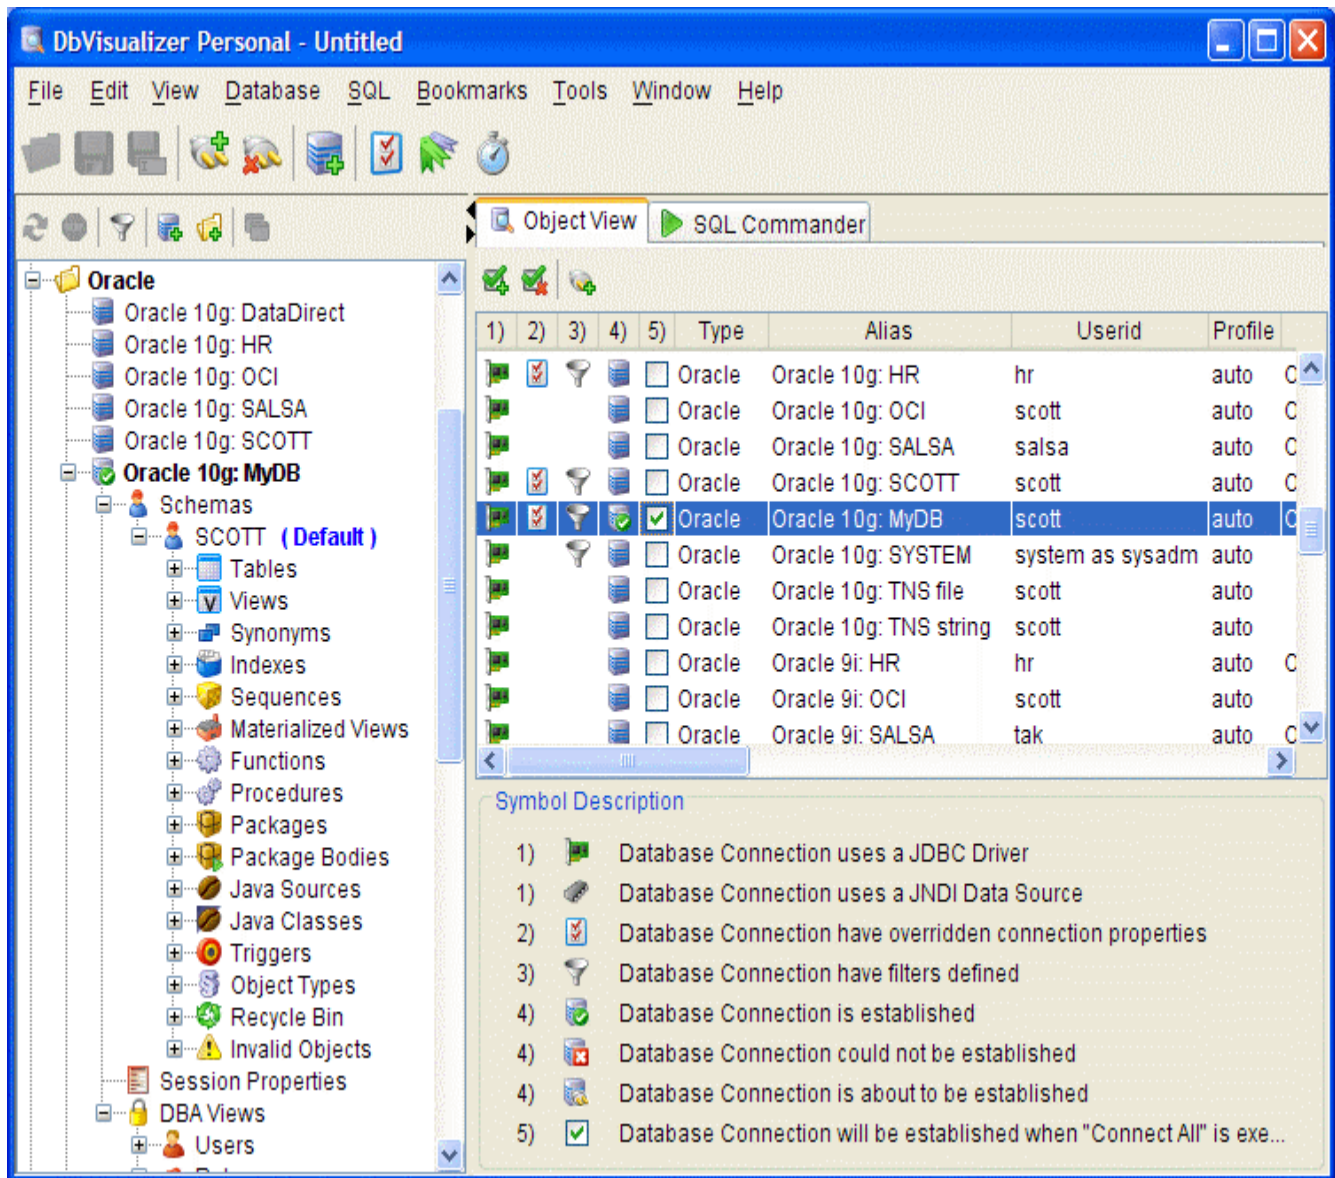


Figure: The Connections Overview

Information for each symbol is provided in the description area below the list. The fifth check symbol is the only editable symbol and is used to set the state of the Connect All property, i.e., whether the database connection should be connected when selecting the Database->Connect All menu choice.

Database Objects Explorer

Introduction

The Database Objects Tree is used to explore databases and browse details about objects. Which objects that may be explored and which object actions that exist is database dependent.

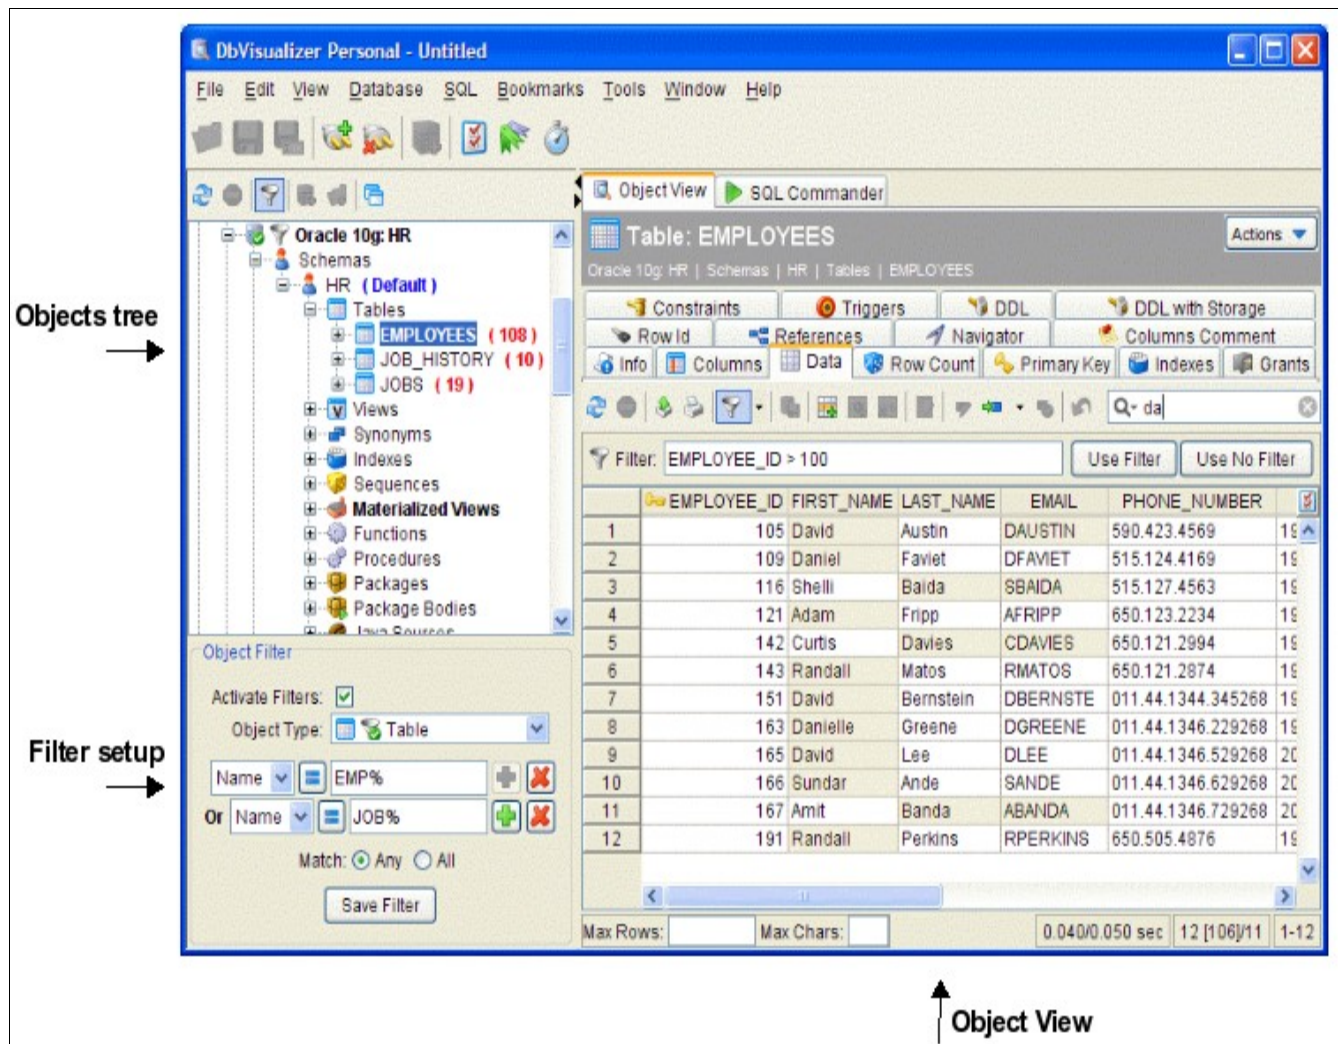


Figure: Database Objects tab

The Database Objects Tree to the left is the place to setup new database connections and establish connections. Once connected, expand the database connection object and explore the objects available. The Object View area to the right displays detailed information about the currently selected object in the tree.

The Filter setup pane below the tree is used to control what objects are displayed in the tree. It comes in handy when you have many schemas or tables in your database and want to limit the number of visible objects.

For some object types, there are *actions* (small dialogs for performing a task) for common operations, such as creating, altering, and dropping database object. Which actions are available depends on the database you are

connected to and the database profile used for the connection. More about this in sections below.

Tip 1: The Database Objects Tree is always visible to the left. If the currently selected main tab is the SQL Commander, you can double click on an object in the tree to automatically switch to the Object View tab.

Tip 2: All object names in the tree can be dragged to any editable text fields, including to the SQL Commander editor.

Create a Database Connection

There are a few objects that always appear in the tree independent of the edition of DbVisualizer and the database profile in use. The most important object is the Database Connection, which is used to setup and establish a database connection. The other two objects are Folder and Connections Overview. The following sections describe these objects in more detail.

Database Connection object

The Database Connection object is the root object for a connection. Before exploring or accessing a database, you need to establish the connection. Create a new database connection using the Database->Add Database Connection main menu choice and the following will appear.

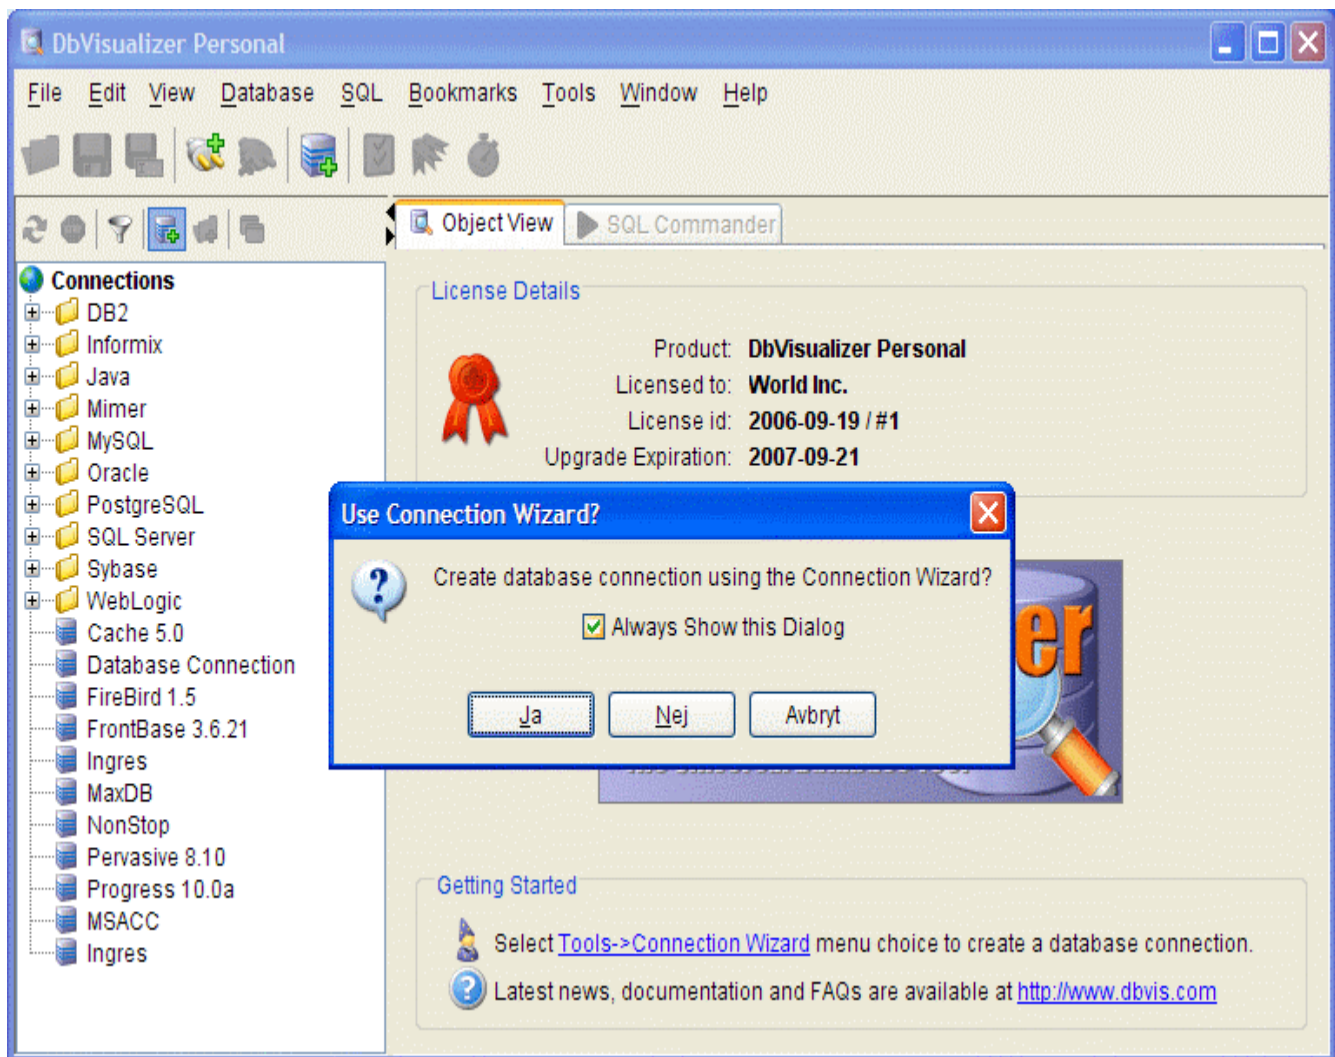


Figure: Add database connection

We recommend that you always use the connection wizard when you create a new database connection, as it hides the complexity of loading drivers and the syntax of database URLs (detailed information on how to establish a connection is provided in the [Load JDBC Driver and Get Connected](#) document).

Tip 1: Once a database connection has been setup properly, you just need to double click on the object to establish the connection.

Tip 2: You can use the Database->Connect All main menu choice to connect all enabled database connections with a single click. You make a database connection "Connect All" -aware in the Database Properties or in the Connections overview.

Alias

The name of the database connection object as it appears in the tree is by default "Database Connection". The Connection Alias can be used to provide a name that is more descriptive. Enter the new name in the Alias field in the Connection sub tab or click on the name in the tree and start editing the name.

Default database and schema

The (Default) indicator after the name of a database or schema in the tree indicates that it is the default database or schema. The default is determined when you connect to the database.

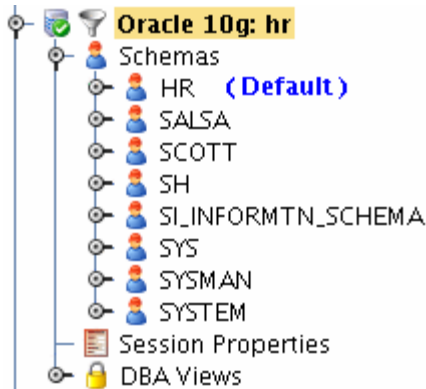


Figure: The (default) indicator for database and schema objects

Tip: In the Connection Properties you can define that only the default database or schema should be visible in the tree.

Remove and copy database connection objects

To remove a database connection, select the Database->Remove Database Connection operation in the main menu. You can copy a database connection with Database->Duplicate Database Connection.

Database Connection detailed information

The following section describes the tabs in the Objects View for a database connection briefly.

Tab	Description
Connection	This tab is always enabled and is used to setup the details for a database connection. You can also connect, disconnect and reconnect using the buttons in this tab.
Database Info	When connected, the Database Info tab shows various information supplied by the driver. Much of this information is low level, even though some of it may be useful.
Data Types	The Data Types tab lists all data types supported by the database.
Search	The Search Tab is used to search among the objects in the tree. Search operates on the content in the tree. See the next section for more information about search.

Search

The Search tab is used to search among the objects in the tree by object name. Note that if you have tree filters or any other property that limits the content of the tree enabled, the search is performed only for those objects that match the filters. The types of objects that are searchable depends on the database you are connected to. For instance, columns are included in the tree for some databases but not for others.

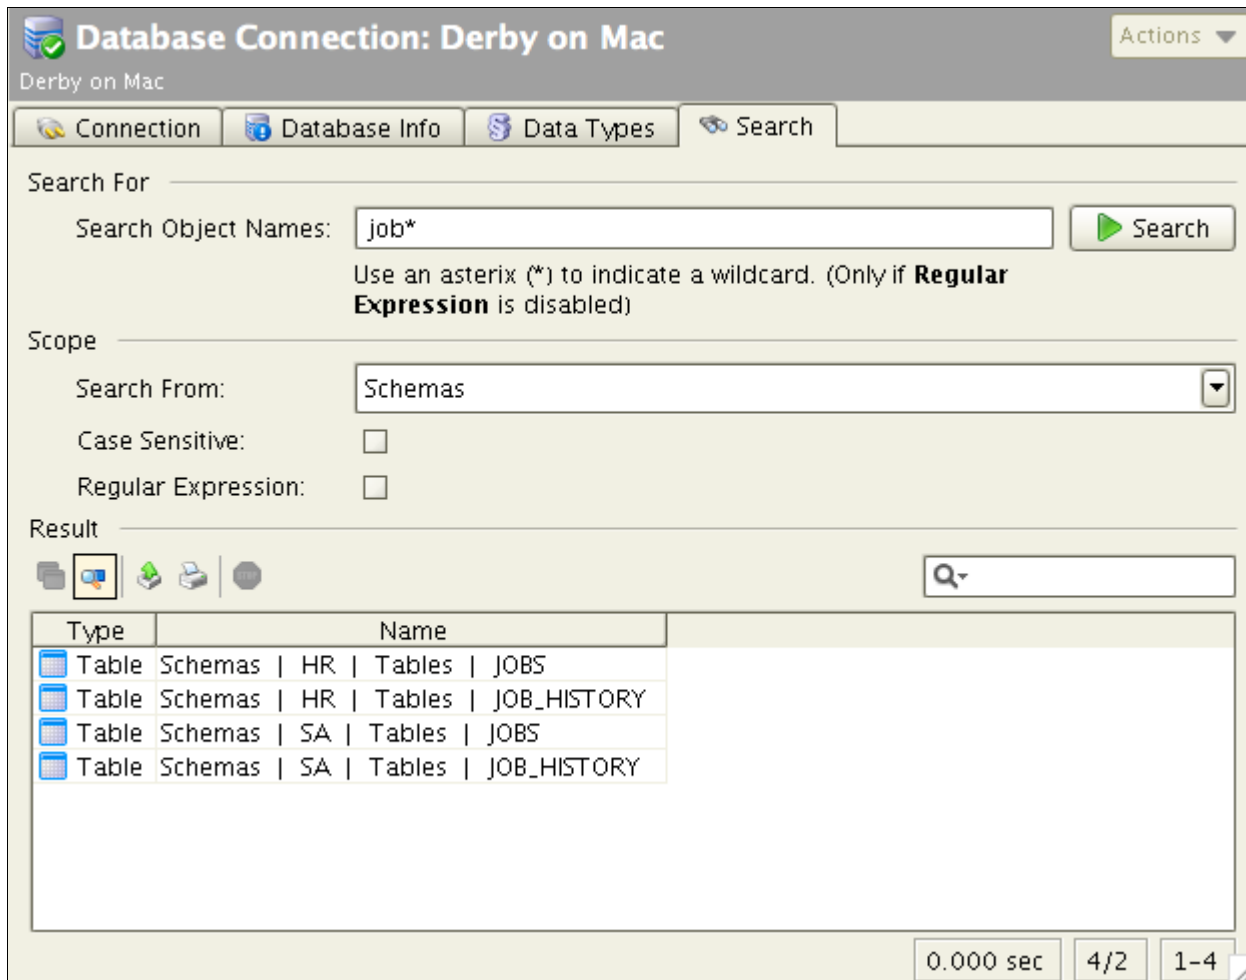


Figure: The Search tab

Search by specifying the name of the object, or name pattern, and press the Search button. You can use asterisk (*) as a wildcard in a pattern, or you can use a regular expression pattern if you enable it by checking the Regular Expression checkbox. You can also specify where in the tree to start the search, and whether to do a case sensitive search.

You can search operation with Stop button in the grid toolbar. Use the Show Object Path toolbar toggle button to include or exclude a column for the complete path for each found object in the grid. This path is the same as if navigating to each object manually in the objects tree. Other grid toolbar buttons let you export and print the search result grid.

Note: The search may take some time to perform the first time since all objects defined in the actual database profile are examined.

Tip: Double click on a row to see detailed information about a specific object. This will display all information about the object in a separate window.

Organizing Database Connections in Folders

If you work with many database connections, you can use folder object to organize and group them in the tree. Folder objects can have child folder objects in an unlimited hierarchy. Use the Database->Create Folder and Database->Remove Folder menu choices to create and remove folder objects. You can use the Database->Move Up/Down main menu choices to move the folders (and database connections) in the tree, or you can just drag

and drop to the nodes to a new location.

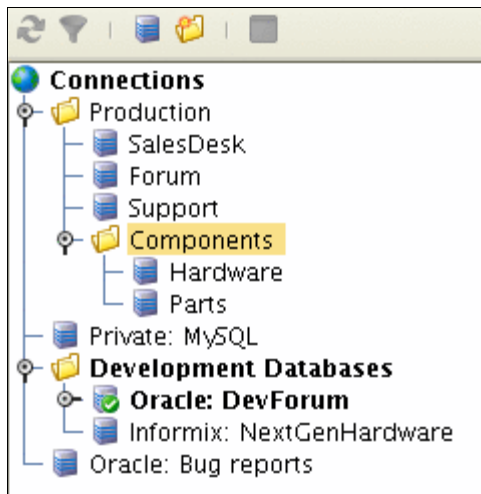


Figure: The database objects tree and the folder object type

Connections overview

The Connections object is the root object in the tree and acts as a holder for all database connections and folders. When selected, it displays an overview of all database connections in the Object Details view. Here you can see the basic settings and states for your database connections. For more information, see the [Load JDBC Driver and Get Connected](#) chapter.




Figure: Connections object

Database Objects Tree

Standard Actions

The Database Objects toolbar buttons are used to do tree related operations. These are individually enabled or disabled based on the currently selected object.



Figure: Objects tree toolbar

Description of the buttons from the left:

Tool bar button	Description
Reload	Reloads the currently selected object by asking for new information about the object from the database. Objects have been created or removed.
Stop	Stops the current tree operation, for instance connecting to a database or expanding a node.
Show/Hide Tree Filter	Is a toggle button that determines whether the Filter management pane will be displayed below the tree.
Create Database Connection	Adds a new Database Connection object in the tree. The location of the new object is determined by the current selection. If no node is selected, the new object is added at the end of the list.
Create Folder	Creates a new folder object.
Show in Window	Request to display the details view for the selected object in a separate window.

The right click menu for an object and the Database main menu lists object specific actions. The following actions are always available for all objects:



Figure: Standard right click menu actions for all objects

Object Actions

An object in the objects tree may have object specific actions attached to it. These actions are accessible via any of:

- Right-click menu in the objects tree
- Via the Database->Selected Object main menu
- Via the Actions menu button in the object view

Here is an example of the actions menu launched via the Actions menu button:

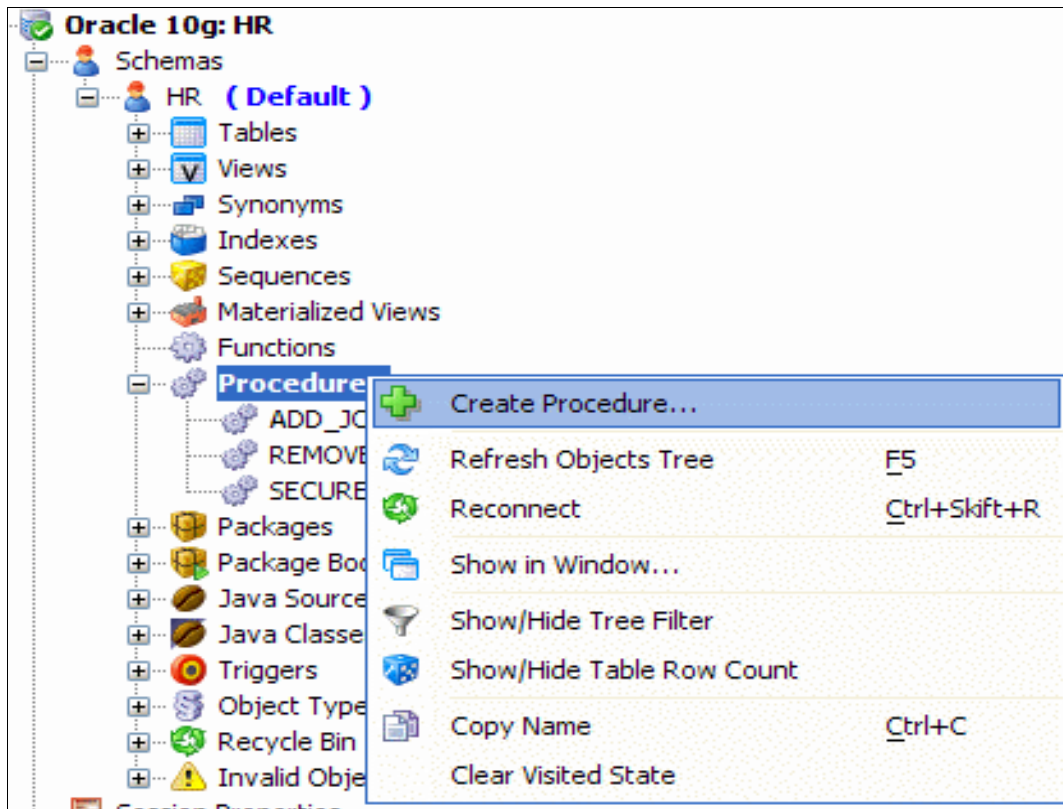


Figure: Object actions menu

Common Object Actions

There are a few actions that appear for some object types in all database profiles. These are most often valid for plain table object types and offer related functionality. Read the following sections for more information.

Create Table

The Create Table action launches the Create Table feature. You use it to create a table, optionally with a primary key, foreign keys and other constraints. Read more about this feature in [Create and Alter Table](#).

Create Index

The Create Index action launches the Create Index assistant dialog, where you can select columns to include in a new index for a table. See the [Create and Alter Table](#) for more information.

Import Table Data

Import Table Data launches a dialog where you can specify a CSV file to be imported into a table. Various configurations for how the source file is organized and data mapping are offered. Read more in [Export, Import and Print](#).

Script Object to SQL Editor

Use this action to create pre-defined SQL statements based on the source table and its columns. The created statement is copied to the current SQL editor in the SQL Commander. Here are a couple of examples:

Script Object to SQL Editor -> Select

```
SELECT
COUNTRY_ID,
COUNTRY_NAME,
REGION_ID
FROM
HR.COUNTRIES
```

Script Object to SQL Editor -> Insert

```
INSERT
INTO
  HR.COUNTRIES
  (
    COUNTRY_ID,
    COUNTRY_NAME,
    REGION_ID
  )
VALUES
  (
    ' ',
    ' ',
    0
  )
```

For databases with DbVisualizer [database specific](#) profiles, the Script Object to SQL Editor action menu also contains an entry for generating the DDL for Table and View objects.

Script Object to New SQL Editor

This is the same as Script Object to SQL Editor, except that the SQL is copied to a new SQL editor instead of to the current editor.

Objects Tree Filtering

The Filtering setup is activated via the Database->Show/Hide Tree Filter menu choice and the filter pane appear below the objects tree. Filtering is useful to limit the number of objects that will appear in the tree.

Tree filters are managed per database connection object. What can be filtered is defined per database profile. The generic database profile supports filtering of database (catalog), schema, table and procedure names.

The unfiltered schema objects for an Oracle connection.

The same objects but now filtered based on all schema names starting with "O" or "S".

Filter defined as all names that do not start with "O" and "S".

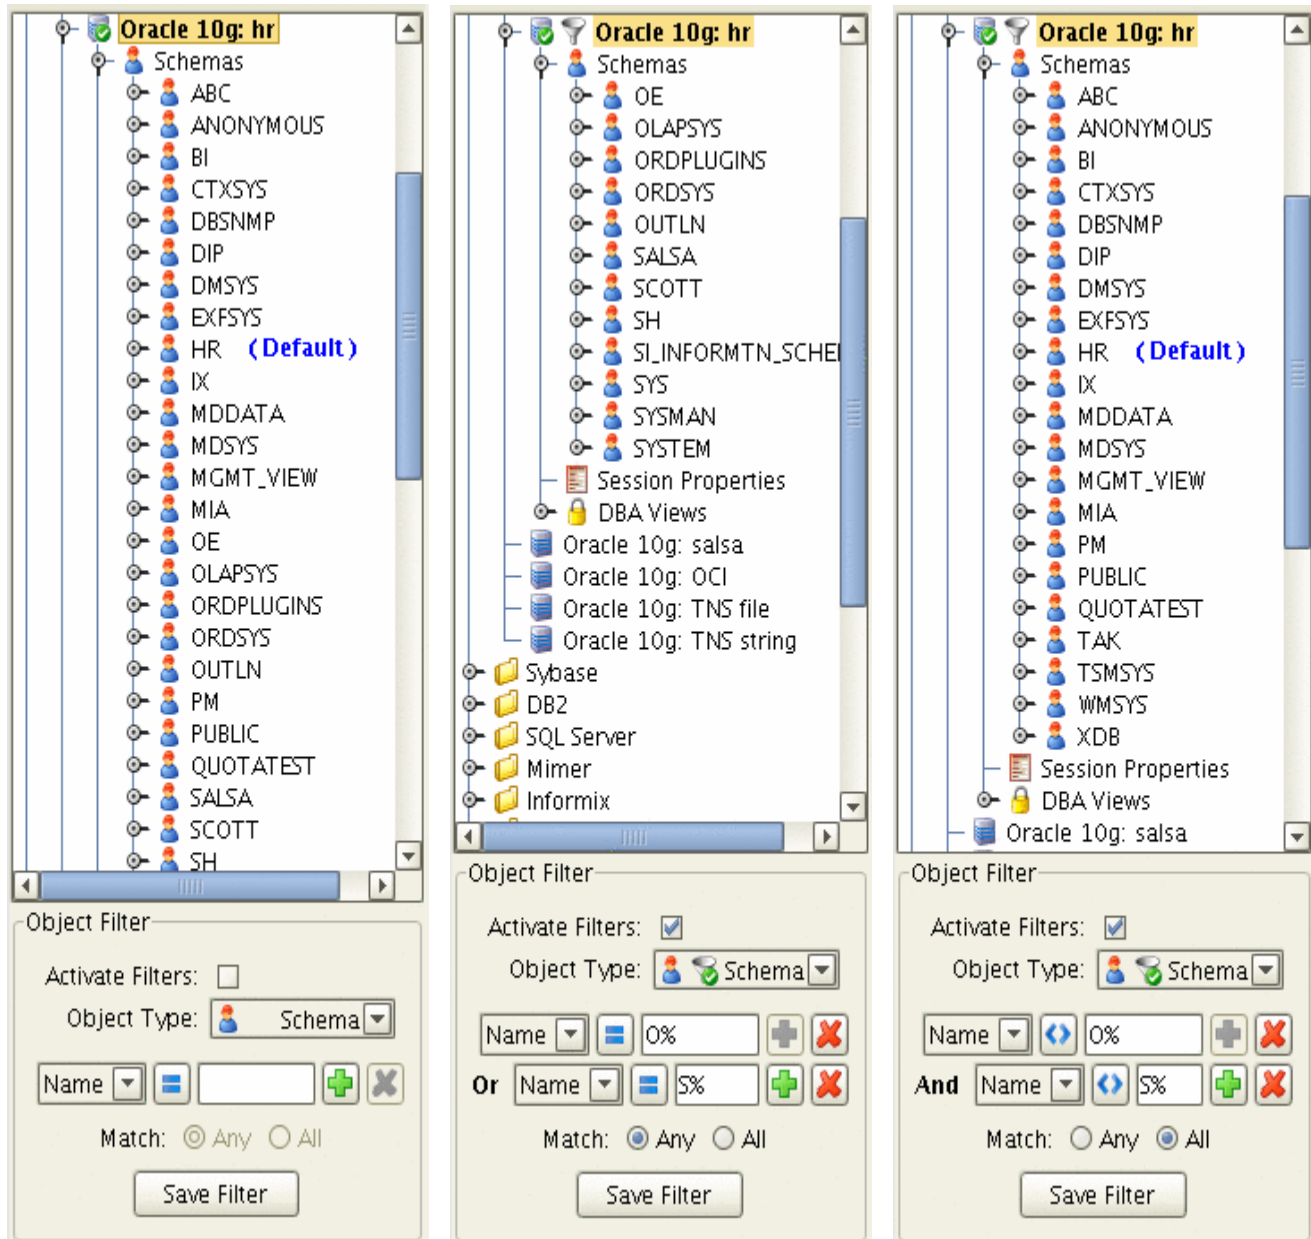


Figure: Examples of tree filter settings

An active filter for a database connection is represented by the funnel icon just before the database connection

name. The active state for a filter is defined using the Activate Filters checkbox in the Object Filter pane. A filter can only be activated if there are any filters defined. Up to 15 filters can be defined per object type.

Tip: A common requirement is to list only the default schema or catalog (database) in the database objects tree. This can be accomplished using the filtering functionality, but the recommended way is to do this with the Show only default Database or Schema property in the Properties tab for the Database Connection object. You can read more about this in the [Tool Properties](#) section.

Show Table Row Count

The Database->Show/Hide Table Row Count menu choice decides whether the number of rows for table objects will be listed after the name of the table in the tree.

Note: Enabling this property results in a performance degradation.

Database Profiles

A [Database Profile](#) is the foundation for database specific support in DbVisualizer. A database profile is, somewhat simplified, a definition of which kind of information that is presented in the database objects tree and in the various object views for a specific database engine. In addition, the profile defines the actions for the object types defined in the profile. DbVisualizer loads the matching database profile when you connect to a database. If no matching profile is found, or if you are running DbVisualizer Free, DbVisualizer uses a [Generic](#) profile with just the general database information and actions included.

Database Specific Support

DbVisualizer Personal currently offer database specific support (database profiles) for the following databases (click links for details):

- [Oracle](#)
- [DB2](#)
- [Sybase ASE](#)
- [SQL Server](#)
- [MySQL](#)
- [PostgreSQL](#)
- [Informix](#)
- [Mimer](#)
- [JavaDB/Derby](#)

The specialized database profiles define different object types, so the database objects tree may look different depending on which database you are connected to. The structure and organization of a database profile is also something that may impact the layout of the tree, even though the provided ones are similar in their structure. There are two root nodes in the majority of the profiles:

- User objects
- DBA objects

User objects are, for example, tables, views, triggers, and functions, while DBA objects most often are objects that require administration privileges in the database in order to access them. DbVisualizer puts all DBA objects under the DBA Views tree node. If you connect to a database using an account with insufficient privileges to access a DBA object, you may see error messages if you try to select nodes under the DBA Views node. This is an example of the DBA sub tree.

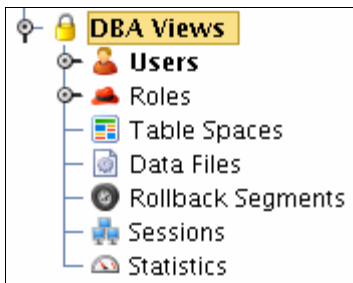


Figure: The DBA Views tree object

Note: Database profiles are defined in XML and it is quite easy to extend and modify them. Read more in the [Plug-in Framework](#) document.

Generic profile

DbVisualizer supports a wide range of databases. The nature of the databases and what they support differ from vendor to vendor, so the appearance and structure of the tree below the Database Connection objects for different databases differ as well. The generic database profile (the only profile available in DbVisualizer Free) displays objects based on what JDBC offers in terms of database information (aka metadata information). DbVisualizer asks the JDBC driver for all schemas, databases, tables and procedures, and then builds the tree based on what the driver returns.

The advantage of using JDBC to get database metadata is that it is a standard way to access the information, independent of the database engine type; the JDBC driver layer hides the proprietary details about where and how the information is really stored. The drawback with using JDBC is that JDBC doesn't offer access to all metadata a database may hold. While the information presented by the generic profile, with its reliance on JDBC, is sufficient for many tasks, a database specific profile offers far more details as well as more features. If you use DbVisualizer Free with one of the databases supported by database specific profiles, you may want to upgrade to the DbVisualizer Personal edition.

The generic database profile when used for an Oracle connection look as follows:




Figure: The generic database profile when applied to an Oracle database connection

The appearance of the generic database profile may include schema objects and/or catalog objects depending on whether the database supports these objects. The Procedures object always appear in the tree, regardless of if the database connection supports procedures or not.

The following sections describe the objects provided by the generic profile.

Catalog/Database object

Catalog is the term used in JDBC and some database engines for a logical grouping of database objects. Other database engines, e.g., Sybase, PostgreSQL, SQL Server and MySQL use the term Database for, more or less, the same purpose. Both terms are used in interchangeably in DbVisualizer.

The Object View for a Catalog object in the generic profile is a pane with two tabs, Tables and References. The Tables tab lists all the tables that are located in the catalog while References shows the exact same list of tables but instead as a referential integrity graph.

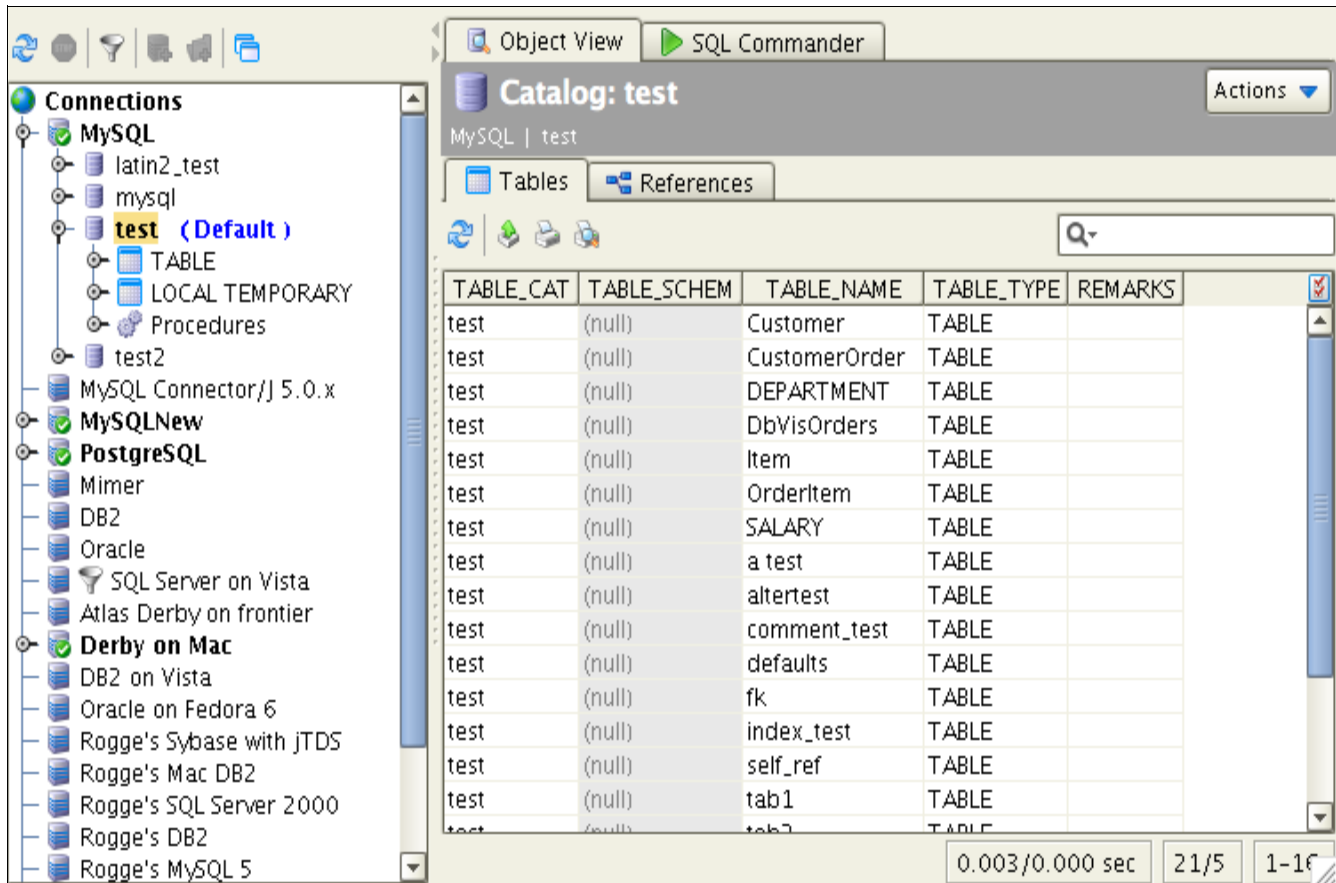


Figure: The view for Catalog objects

The child objects shown for a catalog object depend on the capabilities of the JDBC driver. Typically, a child object represents a type of table that the driver use to categorize the tables in a catalog, e.g., regular tables or system tables. For instance, the example in the figure above shows a MySQL database connection with catalog objects as its child object. The catalog child objects are TABLE and LOCAL TEMPORARY, because these are the table types that the MySQL JDBC driver supports (these table types are the same as those listed in the Table Types tab when selecting a database connection object). For other databases, you may see child objects representing other table types.

Tip 1: You can double click on a catalog object to display the detail view in a separate window.

Tip 2: Select one or several rows (cells) in the tables grid and then choose Script: SELECT ALL to create a select script for the selected tables, copied to the current SQL Commander where it can be executed.

Schema object

The generic profile Schema object tree and view are organized in the same way as for the Catalog objects. There is in fact no difference except that the schema objects are in another level in the tree and is represented by a different icon.

The following screen shot shows the information for the selected schema with the Reference tab selected.

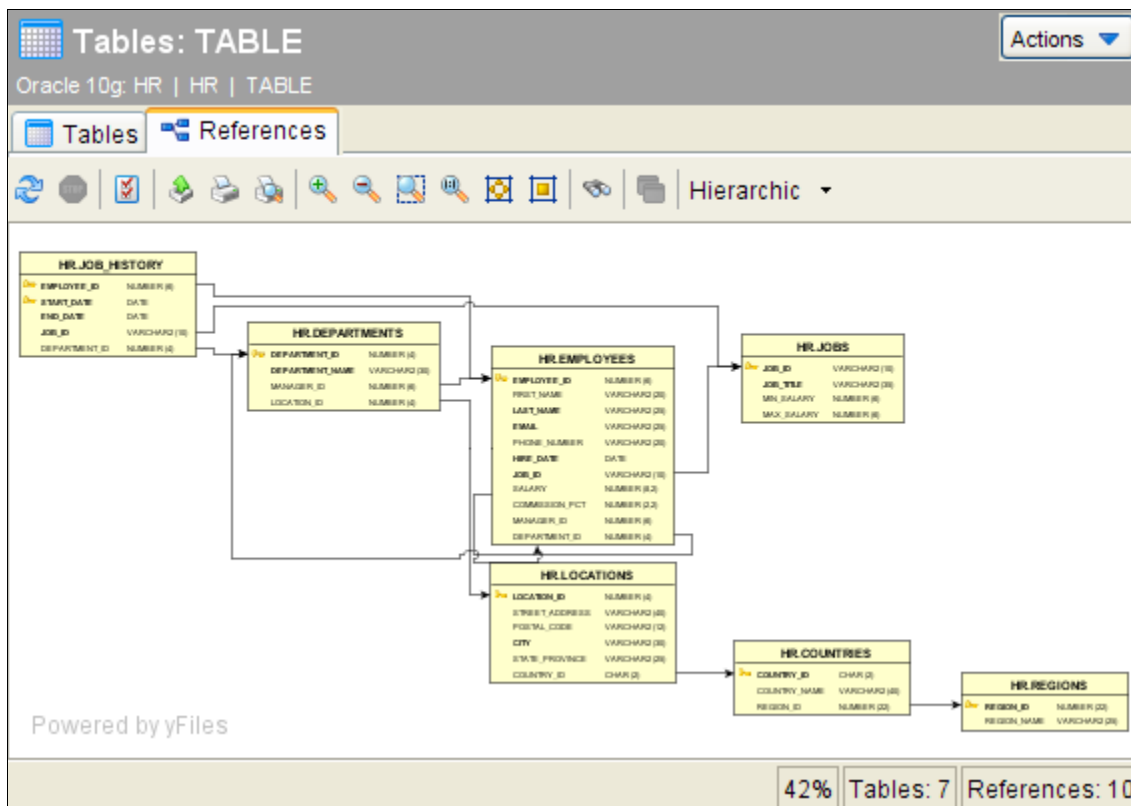


Figure: The view for Schema objects

Table Type object

The generic profile Table Type object has been briefly described earlier. The name and the number of table type objects are determined by the driver as DbVisualizer asks for the supported table types. When DbVisualizer retrieves all tables it checks each table's type and puts them into the matching table type object. The reason is simply to make the tree easier to browse.

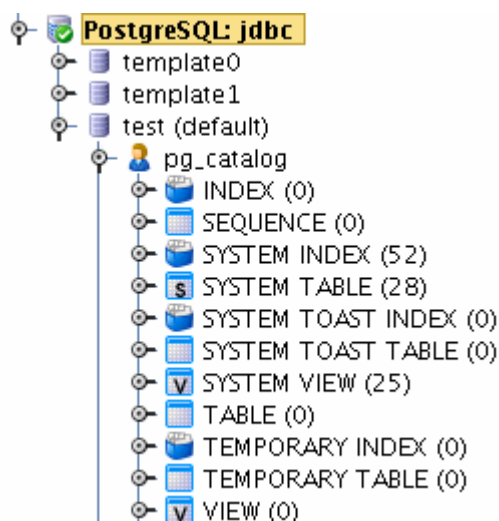


Figure: Example of table type objects for PostgreSQL

Note: Even though the figure above lists objects as INDEX, SEQUENCE, VIEW, etc are all treated as tables by DbVisualizer.

Table object

The Table object is probably the most frequently accessed object in the tree, since its Object View shows not only a lot of information about the table but also the data the table holds. This is also the place where you can edit the table data.

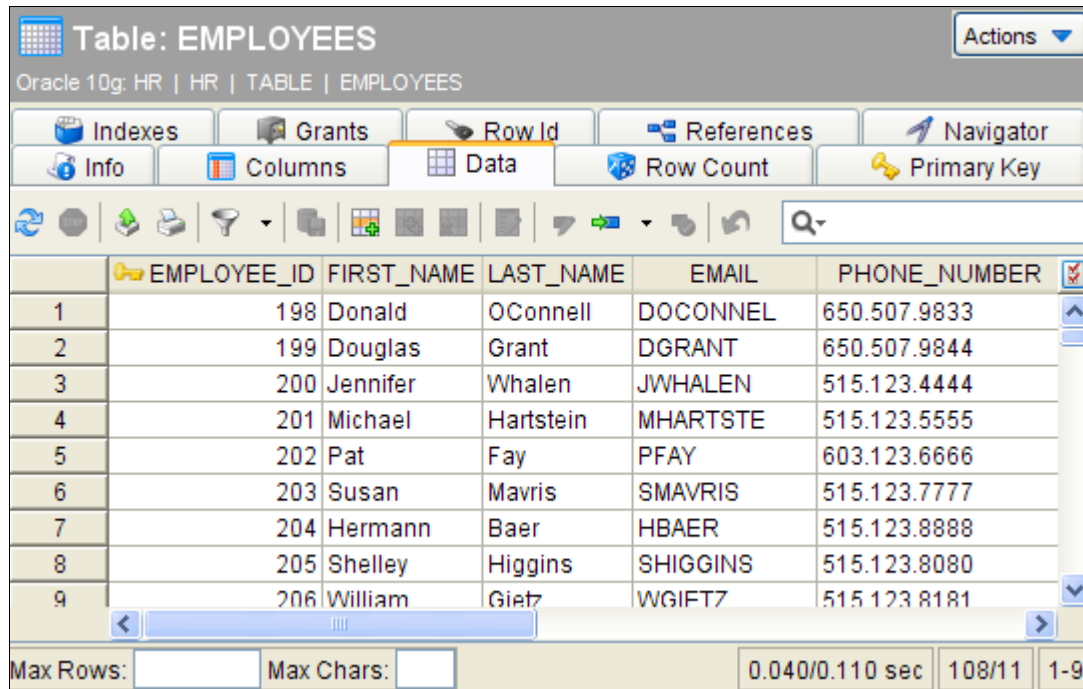



Figure: The view for Table objects

The Object View for a table object contains the following tabs:

Tab	Description
Info	Brief information about the table object
Columns	This tab lists type information about all columns in the table
Data	Read more in Data tab
Row Count	Lists the table row count
Primary Key	Shows the primary key
Indexes	Lists all indexes for the table
Grants	Displays any privileges for the table
Row Id	Displays the optimal set of columns that uniquely identifies a row
References	Read more in References tab
Navigator	Read more in Navigator tab

Procedure object

The Procedure object shows the name of the procedure or function in the tree, and the Object View lists the parameters that are used when calling it.



PROCEDURE_SCHEM	PROCEDURE_NAME	COLUMN_NAME	COLUMN_TYPE	DATA_T
HR	ADD_JOB_HISTORY	P_EMP_ID	1	
HR	ADD_JOB_HISTORY	P_START_DATE	1	
HR	ADD_JOB_HISTORY	P_END_DATE	1	
HR	ADD_JOB_HISTORY	P_JOB_ID	1	
HR	ADD_JOB_HISTORY	P_DEPARTMENT_ID	1	

Figure: The procedure object


The Object View shows a list of column names for the selected procedure.

Object Views

The Object View tab shows detailed information about the selected tree object. The Object View may contain several sub tabs, depending on the current database profile and the type of the object selected in the tree. There may also be several representations of the same information, providing different views of the information. The following sections describe the different views, or visual presentation forms, provided by DbVisualizer.

Grid

The Grid view is the most common one as it displays the data in a standard grid style.



OWNER	INDEX_NAME	INDEX_TYPE	TABLE_OWNER	TABLE_NAME	TAB
HR	COUNTRY_C_ID_PK	IOT - TOP	HR	COUNTRIES	TAB
HR	DEPT_ID_PK	NORMAL	HR	DEPARTMENTS	TAB
HR	DEPT_LOCATION_IX	NORMAL	HR	DEPARTMENTS	TAB
HR	EMP_DEPARTMENT_IX	NORMAL	HR	EMPLOYEES	TAB
HR	EMP_EMAIL_UK	NORMAL	HR	EMPLOYEES	TAB
HR	EMP_EMP_ID_PK	NORMAL	HR	EMPLOYEES	TAB
HR	EMP_JOB_IX	NORMAL	HR	EMPLOYEES	TAB
HR	EMP_MANAGER_IX	NORMAL	HR	EMPLOYEES	TAB
HR	EMP_NAME_IX	NORMAL	HR	EMPLOYEES	TAB
HR	JHIST_DEPARTMENT_IX	NORMAL	HR	JOB_HISTORY	TAB
HR	JHIST_EMPLOYEE_IX	NORMAL	HR	JOB_HISTORY	TAB
HR	JHIST_EMP_ID_ST_DATE_PK	NORMAL	HR	JOB_HISTORY	TAB
HR	JHIST_JOB_IX	NORMAL	HR	JOB_HISTORY	TAB
HR	JOB_ID_PK	NORMAL	HR	JOBS	TAB
HR	LOC_CITY_IX	NORMAL	HR	LOCATIONS	TAB
HR	LOC_COUNTRY_IX	NORMAL	HR	LOCATIONS	TAB

Figure: The Grid view

Form

The Form view extends the Grid view by adding a form below the grid. Click on a row in the grid and the information is displayed in the form.

The screenshot shows the Oracle Database Objects Explorer interface. The title bar reads "Storage" and the breadcrumb is "Oracle on Fedora 6 | DBA Views | Storage". There are three tabs: "Tablespaces", "Data Files", and "Rollback Segments", with "Rollback Segments" selected. Below the tabs are icons for refresh, save, print, and delete, and a search box. A grid view displays a list of segments:

SEGMENT_NAME	OWNER	TABLESPACE_NAME	SEGMENT_ID	FILE_ID	BLOCK_ID	INITIAL
SYSTEM	SYS	SYSTEM	0	1	9	
_ <u>SYSSMU1</u> \$	PUBLIC	UNDO	1	2	9	
_ <u>SYSSMU10</u> \$	PUBLIC	UNDO	10	2	153	
_ <u>SYSSMU2</u> \$	PUBLIC	UNDO	2	2	25	
_ <u>SYSSMU3</u> \$	PUBLIC	UNDO	3	2	41	
_ <u>SYSSMU4</u> \$	PUBLIC	UNDO	4	2	57	

Below the grid is a status bar showing "0.194/0.003 sec", "11/14", and "1-7". There are icons for refresh, save, print, and delete, and another search box. The form view below shows the details for the selected row:

Name	Value
SEGMENT_NAME	SYSTEM
OWNER	SYS
TABLESPACE_NAME	SYSTEM
SEGMENT_ID	0
FILE_ID	1
BLOCK_ID	9
INITIAL_EXTENT	114688

Figure: The Form view

If there is only one row of data, only the form is displayed.

Source

The Source view is typically used to show the source for functions, procedures, triggers, etc. It is based on a read only editor with SQL syntax coloring. The sub toolbar buttons from the left:

- Reload the data from the database
- Stop loading the data from the database
- Export the data to file
- Print the data
- Copy the data to SQL Commander
- Format the SQL

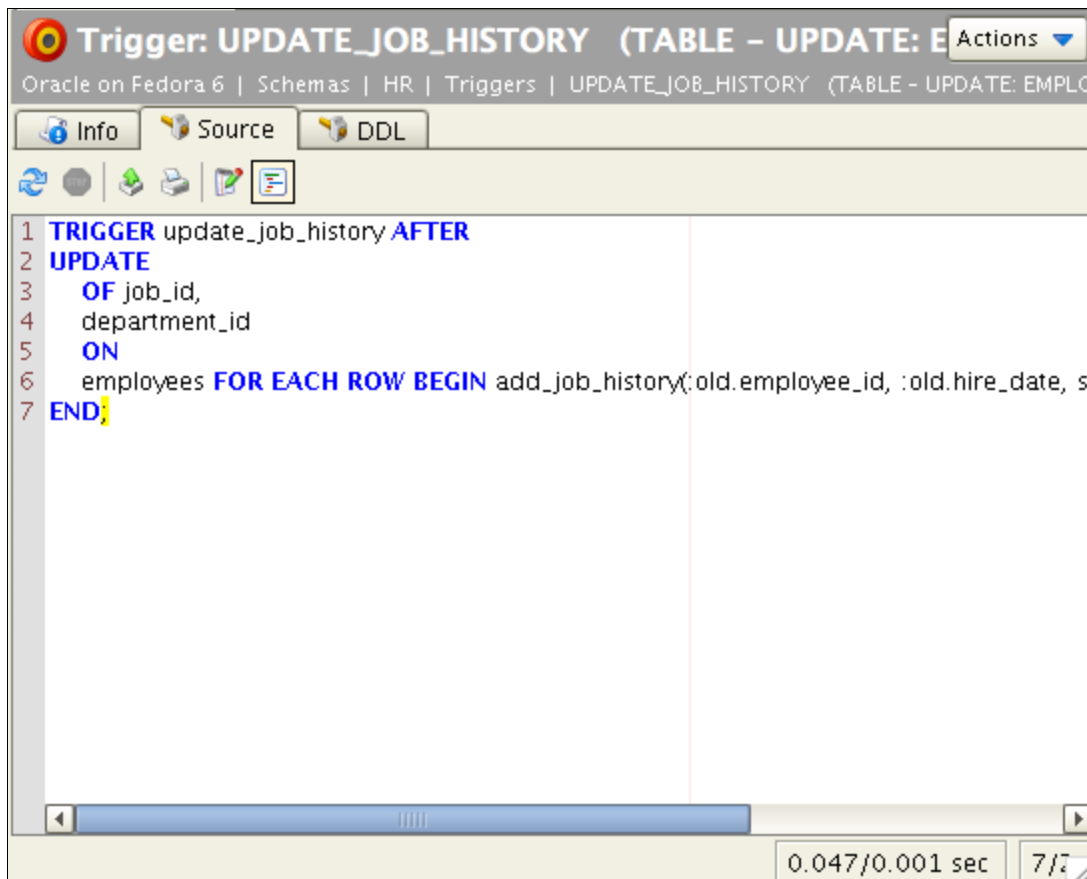


Figure: The Source view

Table Row Count

The row count view is really simple: it only shows the number of rows in the selected object.

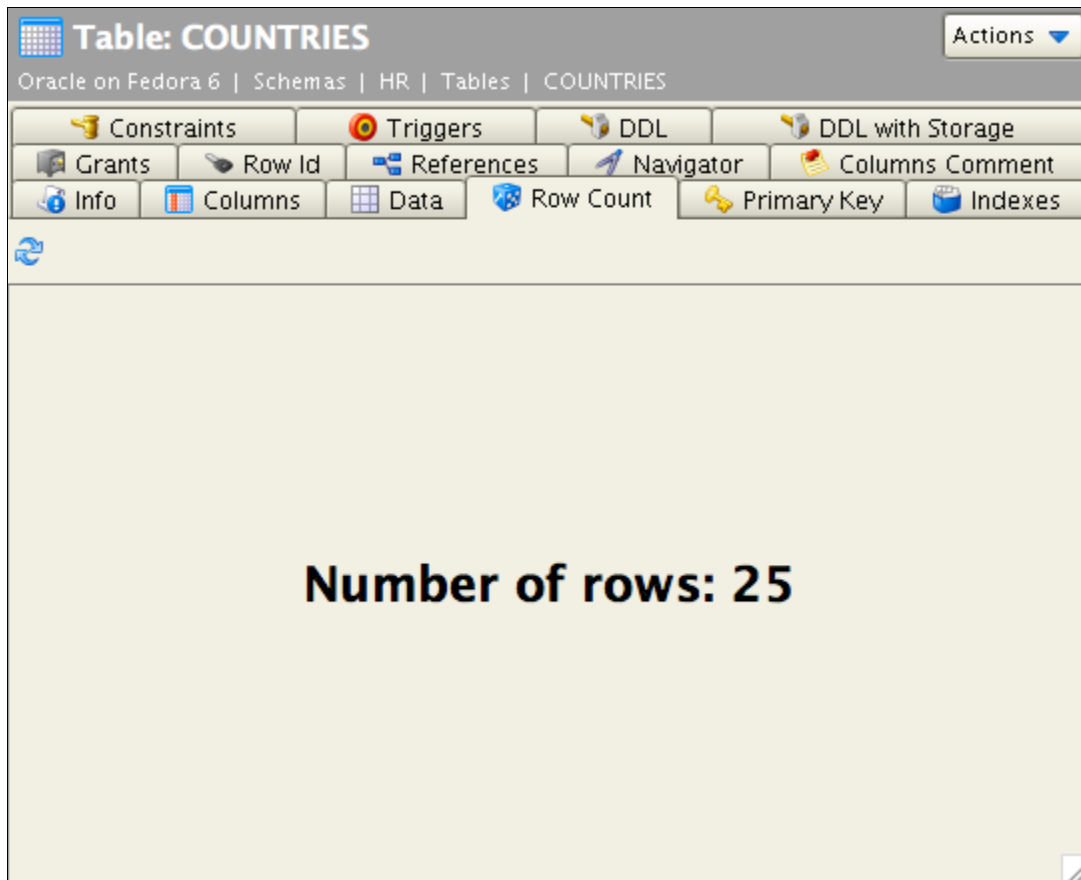



Figure: The Row Count view

Table Data

You use the Data tab to browse the data in the table and to do various data related operations. This view is based on the [generic grid](#), but it adds a few more visual components to limit the max number of rows, the width of text columns and the collection of data tab specific operations in the right-click menu. In addition, you can also use a filter limit the data to the rows that match the filter. The data tab is the place to do [edits](#) in DbVisualizer Personal.



The screenshot shows the DbVisualizer interface for the 'Table: COUNTRIES' in the 'HR' schema. The 'Data' tab is active, displaying a grid of 15 rows. The columns are COUNTRY_ID, COUNTRY_NAME, and REGION_ID. The row for Brazil (COUNTRY_ID: BR, REGION_ID: 2) is highlighted in yellow. The interface includes a toolbar with icons for Constraints, Triggers, DDL, DDL with Storage, Grants, Row Id, References, Navigator, Columns Comment, Info, Columns, Data, Row Count, Primary Key, and Indexes. At the bottom, there are input fields for 'Max Rows: 2500' and 'Max Chars: ' and performance metrics: '0.009/0.004 sec', '25/3', and '1-15'.

	COUNTRY_ID	COUNTRY_NAME	REGION_ID
1	AR	Argentina	2
2	AU	Australia	3
3	BE	Belgium	1
4	BR	Brazil	2
5	CA	Canada	2
6	CH	Switzerland	1
7	CN	China	3
8	DE	Germany	1
9	DK	Denmark	1
10	EG	Egypt	4
11	FR	France	1
12	HK	HongKong	3
13	IL	Israel	4
14	IN	India	3
15	IT	Italy	1

Figure: The Data tab for Table objects

Right-click menu

The Data tab grid right-click menu contains some operations in addition to those the standard grid right-click menu. The additional operations are primarily for creating SQL statements based on the current selection. Choosing any of these creates the appropriate SQL and then switch the view to the SQL Commander tab. You must use these operations to edit table data in the DbVisualizer Free edition. With the DbVisualizer Personal edition, you can instead use inline and form based editing. (Information about the standard right click menu operations are available in the [Getting Started and General Overview](#) document).

You can generate SQL with either static values as they appear in the grid, or with DbVisualizer variables. A variable is essentially a placeholder for a value in an SQL statement. When the statement is executed, DbVisualizer locates all variables and presents them in a dialog where you can enter or modify values for the variables. DbVisualizer replaces the variable placeholders with the new values before executing the statement. Variables can be used in any SQL statement and DbVisualizer relies on them heavily. (Read more about variables in the [SQL Commander](#) document).

Whether to use variables in the SQL statements generated by the right-click menu SQL operations depends on the Table Data->Include Variables in SQL setting in Tool Properties, under the General tab. By default, variables are being used in the statement. If you disable the property, static values are instead used in the generated statement.

Here is an example with the Include Variables in SQL setting enabled and then disabled. The SQL is generated when the select * where operation is selected based on the selection in the previous figure.

Include Variables in SQL is enabled
<pre>select * from HR.COUNTRIES where COUNTRY_NAME = \$\$COUNTRY_NAME (where) Brazil String where nullable ds=40 dt=VARCHAR \$\$</pre>
Include Variables in SQL is disabled
<pre>select * from HR.COUNTRIES where COUNTRY_NAME = 'Brazil'</pre>

The following lists the generated SQL for each of the operations based on the selection of COUNTRY_NAME = Brazil, with variables disabled.

Operation	SQL Example
Script: SELECT ALL	<pre>select * from HR.COUNTRIES</pre>
Script: SELECT ALL WHERE	<pre>select * from HR.COUNTRIES where COUNTRY_NAME = 'Brazil'</pre>
Script: SELECT ALL WITH FILTER	<pre>select * from HR.COUNTRIES where REGION_ID = 1 // If this is the filter, see below</pre>
Script: INSERT INTO TABLE	<pre>insert into HR.COUNTRIES (COUNTRY_ID, COUNTRY_NAME, REGION_ID) values ('', '',)</pre>
Script: INSERT COPY INTO TABLE	<pre>insert into HR.COUNTRIES (COUNTRY_ID, COUNTRY_NAME, REGION_ID) values ('BR', 'Brazil', 2)</pre>
Script: UPDATE WHERE	<pre>update HR.COUNTRIES set COUNTRY_ID = 'BR', COUNTRY_NAME = 'Brazil', REGION_ID = 2 where COUNTRY_NAME = 'Brazil'</pre>
Script: DELETE WHERE	<pre>delete from HR.COUNTRIES where COUNTRY_NAME = 'Brazil'</pre>

Where Filter

The filter capability in the Data tab lets you limit the number of rows in the grid, using the same syntax as for an SQL WHERE clause. The Filter menu button in the grid toolbar contains all operations related to using a filter.

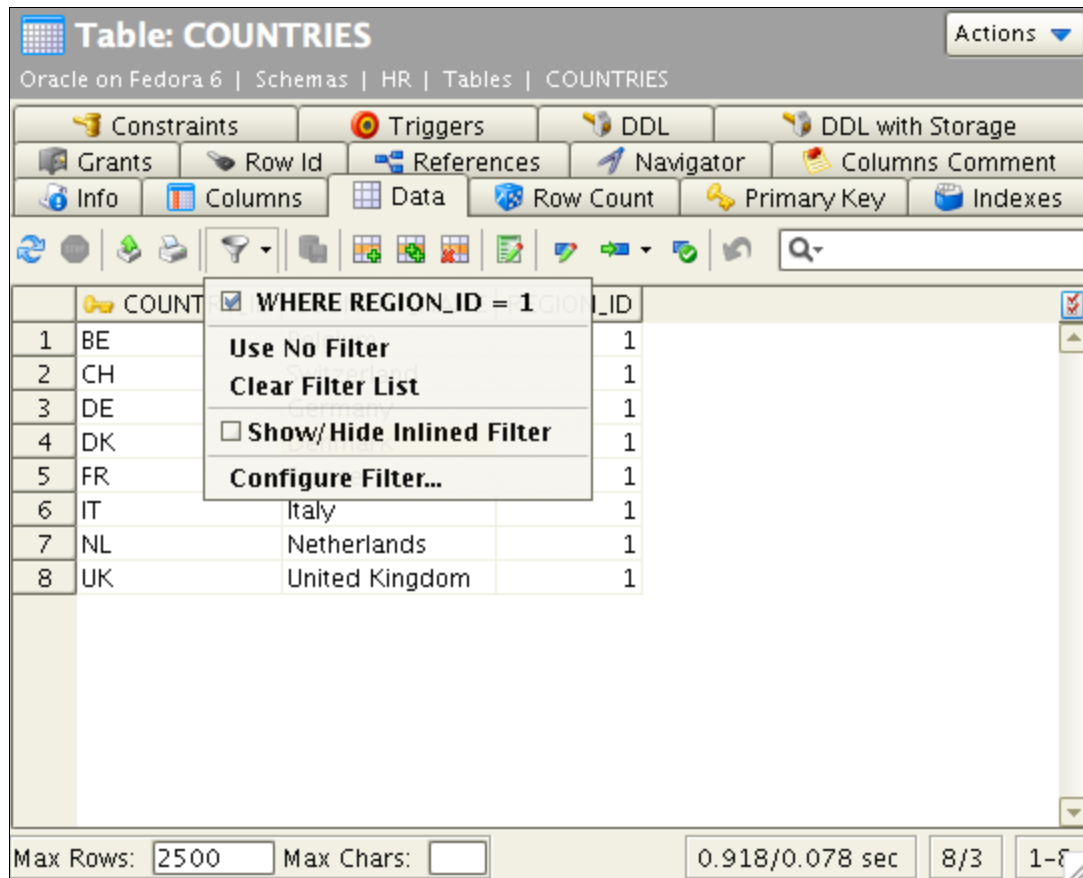


Figure: Filter menu

The top entries in the menu are previously used filters for the table, if any. The checkbox is selected for the filter that is currently in use. The filters are saved between DbVisualizer sessions, and you can toggle between them by selecting them from the menu. You use the Use No Filter choice to disable all filters for the table, and the Clear Filter List to permanently remove all filters for the table.

To create a new filter, select Configure Filter to launch the Filter Configuration dialog.

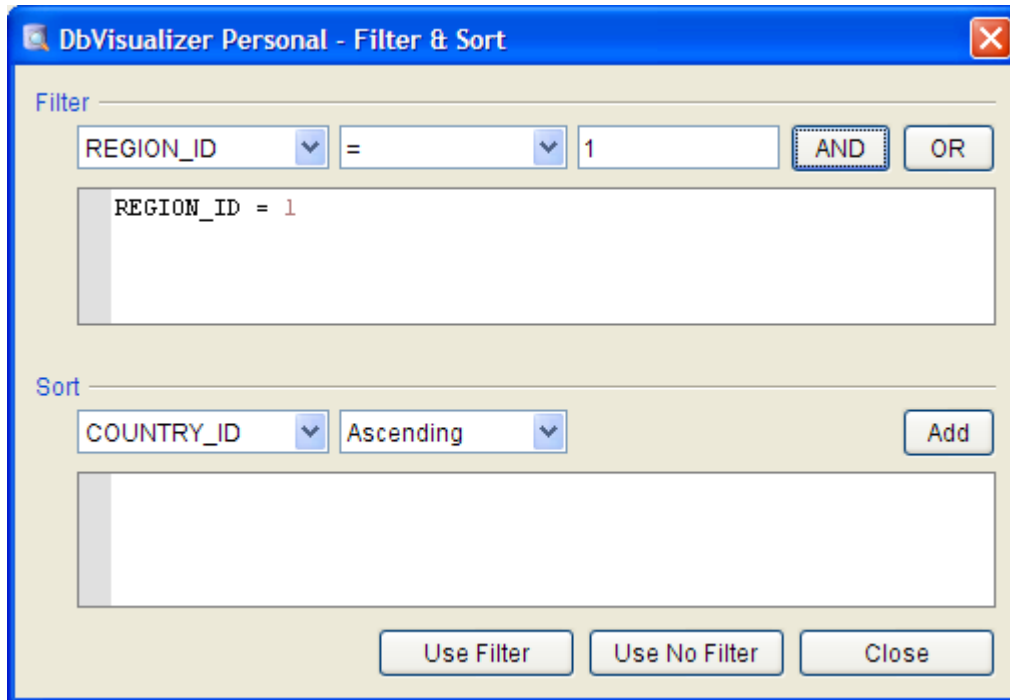


Figure: The Data tab Filter Configuration dialog

The Filter Configuration dialog contains one Filter area and a Sort area.

The Filter area is composed of two parts. The upper one is used to define a condition for a single column. You can use the two lists to select the column name and an operator, and enter the value of the column in the text field. You can use Ctrl-Enter while editing the value to force a reload of the grid based on that single filter. The lower part displays the complete filter and the buttons are used to control whether the newly entered filter will be AND'ed or OR'ed with the complete filter. The buttons change appearance based on whether there is any filter or not. While in the complete filter you can use Ctrl-Enter to force a reload based on the complete filter.

The Sort area is similar to the Filter area. You can select column names and sort order from the two lists, and click the Add button to add the sort criteria for the single column to the complete criteria.

Click the Use Filter button to apply the filter and save it, and close the dialog by clicking the Close button.

If you often need to tweak the filter conditions and want a more compact user interface, you can use the inline filter view. Use the Show/Hide Inline Filter choice in the Filter menu to toggle the visibility of the inline filter.

Table: COUNTRIES Actions ▾

Oracle on Fedora 6 | Schemas | HR | Tables | COUNTRIES

Constraints Triggers DDL DDL with Storage
Grants Row Id References Navigator Columns Comment
Info Columns Data Row Count Primary Key Indexes

Filter: Use Filter Use No Filter

	COUNTRY_ID	COUNTRY_NAME	REGION_ID
1	AR	Argentina	2
2	BR	Brazil	2
3	CA	Canada	2
4	MX	Mexico	2
5	US	United States of America	2

Max Rows: Max Chars: 0.015/0.000 sec 5/3 1-5


Figure: Data tab with the Inline filter enabled

The inline filter is displayed above the grid. You can edit the condition in the text field and use Ctrl-Enter or click the Use Filter button to apply the modified condition. Instead of manually typing column names in the field use the Ctrl-Space key binding to show a list of available columns.

Quick Filter

The quick filter acts on the data that is already in the grid, as opposed of a WHERE filter which is used to limit the number of rows fetched from the database. With a Quick filter, you can easily list only those rows in the grid that match the entered search string.

The following figure shows data that matches the search string "d". Matching cells are highlighted.



The screenshot shows the Oracle Database Objects Explorer interface. The table 'COUNTRIES' is displayed with columns: COUNTRY_ID, COUNTRY_NAME, and REGION_ID. A quick filter 'd' is applied, highlighting rows where the COUNTRY_NAME contains the letter 'd'. The highlighted rows are: DE (Germany), DK (Denmark), and US (United States of America). The interface also shows various toolbars and a status bar at the bottom.

	COUNTRY_ID	COUNTRY_NAME	REGION_ID
1	CA	Canada	2
2	CH	Switzerland	1
3	DE	Germany	1
4	DK	Denmark	1
5	IN	India	3
6	NL	Netherlands	1
7	UK	United Kingdom	1
8	US	United States of America	2

Figure: Using the Quick Filter

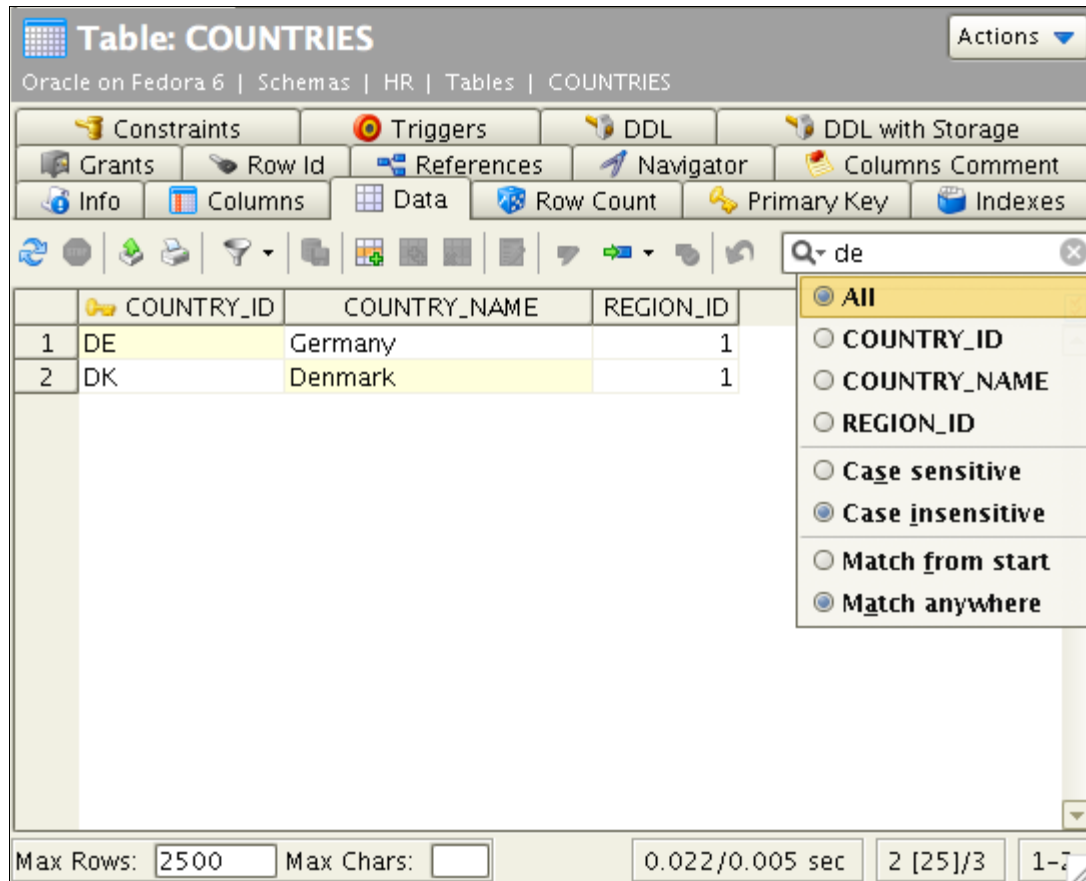
Entering successive characters will narrow the result even further, as in the following figure.

The screenshot shows the Oracle Database Objects Explorer interface. The title bar reads "Table: COUNTRIES" and the breadcrumb path is "Oracle on Fedora 6 | Schemas | HR | Tables | COUNTRIES". The interface includes a toolbar with various options like Constraints, Triggers, DDL, DDL with Storage, Grants, Row Id, References, Navigator, Columns Comment, Info, Columns, Data, Row Count, Primary Key, and Indexes. A search bar at the top right contains the text "de". Below the search bar, a table displays the filtered data for the COUNTRIES table. The table has four columns: an index column, COUNTRY_ID, COUNTRY_NAME, and REGION_ID. Two rows are visible: row 1 with COUNTRY_ID 'DE' and COUNTRY_NAME 'Germany', and row 2 with COUNTRY_ID 'DK' and COUNTRY_NAME 'Denmark'. Both rows have a REGION_ID of 1. At the bottom of the window, there are status indicators: "Max Rows: 2500", "Max Chars: [empty]", "0.022/0.005 sec", "2 [25]/3", and "1-7".

	COUNTRY_ID	COUNTRY_NAME	REGION_ID
1	DE	Germany	1
2	DK	Denmark	1

Figure: Refining the filtering

The Quick Filter pull-down menu (click on the down arrow next to the magnifying glass) lets you choose if the filter should match cells in all columns or just one selected column, case or case insensitive matching, and where in the cell the value must match.



Monitor row count


Read more about the Monitor Row Count and Monitor Row Count Difference in [Monitor and Charts](#).

Editing

Read about data editing in [Edit Table Data](#)

DDL Viewer

The DDL Viewer tabs appear only for Table and View objects and for databases that have [specialized database profiles](#).



```
1 CREATE
2   TABLE EMPLOYEES
3   (
4     EMPLOYEE_ID NUMBER(6,0) NOT NULL,
5     FIRST_NAME VARCHAR2(20),
6     LAST_NAME VARCHAR2(25) NOT NULL,
7     EMAIL VARCHAR2(25) NOT NULL,
8     PHONE_NUMBER VARCHAR2(20),
9     HIRE_DATE DATE NOT NULL,
10    JOB_ID VARCHAR2(10) NOT NULL,
11    SALARY NUMBER(8,2),
12    COMMISSION_PCT NUMBER(2,2),
13    MANAGER_ID NUMBER(6,0),
14    DEPARTMENT_ID NUMBER(4,0),
15    CONSTRAINT EMP_EMP_ID_PK PRIMARY KEY (EMPLOYEE_ID),
16    CONSTRAINT EMP_DEPT_FK FOREIGN KEY (DEPARTMENT_ID) REFERENCES DEPARTMENTS (DEPARTMENT_ID),
17    CONSTRAINT EMP_JOB_FK FOREIGN KEY (JOB_ID) REFERENCES JOBS (JOB_ID),
18    CONSTRAINT EMP_MANAGER_FK FOREIGN KEY (MANAGER_ID) REFERENCES EMPLOYEES (EMPLOYEE_ID),
19    CONSTRAINT EMP_EMAIL_UK UNIQUE (EMAIL),
20    CONSTRAINT EMP_JOB_NUM_CHECK CHECK ("JOB_ID" IS NOT NULL)
```

The screenshot shows a window titled "Untitled*" with a text area containing the DDL for the EMPLOYEES table. The text is color-coded: CREATE, TABLE, and CONSTRAINT are in blue; EMPLOYEES, PRIMARY KEY, FOREIGN KEY, REFERENCES, and UNIQUE are in red; and the rest of the text is in black. The window has a standard toolbar at the top and a status bar at the bottom showing "1:1" and "INS".

Figure: The DDL viewer for a table

References

The References tab for a Table object shows how the table references other tables (e.g., Imported Keys) and how other tables reference the selected table (i.e., Exported Keys), based on primary and foreign key declarations. Use the sub tabs at the bottom of the display to show either view. The following shows the references from the table.

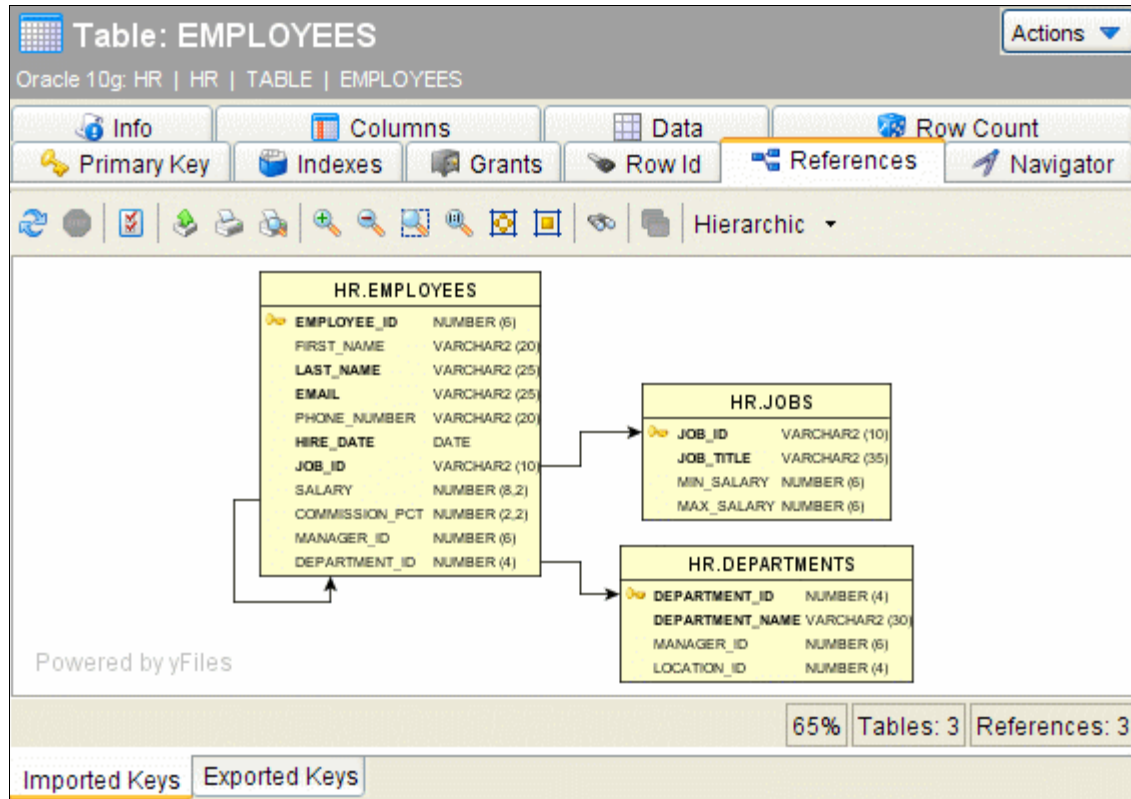


Figure: The references graph showing imported keys for a table

Navigator

The Navigator tab provides an interactive way to navigate in data by following primary key and foreign key references.

The screenshot shows the 'Table: DEPARTMENTS' interface. The 'Navigator' tab is active, displaying a graph with two paths from 'HR.DEPARTMENTS' to 'HR.EMPLOYEES'. The first path is labeled 'DEPARTMENT_ID DEPARTMENT_NAME Human Resources' and leads to a box containing 'DEPARTMENT_ID 40'. The second path is labeled 'DEPARTMENT_ID DEPARTMENT_NAME IT' and leads to a box containing 'DEPARTMENT_ID 60'. Below the graph is a data grid for 'HR.EMPLOYEES' with columns: EMPLOYEE_ID, FIRST_NAME, LAST_NAME, EMAIL, PHONE_NUMBER, HIRE_DATE, JOB_ID, SALARY, and COMMISSION_PCT. The grid contains five rows of data. At the bottom, there are controls for 'Max Rows: 1000', 'Max Chars: []', and performance metrics '0.002/0.002 sec', '5/11', and '1-5'.

EMPLOYEE_ID	FIRST_NAME	LAST_NAME	EMAIL	PHONE_NUMBER	HIRE_DATE	JOB_ID	SALARY	COMMISSION_PCT
103	Alexander	Hunold	AHUNOLD	590.423.4567	1990-01-03	IT_PROG	9000.00	(null)
104	Bruce	Ernst	BERNST	590.423.4568	1991-05-21	IT_PROG	6000.00	(null)
105	David	Austin	DAUSTIN	590.423.4569	1997-06-25	IT_PROG	4800.00	(null)
106	Valli	Pataballa	VPATABAL	590.423.4560	1998-02-05	IT_PROG	4800.00	(null)
107	Diana	Lorentz	DLORENTZ	590.423.5567	1999-02-07	IT_PROG	4200.00	(null)

Figure: The Navigator tab showing two navigation cases

The tab contains a graphic view showing navigation cases (paths through the data) at the top and a data grid showing the data for the navigation case selected in the graph. You navigate in the data by selecting the row in the grid that holds the key value you want to follow, e.g., a specific department in the example shown in the figure, and then select a primary or foreign key relationship from the Related Table list above the grid. This creates a new navigation case in the graph and updates the grid with the corresponding data.

How to use the navigator is described in more detail in the [Data Navigation](#) section.

Procedure Editor

The procedure editor is used to browse, edit and compile procedures, functions, packages and package bodies. The procedure contains the source editor and options to display parameter info and compilation error list. If error(s) occur during compilation, click the error to highlight the related row in the source editor.

Note that the procedure editor currently only support Oracle database.

More information is can be found in the [Procedure Editor](#) document.

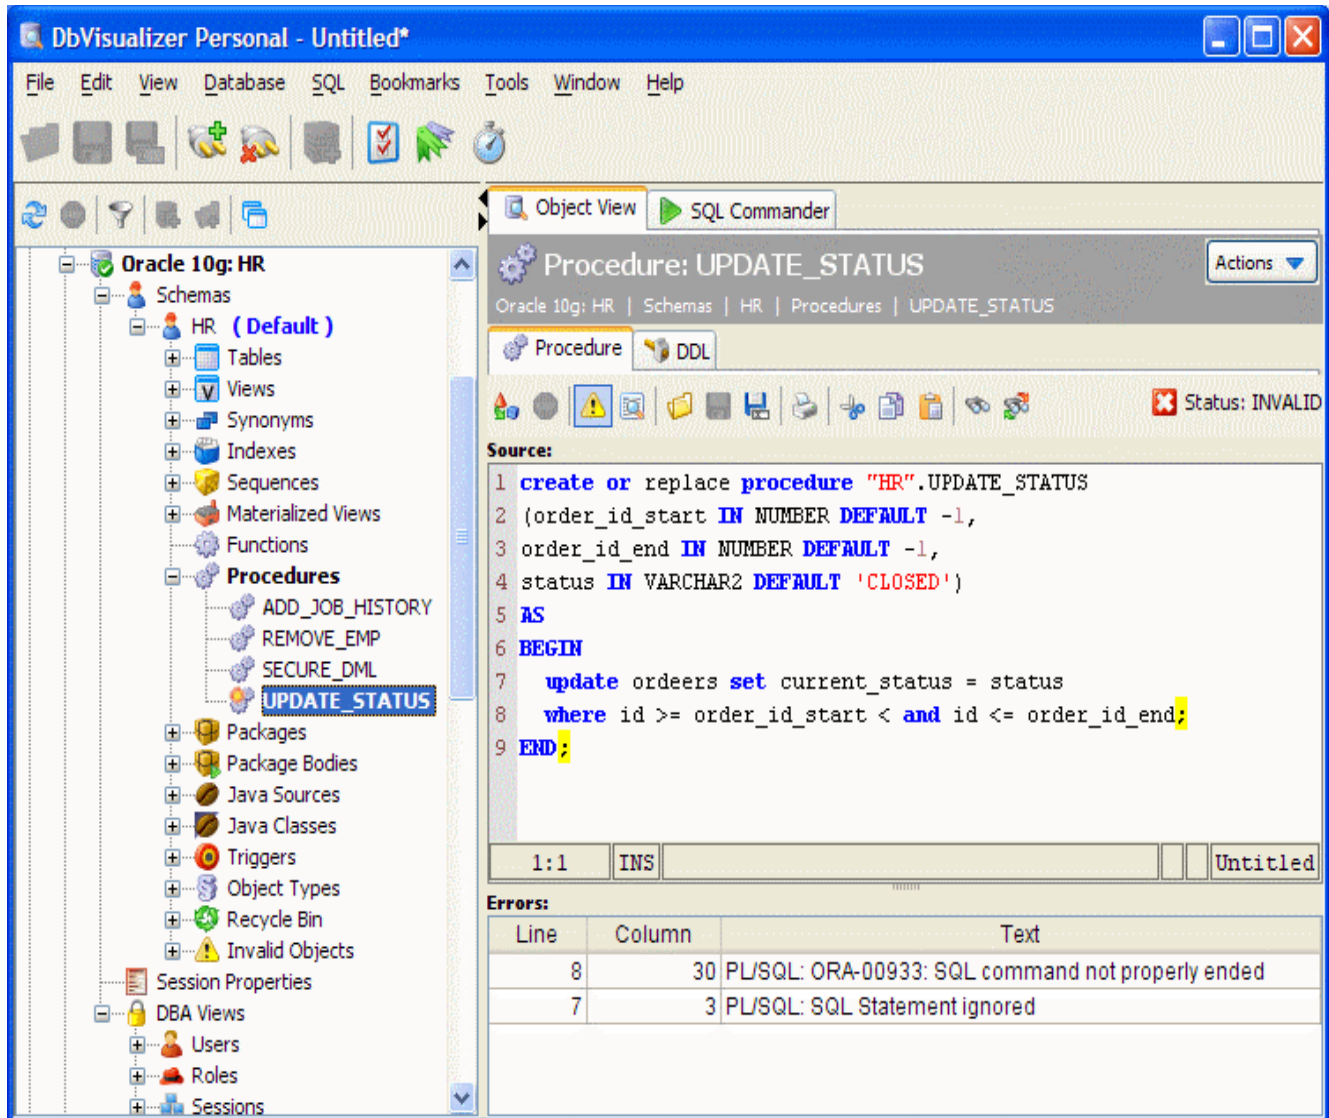


Figure: The procedure editor for functions, procedures, packages and package bodies

SQL Commander

Introduction

The SQL Commander is used to edit, format and execute SQL statements or SQL scripts. Multiple editors may be open at the same time, each controlling its own SQL log and result sets. Result sets can be displayed in grid, text or chart formats.

The SQL Commander supports the following features:

- Syntax coloring
- Auto completion
- Multiple SQL editors
- Multiple result sets
- SQL editors displayed as tabs or windows
- Result sets displayed as tabs or windows
- Support for stored procedures producing multiple result sets
- SQL formatter with extensive customization options
- Execution control (stop on error/warning)
- View result sets as grid, text or chart
- Editable result sets with the inline or form editors
- Support for BLOB, CLOB and binary data
- View BMP, TIFF, PNG, GIF and JPEG images
- View XML data in tree or text format
- Export result sets as CSV, HTML, Excel, XML, SQL or text
- Batch execution enabling export of unlimited sized result sets
- SQL history saved between sessions
- Bookmark management (save favorite SQLs)
- Sort, quick filter and basic calculations of result sets
- Parameterized queries
- Drop objects dragged from the Objects Tree
- Auto Commit on/off support with confirmation checks if uncommitted updates
- Full key binding support with predefined key maps for for Windows, Mac OS X, Linux-UNIX, SQL Query Analyzer and TOAD users

Database specific support:

- [Oracle](#), [DB2](#) and [SQL Server](#): Explain Plan queries presented either in tree or graph format
- [Oracle](#): support for TIMESTAMPLTZ, TIMESTAMPTZ and XML data types
- [Oracle](#): support for DBMS Output

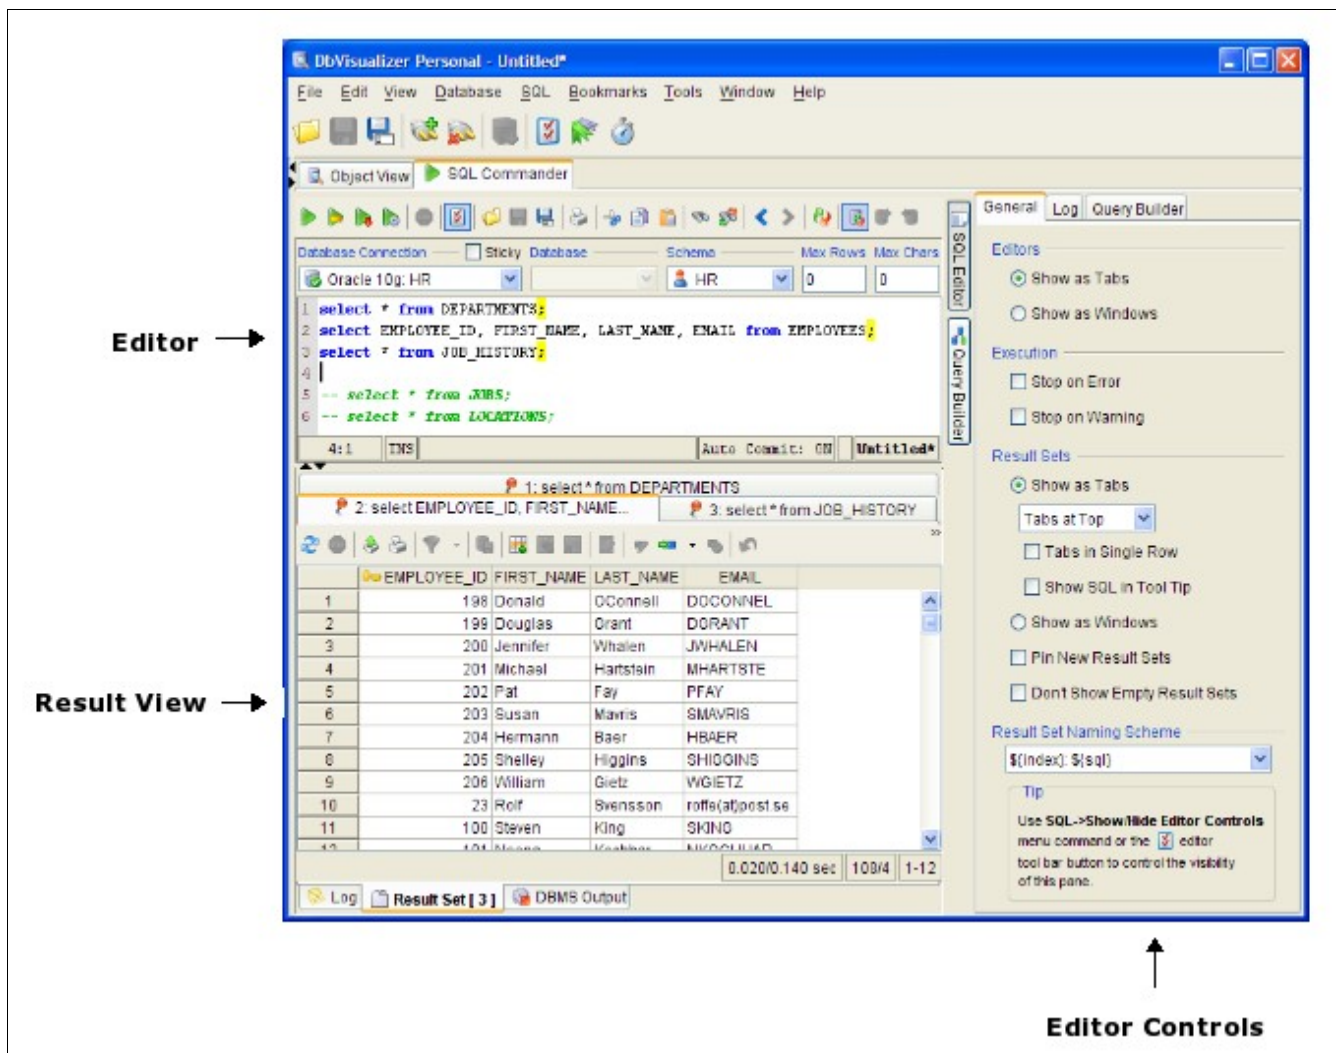


Figure: SQL Commander overview

The figure shows the editing area with its controls above and the output view in the lower part of the screen. The following sections describe all features and controls in the SQL Commander in detail.

Physical Database Connections and Transactions

The SQL Commander supports database transaction control via Auto Commit or manually using commit or rollback. The Use Single Physical Database Connection setting in connection properties specifies whether DbVisualizer will use one or multiple physical database connection. This setting is disabled by default and DbVisualizer will then always use at least one physical connection and one for every SQL editor that is created. Running a statement or sequence of statements in one SQL editor will not lock the rest of the user interface while it is executing. If however using a single physical database connection all of the UI is locked until the execution in the SQL Editor has completed. The reason for this behavior is that otherwise it could lead to data corruption if using the same physical database connection.

Another important feature is that the editor status bar show the number of uncommitted requests if auto commit is off. Pay extra attention to this as it indicates that you should complete the current transaction with either commit or rollback.

Editor

The SQL Commander always have at least one editor. It is called the primary editor and cannot be removed. To create additional editors use the File->Create SQL Editor menu choice or the appropriate key binding. To close an editor use the right-click menu on the editor tab or the close operations in the File menu.

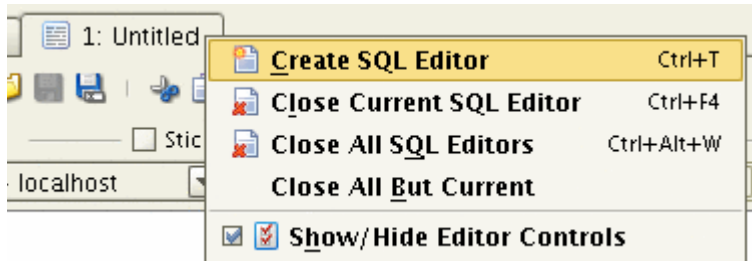


Figure: Editor tab menu

The SQL editor in DbVisualizer is based on the [NetBeans](#) editor module and supports all standard editing features. The right-click menu have the following operations:

	Execute	Ctrl+Enter
	Execute Current	Ctrl+.
	Execute Buffer	
	Execute Explain Plan	Ctrl+Alt+Enter
	Undo	Ctrl+Shift+Minus
	Redo	Ctrl+Shift+Z
	Cut	Ctrl+X
	Copy	Ctrl+C
	Paste	Ctrl+Y
	Clear All	Ctrl+Shift+Delete
	Find...	Ctrl+F
	Find Next	F3
	Find Previous	Shift+F3
	Replace...	Ctrl+H
	Goto Line...	Ctrl+G
	Lower Case	Ctrl+Shift+L
	Upper Case	Ctrl+Shift+U
	Comment Line	Ctrl+Shift+R
	Comment Block	Ctrl+Shift+B
	Format SQL	Ctrl+Shift+F
	Select All	
	Select Current Statement	Ctrl+Shift+.

Figure: The SQL editor right click menu

The SQL editor is also used in the Bookmark Editor and when editing CLOBs in the form editor.

Database Connection, Catalog and Schema

You use the Database Connection and Database (or Catalog) lists above the editor to specify which connection and database to use when executing the SQL in the editor. The list of connections shows all connections as they are ordered in the Database Objects tree, except that all currently active connections are listed first.

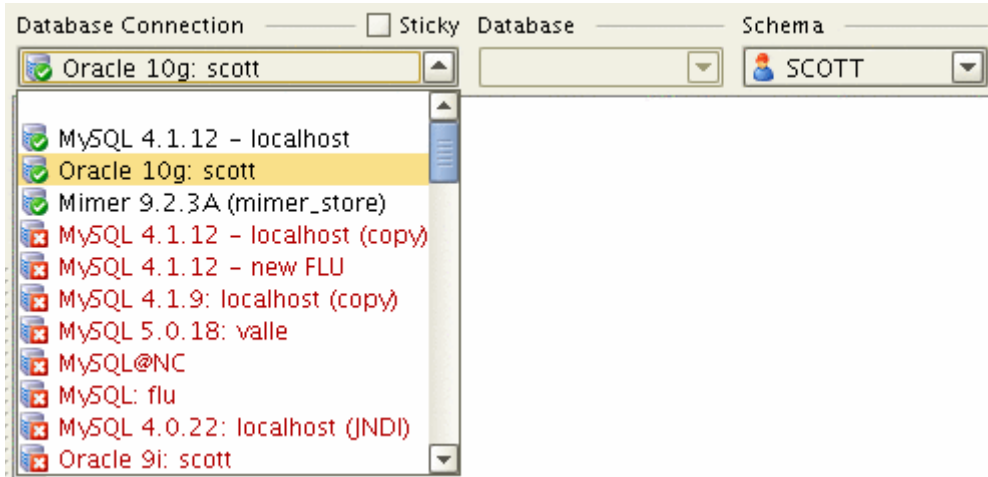


Figure: Database Connection, Database and Schema lists

If you check the Sticky box above the Database Connection, the current connection selection will not change automatically when passing SQL statements from other parts of DbVisualizer, for instance, when passing an SQL bookmark from the Bookmark Editor. Consider an SQL bookmark defined for database connection "ProdDB". If the Sticky checkbox is not checked (i.e., disabled), the database connection is automatically changed to ProdDB when you pass the SQL from the Bookmark Editor. However, if the Sticky checkbox is checked (i.e., enabled), the current database connection setting is unchanged. The Sticky setting is per SQL editor instance.

The Database list (or Catalog) defines which catalog in the connection is the target for the execution. Since not all databases use catalogs, this list may be disabled.

For most databases, the schema selected in the Schema list is used *only* to limit the tables the auto completion feature shows in the completion pop-up; it does not define a default schema for tables referenced in the SQL, because most databases do not allow the default schema to be changed during a session. For the databases that do allow the default schema to be changed, however, the selected schema is also used as the default schema, i.e., the schema used for unqualified table names in the SQL. Currently, the databases that support setting a default schema are DB2, JavaDB/Derby and Oracle. If you don't want the selected schema to be used as the default schema for these database, you can disable this behavior in the Tool Properties, under the database node's SQL Editor settings.

Limiting Result Set size (Max Rows/Chars)

The Max Rows field is used to control how many rows that DbVisualizer will fetch for each result set. If there are more rows available than presented in the result set, you will see a warning indicator in the grid status bar.

Note: Setting Max Rows technically means that it is the JDBC driver limiting the rows. This may for some databases also affect non result set operations such as DELETE. MS SQL Server is one example.

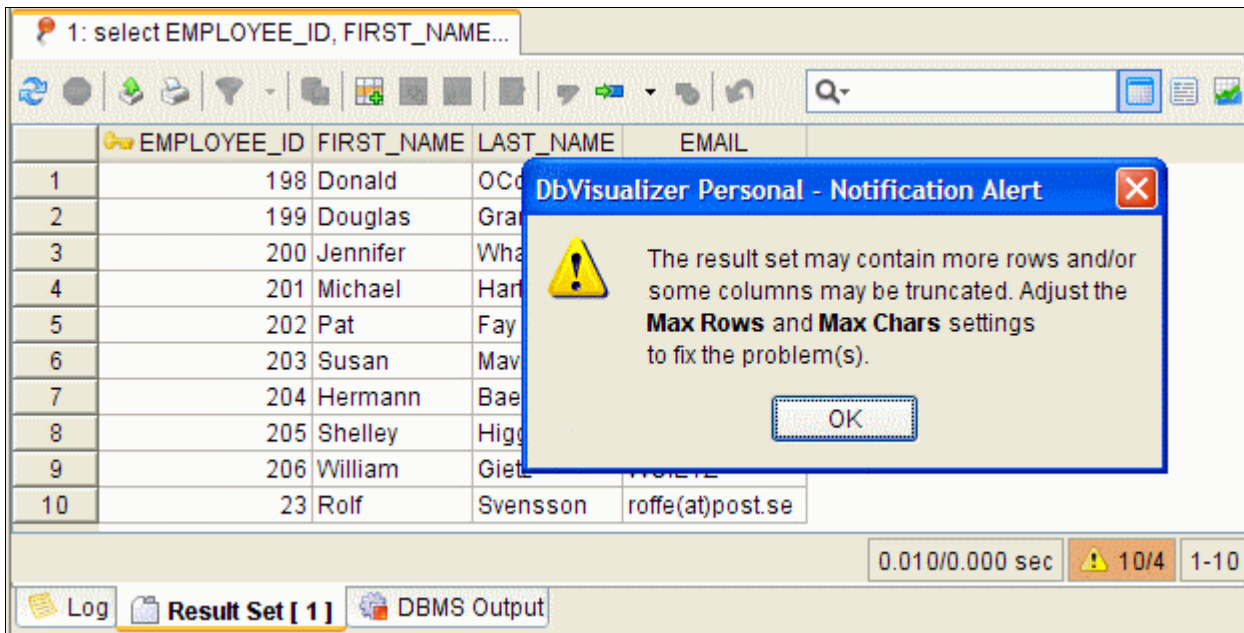


Figure: Max Rows exceeded warning

Click on the icon below the grid shows more information about the warning.

Setting Max Chars limits the number of characters that are presented for text data. A column that contains values with more characters than the specified Max Chars is shown with a different background color to highlight that it is truncated.

Load from and save to file

The SQL editor supports loading statements from a file and saving the content of the editor to a file. Use the standard file operations, Load, Save and Save As in the File main menu to accomplish this. Loading a file always loads it into the currently selected editor.

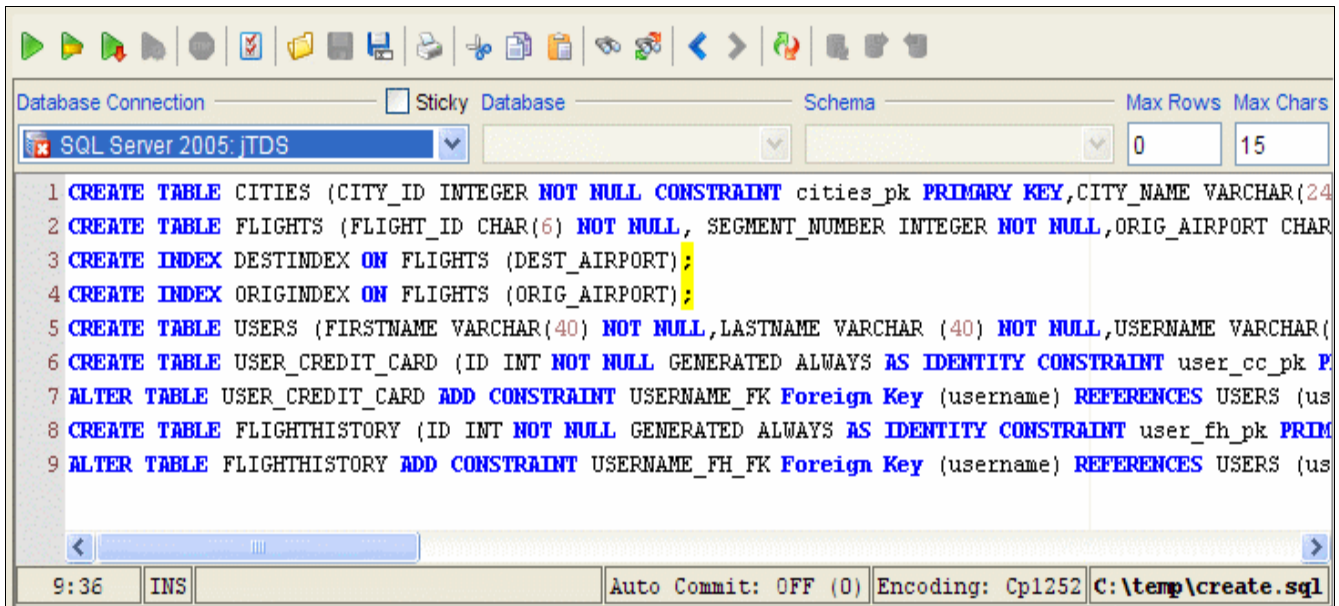


Figure: Loading a file into the SQL Commander

The name of the loaded file is listed in the status bar of the editor. The editor tracks any modifications and indicates changes with an asterisk (*) after the filename.

When you exist DbVisualizer, you are asked what to do if there are any pending edits that need to be saved.

Load Recent

The File->Load Recent sub menu lists the recently loaded files. When you choose an entry, a file chooser dialog is displayed with that file selected. The file chooser lets you to select the target SQL editor for the file:

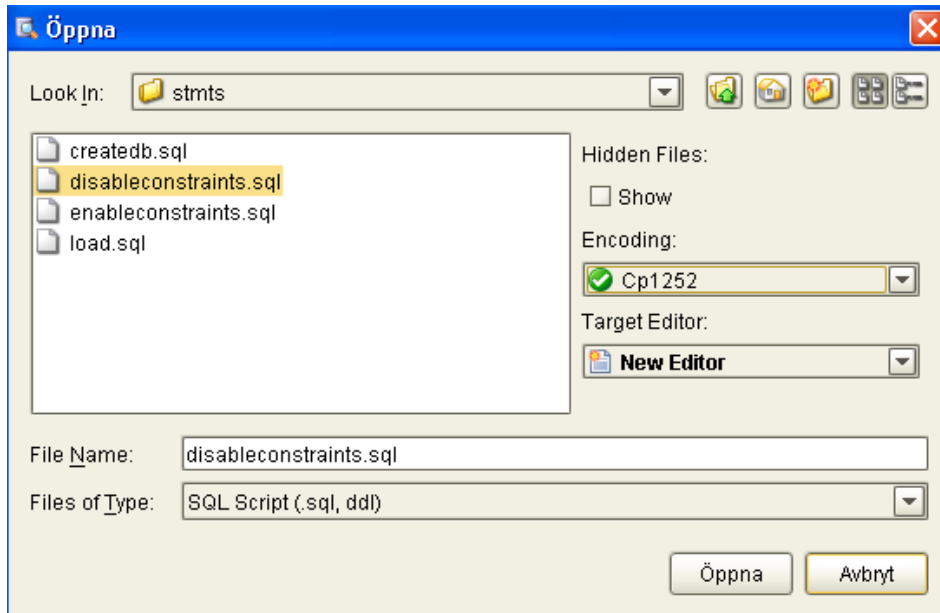


Figure: File Chooser

Editor Preferences

The Editor preferences pane is activated via the SQL->Show/Hide Editor Controls menu option. It keeps settings that control the appearance of the SQL editor, result sets and the log.

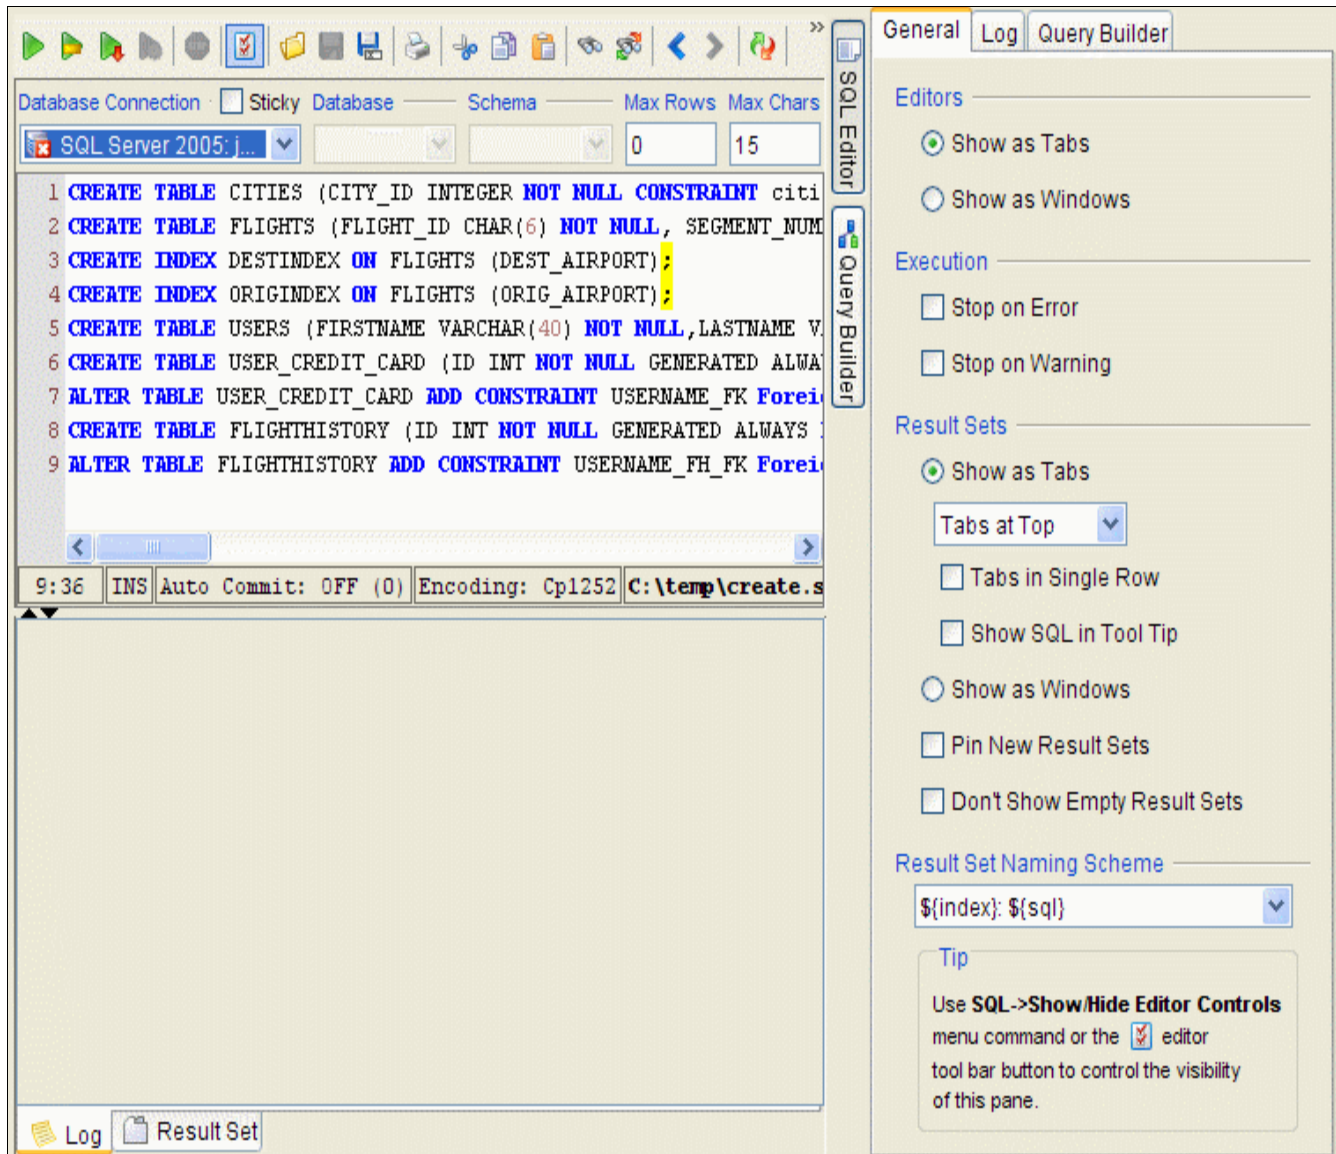


Figure: Editor preferences pane

All settings made in the editor preferences pane are saved between invocations.

Tip: The Result Set Naming Scheme may include HTML code, typically used to change the style of the elements.
Example: `<html>${index}: ${sql} (${rows})</html>`

Multiple editors

There is always one default editor named Main Editor. This editor is used when passing SQL bookmarks from the


Bookmarks Editor or when issuing requests from other parts of DbVisualizer that activate the SQL Commander. You can open additional SQL editors with the File->Create SQL Editor main menu operation. Editors can be organized as tabs or internal windows using the View buttons. To remove all but the Main Editor select the File->Close all SQL Editors menu operation.

Permissions

All SQL commands executed in the SQL Commanded are checked with the DbVisualizer Permission verifier before being executed by the database server. The permission verifier use various rules to determine if a specific SQL is allowed, denied or need confirmation before being executed. Specify in [Tool Properties->Permissions](#) the rules for the verifier.

Charsets and Fonts

You can change the SQL editor font, which is useful and necessary in order to display characters for languages like Chinese, Japanese, etc.



```
1 insert into RESTAURANTS (ID, NAME, SPECIALITY, STYLE, LOCATION) values (
2 1, 'すべてのための肉', 'ミートボール', '現代', '東京');
3
4 insert into RESTAURANTS (ID, NAME, SPECIALITY, STYLE, LOCATION) values (
5 2, 'どこでもすべてのための魚', '魚', '古典的', '北京');|
```

5:38 INS

Figure: SQL Editor with another font

Open [Tool Properties](#) and select the Font category to set the font for the SQL Editor. (It is a good idea to set the same font for both the SQL editor and the grid components).

Note: Displaying data correctly is not just a matter of setting the font, because the character encoding on the client side (in which DbVisualizer runs) and in the database server may not be compatible. There is experimental support to set encodings to accomplish proper conversation between different encodings. Please see the [Getting Started and General Overview](#) document for more information.

Key Bindings

The editor shortcuts, or key bindings, can be redefined in the Tool Properties Key Bindings category. Select the Editor Commands folder to browse all editor actions.

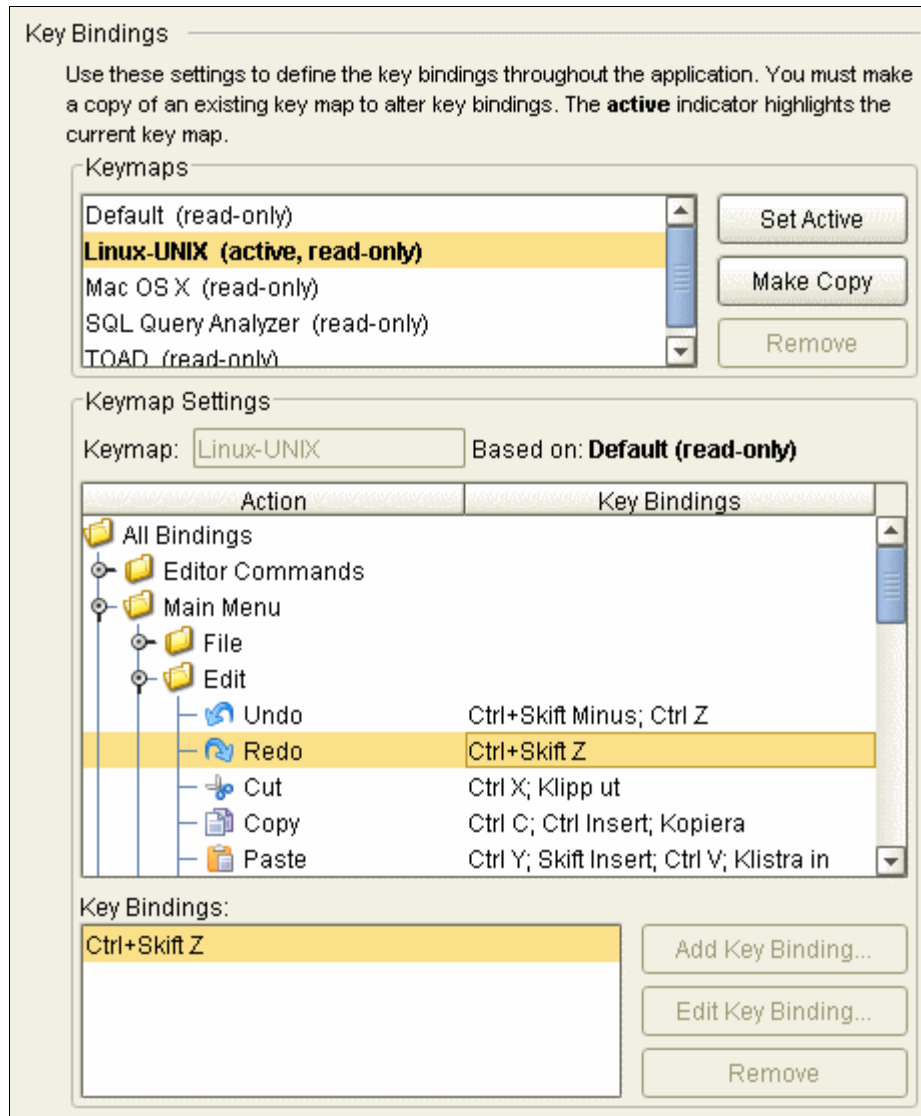


Figure: The Key Bindings editor in Tool Properties

Read more about configuring key bindings in the [Tool Properties](#) document.

Client side Comments

Comments in the SQL editor are identified by the comment identifiers in Tool Properties. These are client side comments and are removed by DbVisualizer before execution.

Sometimes the comments need to be passed to the database. Oracle, for example, uses the block comment identifier to express "hints" for the database. These must be passed to the database for processing. To enable

this, just change the delimiters for the block comment to something that doesn't interfere with the /*+ ... */ notation that Oracle uses.

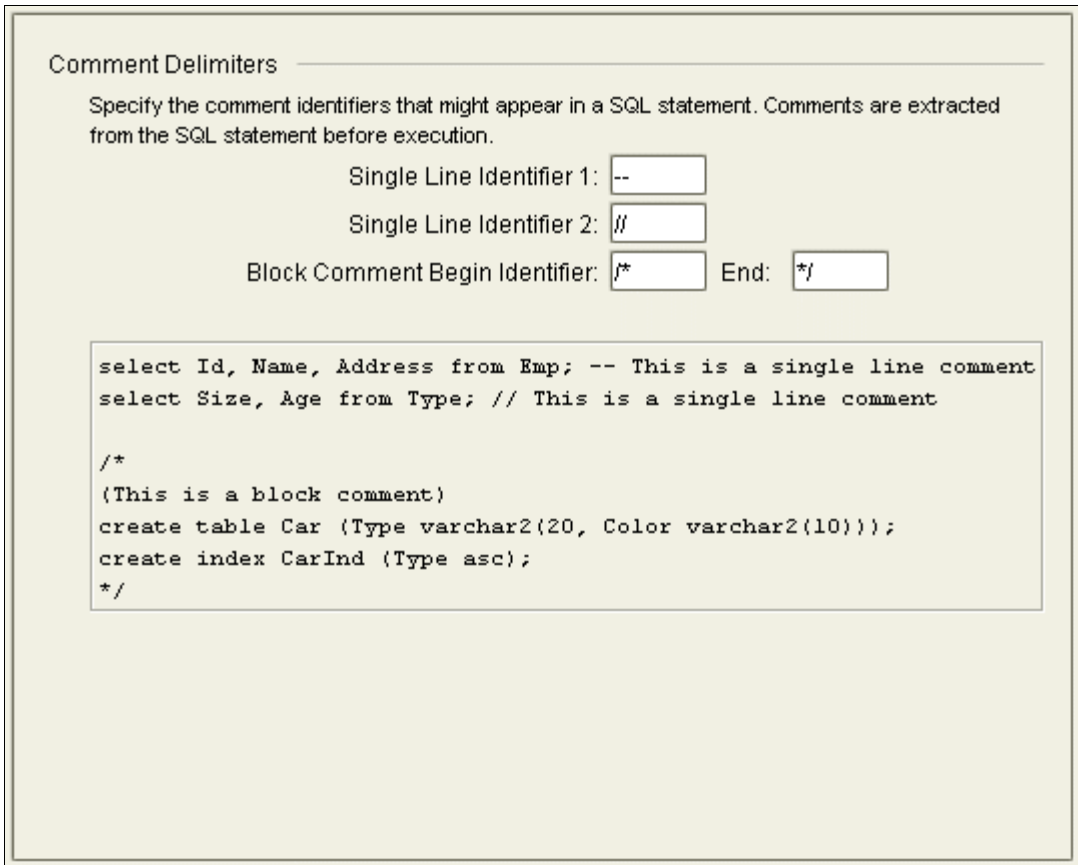


Figure: The Comments category in Tool Properties

Auto Completion

Auto completion is a convenient feature used to assist you when editing SQL statements.

The following figure shows the completion pop up with table names.

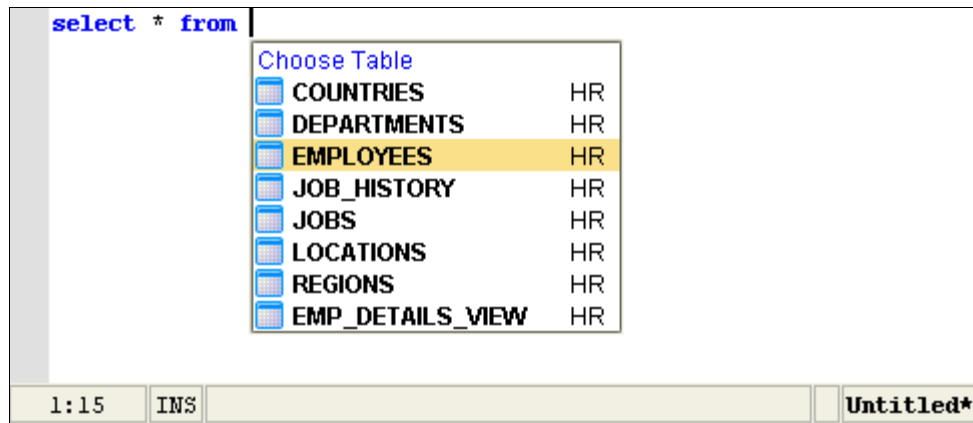


Figure: Auto completion pop up showing table names

Here is another completion pop-up showing column names.

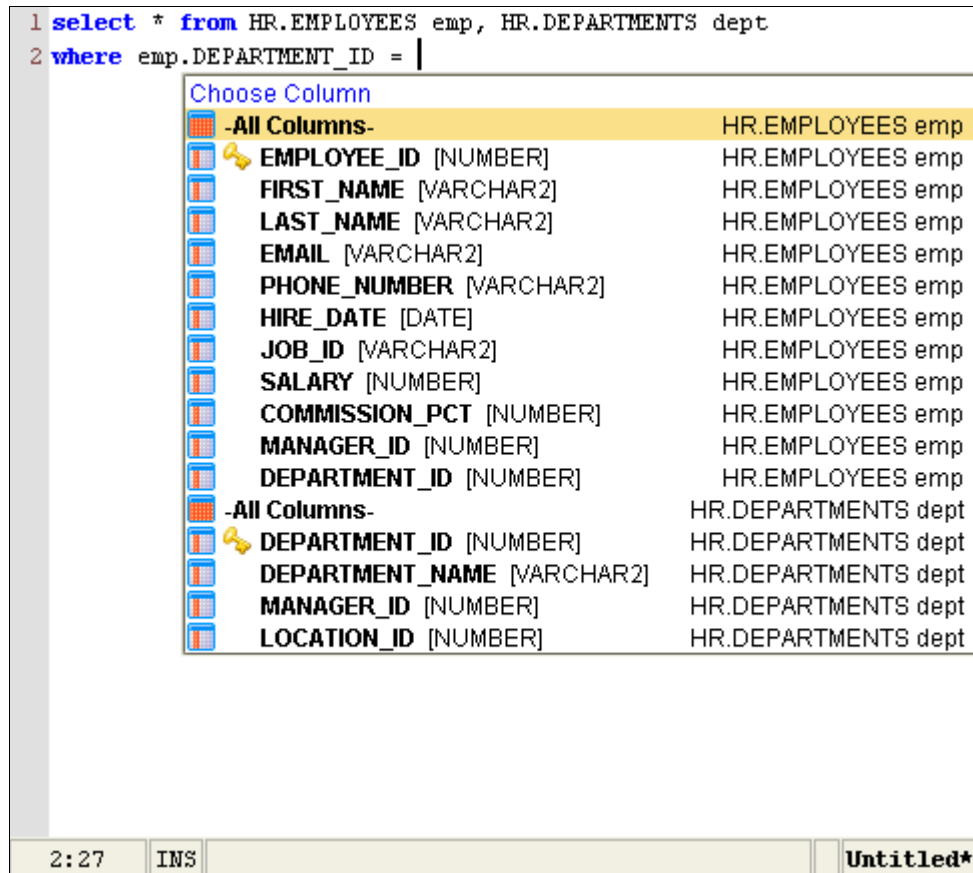


Figure: Auto completion pop up showing column names

DbVisualizer currently provides auto completion for table and columns names for the following DML commands:

- SELECT
- INSERT
- UPDATE
- DELETE

To display the completion pop-up, use the key binding Ctrl-SPACE. You select an entry in the pop-up menu with a mouse double-click, the ENTER key, or the TAB key. To cancel the pop-up, press the ESC key.

Tip: The SPACE key can be configured to select entries in the pop up. Do this in the Tool Properties General->Key Bindings category. Select the Editor Commands key bindings and add the SPACE key for the Insert Newline editor action.

Note 1: If there are several SQL statements in the editor then make sure to separate them using the statement delimiter character (default to ";").

Note 2: In order for the column name completion pop up to appear then you must first make sure there are table names in the statement.

Note 3: All table names that has been listed in the completion pop up are cached by DbVisualizer to make sure subsequent displays of the pop up is performed quickly without asking the database. The cache is cleared only when doing a Refresh in the database objects tree or reconnecting the database connection.

Note 4: The Schema list above the editor is used only to assist the auto completion feature to limit what tables to list in the pop up.

General display settings for the auto completion feature are managed in [Tool Properties](#).

Here are some examples of how the auto completion works depending on when it is activated. The <AC> symbol indicates the position where the auto completion pop-up is requested. The currently selected catalog is empty and the selected schema is HR. (These examples are when accessing an Oracle database).

<pre>select * from <AC></pre>	Shows all tables in the HR schema (since HR is the selected schema)
<pre>select * from SYS.<AC></pre>	The pop up displays all tables in the SYS schema independent of the schema list selection
<pre>select * from SYS.a<AC></pre>	Lists all tables in the SYS schema beginning with the A character
<pre>select <AC> from SYS.all_objects</pre>	Lists all column in the SYS.all_objects table
<pre>select <AC> from SYS.all_objects all, EMPLOYEES</pre>	Lists all columns in the SYS.all_objects and EMPLOYEES table (in the HR schema)
<pre>select emp.<AC> from EMPLOYEES emp</pre>	Lists all columns in the EMPLOYEES table, here identified by the alias emp
<pre>select emp.N<AC> from EMPLOYEES emp</pre>	Lists all columns in the EMPLOYEES table identified by alias emp starting with the N character
<pre>insert into EMPLOYEES (<AC></pre>	Lists all columns in the EMPLOYEES table. Selecting the -All Columns- in the pop-up results in all columns being added, comma separated.

It is possible to fine tune how auto completion shall work in the connection properties. The following settings can be used to adjust if table and column names should be qualified.

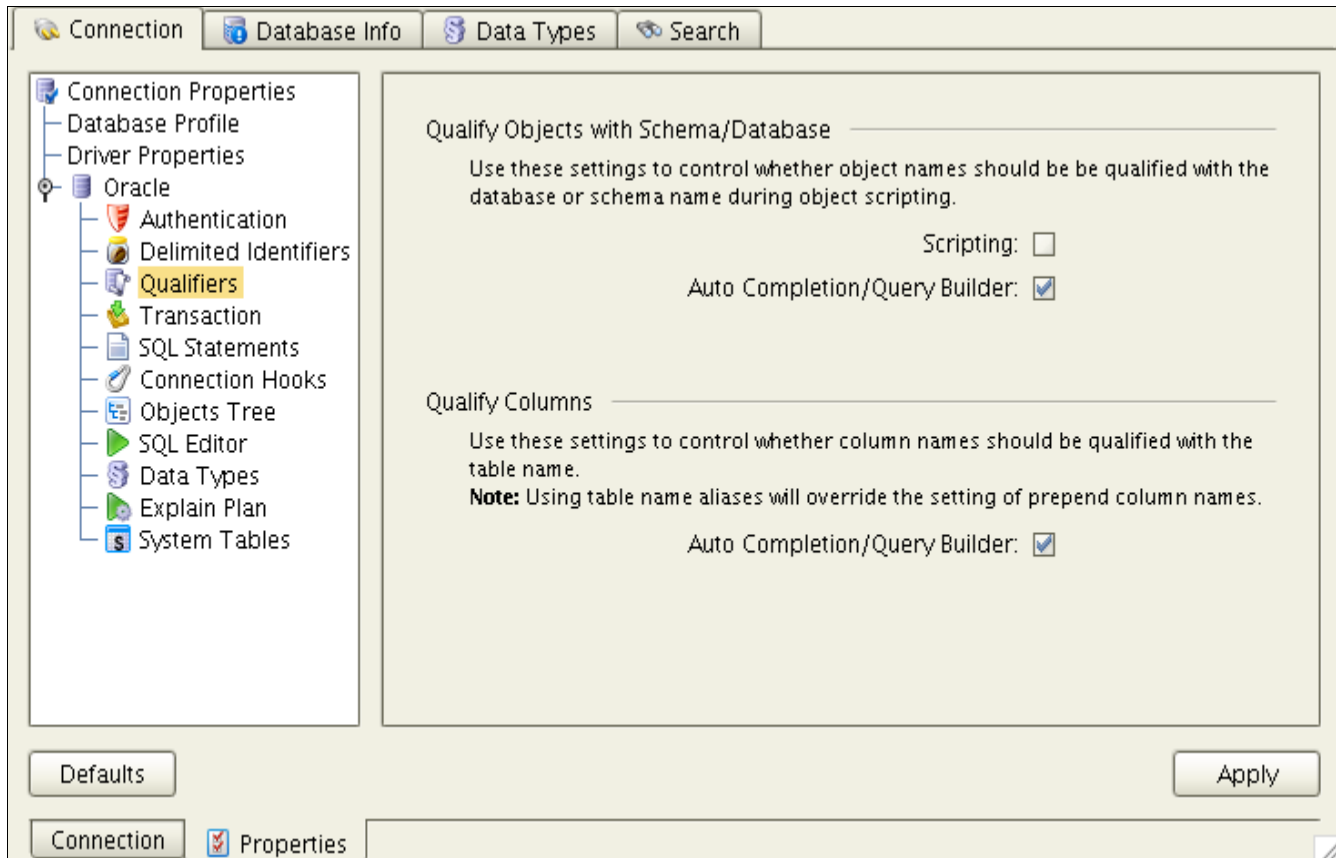


Figure: Properties controlling auto completion qualifiers

With Qualify disabled (for both table names and columns):

```
select Name, Address from EMPLOYEE where Id > 240
```

With Qualify enabled:

```
select EMPLOYEE.Name, EMPLOYEE.Address from HR.EMPLOYEE where EMPLOYEE.Id > 240
```

(The setting of Qualify Columns is ignored when an alias is used for a table name in the SQL).

The property settings in the figure below define whether delimited identifiers should be part of the completed SQL.

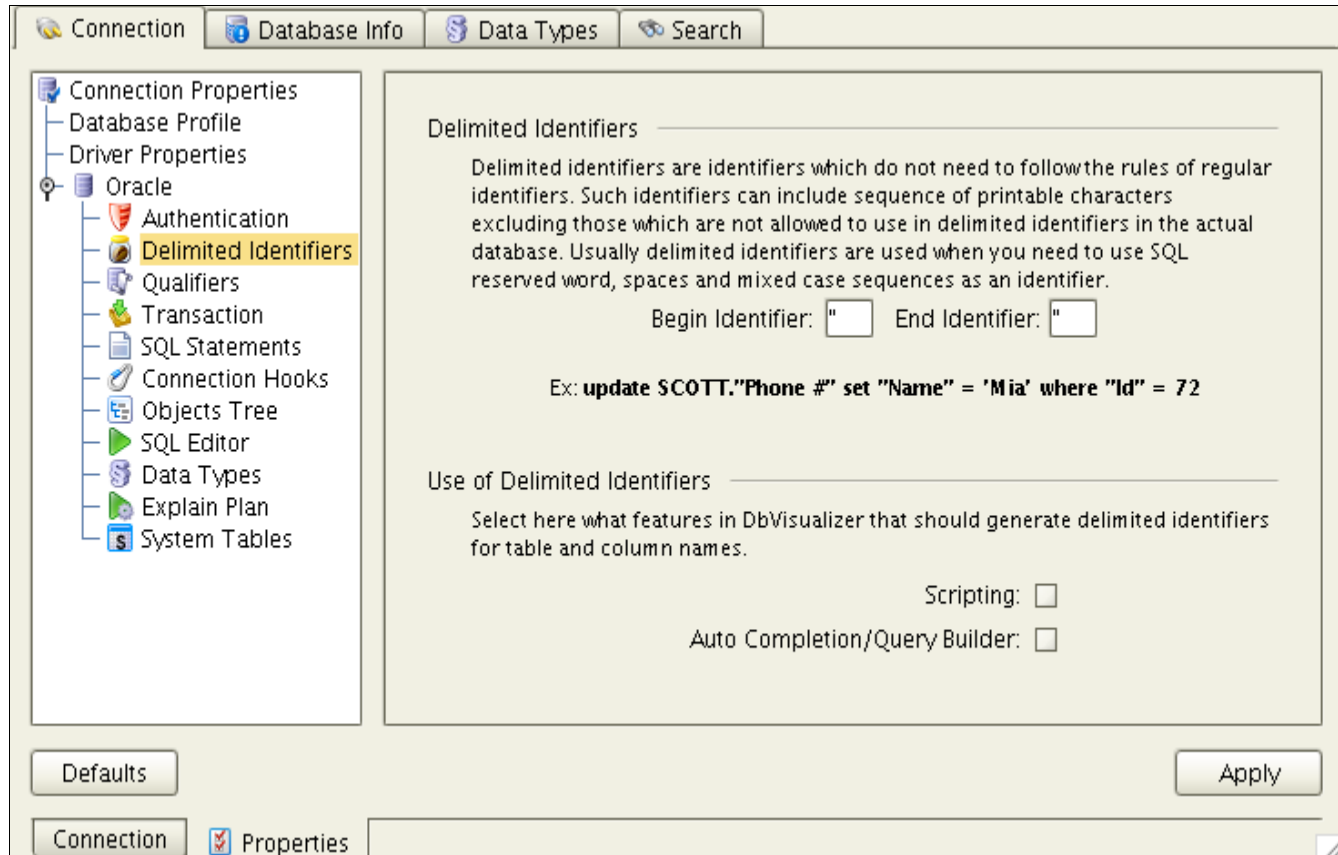


Figure: Properties controlling delimited identifiers for auto completion

With Delimited Identifiers disabled:

```
select Name, Address from HR.EMPLOYEE where Id > 240
```

With Delimited Identifiers enabled:

```
select "Name", "Address" from HR."EMPLOYEE" where "Id" > 240
```

SQL Formatter

The SQL->Format SQL feature is used to format the editor buffer according to the settings defined in the Tool Properties SQL Editor->SQL Formatting category. There are many things you can configure. After making some

changes, press Apply and format again to see the result. The formatter can work with the source SQL enclosed in quotes (e.g., when copied from a program), and it can format the final SQL for inclusion in a program written in languages like Java, C#, PHP, VB, etc.

Example of the SQL before formatting:

```
select CompanyName, ContactName, Address,
City, Country, PostalCode from
Northwind.dbo.Customers OuterC
where CustomerID in (select top 2 InnerC.CustomerId
from Northwind.dbo.[Order Details] OD
join Northwind.dbo.Orders O on OD.OrderId = O.OrderID
join Northwind.dbo.Customers InnerC
on O.CustomerID = InnerC.CustomerId
Where Region = OuterC.Region
group by Region, InnerC.CustomerId
order by sum(UnitPrice * Quantity * (1-Discunt)) desc)
order by Region
```

And after formatting has been applied:

```
SELECT
    CompanyName,
    ContactName,
    Address,
    City,
    Country,
    PostalCode
FROM
    Northwind.dbo.Customers OuterC
WHERE
    CustomerID in
    (
        SELECT
            top 2 InnerC.CustomerId
        FROM
            Northwind.dbo.[
        ORDER
            Details] OD
        JOIN
            Northwind.dbo.Orders O
            ON
            OD.OrderId = O.OrderID
        JOIN
            Northwind.dbo.Customers InnerC
            ON
            O.CustomerID = InnerC.CustomerId
        WHERE
            Region = OuterC.Region
        GROUP BY
            Region,
            InnerC.CustomerId
        ORDER BY
            sum(UnitPrice * Quantity * (1-Discunt)) desc
    )
ORDER BY
    Region
```

SQL History

The SQL Editor keeps track of all executed SQL statements. You can use the Previous and Next buttons in the editor toolbar to walk forward and backward through the statements. They insert the previously executed SQL, with accompanying settings for Database Connection, Catalog and Schema (if Sticky is disabled). The SQL main menu also contains operations for walking through the statement history.

The history entries are also managed as SQL Bookmarks, collected in the History root folder in the [Bookmark Editor](#).

SQL Bookmarks

SQL Bookmarks are used to manage favorite SQL statements between invocations of DbVisualizer. These are handled by the Bookmark Editor but the execution is performed in the SQL Commander. Please refer to the [SQL Bookmarks](#) document for how to use the Bookmarks main menu operations in the SQL Commander.

Execution

The execution of multiple SQL statements can be controlled using the Stop Execution On controls. These define whether the execution of the following SQL statements will be stopped based on two states:

- Errors
Stop the execution if the SQL resulted in an error
- Warnings
Stop the execution if the SQL executed successfully but no rows were affected

Note: The Stop Execution On controls are only effective when executing multiple SQL statements

SQL->Execute

Use the SQL->Execute main menu operation to execute the SQL in the current (selected) SQL editor. The SQL Commander executes the statements one by one and indicates the progress in the log area. The currently selected Database Connection is used for all statements. The SQL Commander does not support executing SQLs for multiple database connections in one batch.

The result of the execution is displayed in the output view based on what result(s) are returned. If there are several results and an error occurred in one of them, the [Log](#) view is automatically displayed to indicate the error.

If you select a statement in the SQL editor and choose SQL->Execute, only the selected statement is executed. This is a useful feature when you have several SQL statements in the SQL editor and you just want to execute one or a few of the statements.

```
1 select * from DBA.sales_order;
2 //select * from DBA.sales_order_items;
3 select * from DBA.contact;
4 select * from DBA.customer;
5 //select * from DBA.fin_code;
6 /* select * from DBA.fin_data;
7 select * from DBA.product; */
8 select * from DBA.department;
9 select * from DBA.employee;
```

Figure: Selection execute

In the above figure, only the highlighted statement is being executed.

SQL->Execute Current

The Execute Current operation is useful when you have a script with several SQL statements. It lets you execute the statement at the cursor position without first having to select the SQL statement. The default key binding for execute current is Ctrl-PERIOD (Ctrl-.).

Note: Execute Current determines the actual statement by parsing the editor buffer using the standard statement delimiters.

Tip: If you are unsure what the boundaries are for the current statement then use Edit->Select Current Statement. This will highlight the current statement without executing it.

SQL->Execute Buffer

Execute Buffer sends the complete editor buffer for execution as one statement. No comments are removed or parsing of individual statements based on any delimiters is made. This operation is useful when executing anonymous SQL blocks or SQLs used to create procedures, functions, etc.

SQL->Execute Explain Plan (Oracle, SQL Server and DB2)

Explain Plan is supported for Oracle, DB2 and SQL Server. Explain Plan executes your query and records the plan that the database devises to execute it. By examining this plan, you can find out if the database is picking the right indexes and joining your tables in the most efficient manner. The explain plan feature works much the same as executing SQLs to present result sets; you may highlight statements, run a script or load from file. The explain plan results can easily be compared by using the pin feature in combination with window style presentation.

DbVisualizer presents the plan either in a tree style format or in a graph. What information is shown depends on what database it is. In the tree view put the mouse pointer on the column header for a tooltip description what that column represents. The following screenshot shows the SQL in the editor at top and the resulting explain plan as result.

```

1 SELECT d."DEPARTMENT_NAME",
2       l."CITY",
3       c."COUNTRY_NAME",
4       r."REGION_NAME"
5 FROM HR."DEPARTMENTS" d,
6       hr."LOCATIONS" l,
7       hr."COUNTRIES" c,
8       hr."REGIONS" r
9 WHERE d."LOCATION_ID" = l."LOCATION_ID"
10      and l."COUNTRY_ID" = c."COUNTRY_ID"
11      and c."REGION_ID" = r."REGION_ID"
12      and d."MANAGER_ID" in
13      (SELECT "EMPLOYEE_ID"
14       FROM HR."EMPLOYEES"
15       WHERE "FIRST_NAME" like 'A%'
16      )

```

16:6 INS Untitled*

EXPLAIN 1: SELECT d."DEPARTMENT_NAME", l...

Tree View Graph View

Operation	Node Cost (%)	Cost	CPU Cost	I/O Cost	Optimizer	Cardinality
SELECT STATEMENT	0.0 %	6	363492	6	ALL_ROWS	3
NESTED LOOPS	0.0 %	6	363492	6		3
NESTED LOOPS	16.7 %	5	338468	5		3
NESTED LOOPS	0.0 %	3	244292	3		11
NESTED LOOPS	0.0 %	2	150072	2		23
HR.COUNTRY_C_ID_PK INDEX (FULL SCAN)	16.7 %	1	12121	1	ANALYZED	25
HR.LOCATIONS TABLE ACCESS (BY INDEX ROWID)	16.7 %	1	8871	1	ANALYZED	1
HR.LOC_COUNTRY_IX INDEX (RANGE SCAN)	0.0 %	0	1250	0	ANALYZED	2
HR.DEPARTMENTS TABLE ACCESS (BY INDEX ROWID)	16.7 %	1	9689	1	ANALYZED	1
HR.DEPT_LOCATION_IX INDEX (RANGE SCAN)	0.0 %	0	1650	0	ANALYZED	4
HR.EMPLOYEES TABLE ACCESS (BY INDEX ROWID)	16.7 %	1	8561	1	ANALYZED	1
HR.EMP_EMP_ID_PK INDEX (UNIQUE SCAN)	0.0 %	0	1050	0	ANALYZED	1
HR.REGIONS TABLE ACCESS (BY INDEX ROWID)	16.7 %	1	8341	1	ANALYZED	1
HR.REG_ID_PK INDEX (UNIQUE SCAN)	0.0 %	0	1050	0	ANALYZED	1

Figure: Explain Plan presented as a tree

The Graph View shows the plan as a graph. The graph can be exported to an image file or printed. Use the File menu choices to export and print.

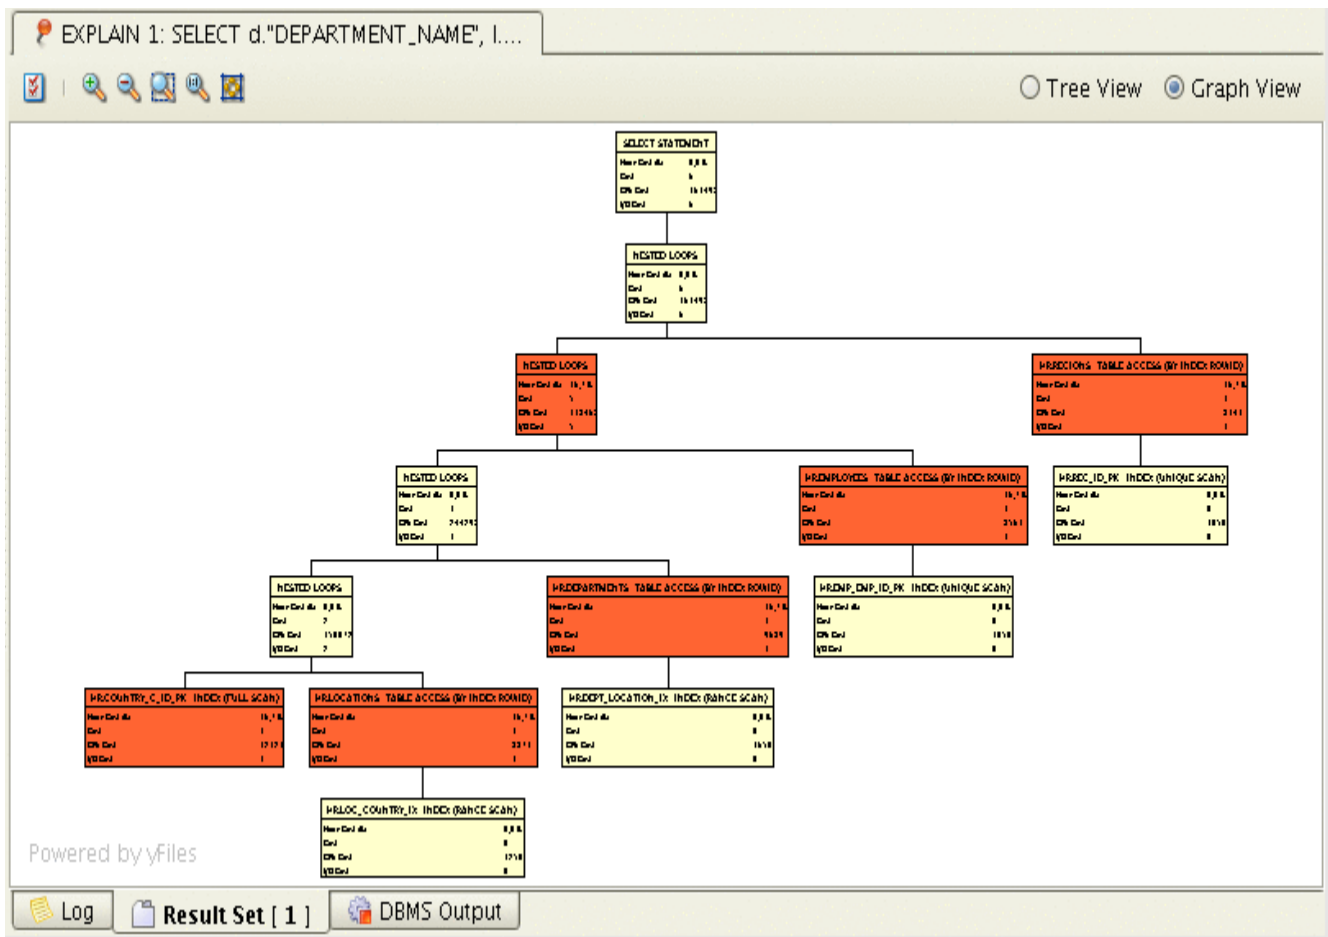


Figure: Explain Plan presented as a graph

Each of the supported databases use different techniques to manage their explain plan support. To control this, either click the Preferences toolbar button or go to Connection Properties->[database]->Explain Plan.

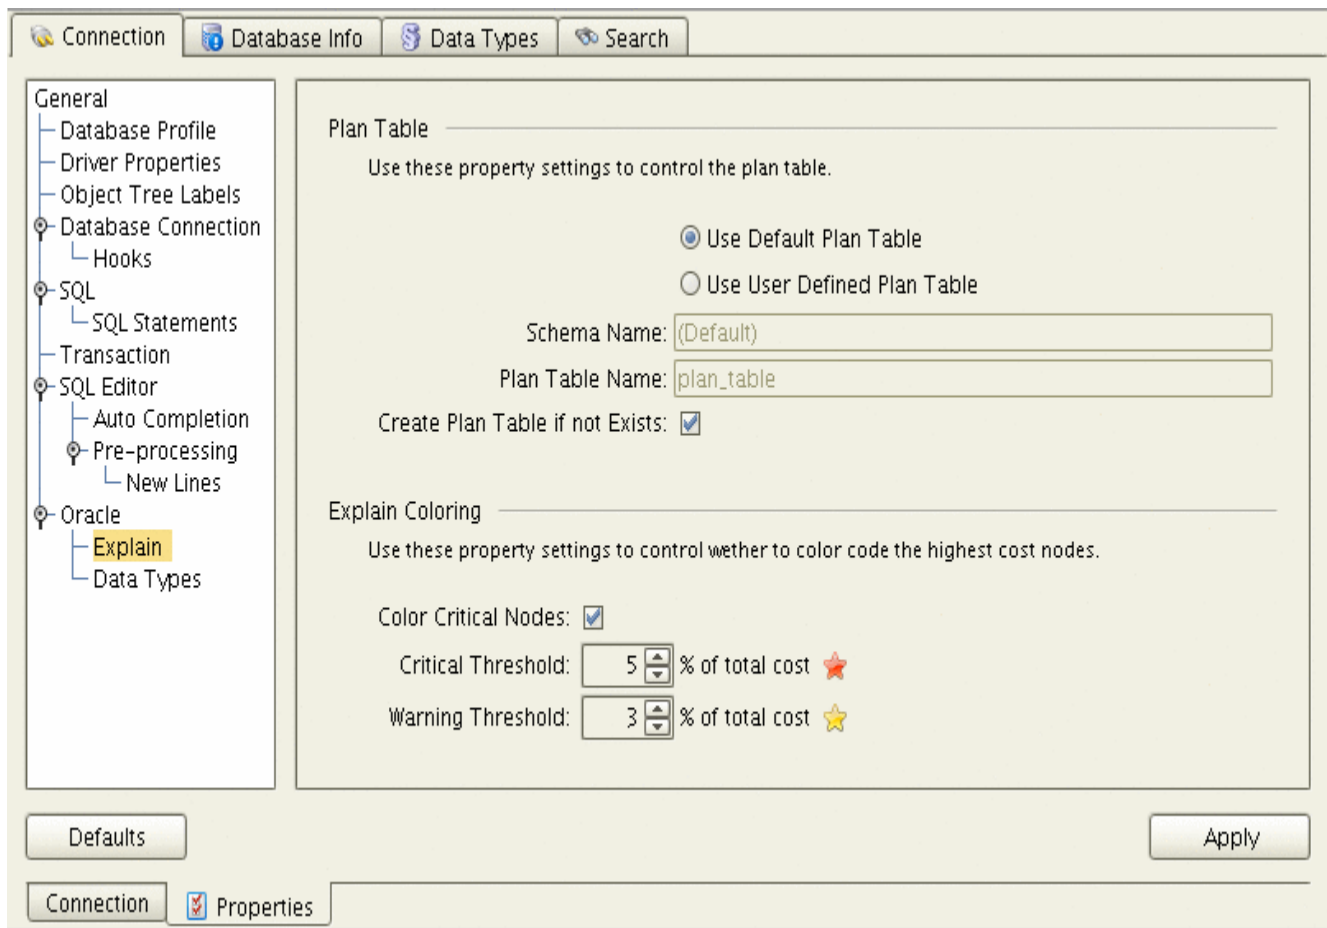


Figure: Explain Plan configuration

The configuration options are different for each of the supported databases.

Auto Commit, Commit and Rollback

The commit and rollback SQL commands and the accompanying operations in the Database main menu are enabled if the Auto Commit setting is off for the actual SQL editor. The default setting for Auto Commit is on, which means that the driver/database automatically commits each SQL that is executed. If Auto Commit is disabled, it is very important to manually issue the commit or rollback operations when appropriate.

The Auto Commit setting is enabled by default and can be adjusted in the Connection Properties.

SQL Scripts

An SQL script is composed of several SQL statements and can be executed in a batch. Each SQL statement is separated by a single character, a sequence of characters, or the word "go" on a single line. The default settings for the separator characters are defined in Tool Properties and can be modified to match your needs.

The screenshot shows a dialog box titled "SQL Statement Delimiter" with a description: "The character(s) that delimits one SQL statement from another during execution. The delimiter is ignored if found within apostrophes (' '), double quotes('\"'), in a single line comment or in a block comment. It is also ignored if found in a Variable." Below this are two input fields: "SQL Statement Delimiter 1:" with a text box containing ";" and "SQL Statement Delimiter 2:" with a text box containing ";".

The second section is titled "Allow 'go' as Delimiter" with a description: "Check to enable go as the only word on a line as an additional SQL statement delimiter." Below this is a checkbox labeled "Allow 'go' as Delimiter:" which is checked.

The third section is titled "Anonymous SQL Block Identifiers" with a description: "The character(s) used in the SQL Commander to identify the begin and end of an anonymous SQL block." Below this are two input fields: "Begin Identifier:" with a text box containing "--/" and "End Identifier:" with a text box containing "/".

Figure: Statement Delimiters

The following SQL script illustrates some uses of the SQL statement delimiters based on the settings in the previous figure:

```
select * from MyTable;                               /* Stmt 1 */
insert into table MyTable
(Id, Name) /* This is a comment */ values (1, 'Arnold') /* Stmt 2 */
go
update MyTable set Name = 'George' where Id = 1;     /* Stmt 3 */
select * from
MyTable; // This is a comment                       /* Stmt 4 */
```

Anonymous SQL blocks

An anonymous SQL block is a block of code which contains not only standard SQL but also proprietary code for a specific database. The anonymous SQL block support in the SQL Commander uses another technique in the JDBC driver to execute these blocks. The way you tell the SQL Commander know that a SQL block is to be

executed is to insert a begin identifier just before the block and an end identifier after the block. The figure in the previous section shows these settings and the default values:

```
Begin Identifier: --/  
End Identifier: /
```

Here follows an example of an anonymous SQL block for Oracle:

```
--/ script to disable foreign keys  
  
declare cursor tabs is select table_name, constraint_name  
from user_constraints where constraint_type = 'R' and owner = user;  
  
begin  
for j in tabs loop  
execute immediate ('alter table '||j.table_name||' disable constraint '||j.constraint_name);  
end loop;  
end;  
/
```

If you want to execute the complete editor buffer as an anonymous SQL block, use the SQL->Execute Buffer operation. In this case, you do not need the begin and end identifiers.

Stored Procedures

Executing stored procedures is not officially supported by DbVisualizer even though it works for some databases. The best way to figure it out is to test.

Our internal tests show that the Sybase ASE and SQL Server procedure calls work okay in the SQL Commander. DbVisualizer also presents multiple result sets from a single procedure call as of version 4.0 for these databases.

Client Side Commands

The SQL Commander supports a number of DbVisualizer specific editor commands. An editor command begins with the at sign, "@". The following sections describe the available commands.

@run - run SQL script from file

@cd <directory> - change directory

@<file> - run SQL script from file

Use the following commands to locate and execute SQL scripts directly from a file without first loading the file into the SQL editor. This is useful if you are using an external editor or a development environment to edit the SQL but use DbVisualizer to execute it.

- **@run <file>**
Request to execute the file specified as parameter
- **@cd <directory>**

Change the working directory for the following @run or @<file> commands

- @<file>
Same as @run <file>

Example of a script utilizing the file referencing commands:

```
select * from MyTable;      -- Selects data from MyTable
                             -- Execute the content of the
@run createDB.sql;         -- createDB.sql file. The location
                             -- of this file is the same as the working
                             -- directory for DbVisualizer.

@cd /home/mupp;           -- Request to change directory to /home/mupp

@loadBackup.sql;          -- Execute the content in the
                             -- loadBackup.sql file. This file will now
                             -- be read from the /home/mupp directory.
```

@export - export result sets to file

The @export commands are used to declare that any result sets from the SQL statements that follows should be written to a file instead of being presented in the DbVisualizer tool. This is really useful, since it enables dumping very large tables to a file for later processing or, for example, to perform backups. The following commands are used to control the export:

- @export on
Defines that the SQL statements that follows will be exported rather than being presented in DbVisualizer
- @export set parm1="value1" parm2="value2"
The set command is used to customize the export process. Check the table below for the complete set of parameters.
- @export off
Defines that SQL statements that follows will be handled the normal way, i.e., the result sets are presented in the DbVisualizer tool

These parameters are supported:

Parameter	Default Value	Valid Values
AppendFile	false	true, false, clear
BinaryFormat	Don't Export	Don't Export, Value, Hex, Base64
CsvColumnDelimiter	\t (TAB)	
CsvIncludeColumnHeader	true	true, false
CsvIncludeSqlCommand	false	true, false
CsvRowCommentIdentifier		
CsvRowDelimiter	\n	\n (UNIX/Linux/Mac OS X), \r\n (Windows)
DateFormat	yyyy-MM-dd	See valid formats in Tool Properties document
DecimalNumberFormat	Unformatted	See valid formats in Tool Properties document
Destination	File	File
Encoding	UTF-8	
Filename	REQUIRED	

Format	CSV	CSV, HTML, XML, SQL
HtmlIncludeSqlCommand	false	true, false
HtmlIntroText		
HtmlTitle	DbVisualizer export output	
NumberFormat	Unformatted	See valid formats in Tool Properties document
QuoteTextData	None (ANSI if Format="SQL")	None, Single, Double, ANSI
Settings		
ShowNullAs	(null)	
SqlIncludeSqlCommand	false	true, false
SqlRowCommentIdentifier	--	
SqlSeparator	;	
TimeFormat	HH:mm:ss	See valid formats in Tool Properties document
TimeStampFormat	yyyy-MM-dd HH:mm:ss.SSSSSS	See valid formats in Tool Properties document
XmlIncludeSqlCommand	false	true, false
XmlIntroText		

Example 1: @export with minimum setup

The following example shows the minimum commands to export a result set. The result set produced by the select * from Orders will be exported to the C:\Backups\Orders.csv file, using the default settings.

```
@export on;
@export set filename="c:\Backups\Orders.csv";

select * from Orders;
```

Example 2: @export with automatic table name to file name mapping

This example shows how to make the filename the same as the table name in the select statement. The example also shows several select statements. Each will be exported in the SQL format. Since the filename is defined to be automatically set, this means that there will be one file per result set and each file is named by the name of its table.

Note: There must be only one table name in a select statement in order to automatically set the filename, i.e if the select joins from several tables or pseudo tables are used, you must explicitly name the file.

```
@export on;
@export set filename="c:\Backups\${table}" format="sql";

select * from Orders;
select * from Products;
select * from Transactions;
```

Example 3: @export all result sets into a single file

This example shows how all result sets can be exported to a single file. The AppendFile parameter supports the following values.

- true
The following result sets will all be exported to a single file
- false
Turn off the append processing
- clear
Same as the true value but this will in addition clear the file before the first result set is exported

```
@export on;
@export set filename="c:\Backups\alltables.sql" appendfile="clear" format="sql";

select * from Orders;
select * from Products;
select * from Transactions;
```

Example 4: @export using predefined settings

The Export Grid Wizard lets you save export settings to a file for later use. Such an export settings file can be referenced in the @export set command.

```
@export on;
@export set settings="c:\exportsettings\htmlsettings.xml" filename="c:\Backups\${table}";

select * from Orders;
select * from Products;
select * from Transactions;
```

The example shows that all settings will be read from the c:\exportsettings\html.xml file.

@exit [nocheck] - Exit DbVisualizer

The @exit command is the same as selecting the File->Exit operation. This command can be useful if you start DbVisualizer using the -invisible, -sql and -execute program arguments. Having @exit last in the loaded SQL file forces DbVisualizer to exit once the script has been executed. The nocheck argument defines that no confirmation dialogs should be displayed during exit.

@echo - Echo text

The @echo command simply echos the supplied text in the output.

@window iconify - Iconify the main window

This command results in the main window being lowered (iconified).

@window restore - Raise the main window

This command results in the main window being raised (if iconified).

@desc table - Describe the columns in table

Use the @desc command to show column information for a table. For tables that are not in the current database or schema, you need to qualify the table name properly.

```
@desc table;  
@desc database.table;  
@desc schema.table;
```

@spool log - Save log to file

The @spool log command is used to save the log to file. (The log is not cleared after being saved).

```
@spool log mylog.txt
```

@stop on error - Stop execution if any error occurs

@stop on warning - Stop execution if any warning occurs

The @stop on error and warning can be used to control that the script processing should stop if any error or warning occurs. The corresponding @continue on xxx command is used to ignore any error or warning conditions.

```
@stop on error;  
@stop on warning;  
  
@continue on error;  
@continue on warning;
```

Parameterized queries (variables)

Variables can be used to build parameterized SQL statements. When you execute a statement, the SQL Commander looks for variables and prompt you to enter or replace the variable values. Variables are also used internally in DbVisualizer. The SQL templates that are listed in the Tool Properties SQL->SQL Statements category are used inside DbVisualizer in various situations. The difference compared to variables you enter in the SQL editor is that DbVisualizer automatically substitutes the predefined variable names with the correct values when the templates are used instead of prompting you for values.

A variable has the following format in its simplest form:

```
$$FullName$$
```

A variable must begin and end with the character(s) identified by the Variable Identifier property in the Tool Properties SQL category (default is \$\$, as in the example above). During execution, the SQL Commander searches for variables and displays a window with the name of each variable and an input (value) field. Enter the value for each variable and then press Execute. This replaces the variable with the value and finally lets the database execute the statement.

Tip: Use the Ctrl->Enter key binding as a shortcut for Execute.

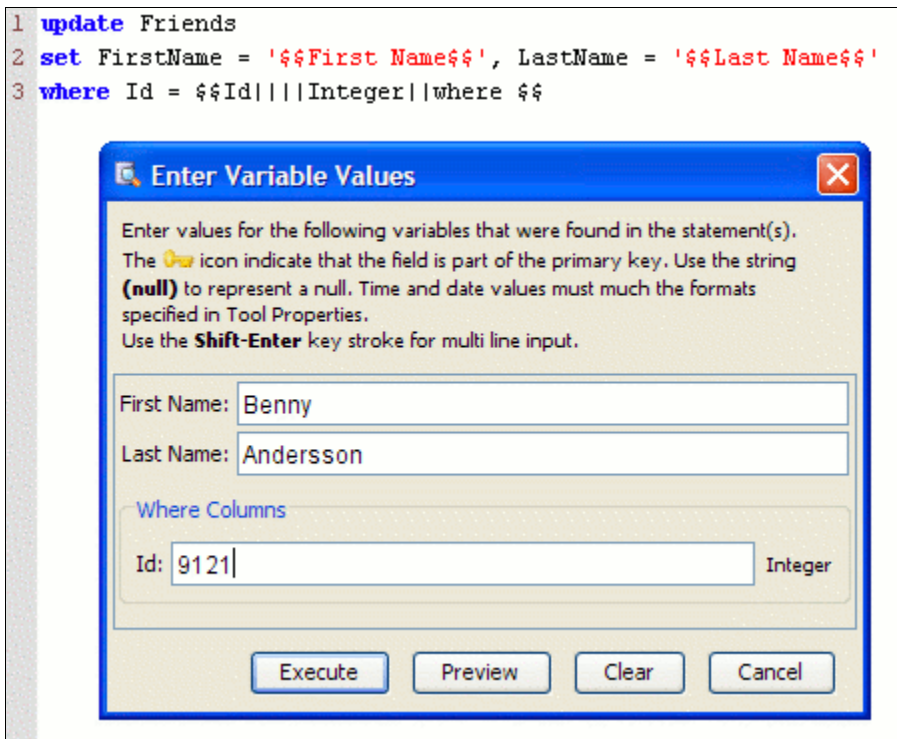


Figure: The substitute variables window

The above example is the simplest case, as it only contains the variable name. In this case it is also necessary to place the text value within quotes since the substitution window cannot determine from the variable itself if it is a number or text variable.

The final substituted SQL statement that results from the initial SQL and variable value is:

```
update Friends set FirstName = 'Benny', LastName = 'Andersson' where Id = 9121;
```

Variable Syntax

The variable format supports setting a default value, data type and a few options as in the following example:

```
$$FullName||Andersson||String||where pk $$
```

The full format of the variable syntax is:

```
$$variableName [|| defaultValue [|| type [|| options]]]
```

- **variableName**
Required. This is the name that appears in the substitution dialog. If several variables have the same name, the substitution dialog shows only one and the entered value will be applied to all variables of that name.
- **defaultValue**
The default value that appears in the substitution dialog
- **type**
The type of variable - String, Integer, BinaryData, etc. This is used to determine if the value should be enclosed by quotes or not. If no type is specified, it is treated as an Integer (no quotes).
- **options**
The options part is used to express various things. Most interesting are the pk and where keywords. (Note: There must be a whitespace character following a keyword).

- pk
Declares that an icon should appear before the variable name in the substitution dialog to indicate that it is a primary key field.
- where
Declares that the variable is part of the WHERE clause and so should appear last in the list of variables.

Output View

The Output View in the lower area of the SQL Commander is used to display the result of the SQLs being executed. How the results are presented is based on what type of result it is. A log entry is always produced in the Log view for each SQL statement that is executed. This entry shows at a minimum the execution time and how many rows were affected by the SQL. There may also be a result set if the SQL returned one. These result sets are presented either as tabs or windows based on your choice.

The screenshot shows the SQL Commander interface with three tabs: '1: select * from DEPARTMENTS', '2: select EMPLOYEE_ID, FIRST_NAME...', and '3: select * from JOB_HISTORY'. The '3: select * from JOB_HISTORY' tab is active, displaying a table with 8 rows. The table has columns: EMPLOYEE_ID, START_DATE, END_DATE, JOB_ID, and DEPARTMENT_ID. Below the table, there is a status bar showing '0.010/0.010 sec', '10/5', and '1-9'. At the bottom, there are tabs for 'Log', 'Result Set [3]', and 'DBMS Output'.

	EMPLOYEE_ID	START_DATE	END_DATE	JOB_ID	DEPARTMENT_ID
1	102	1993-01-13 00:00:00.0	1998-07-24 00:00:00.0	IT_PROG	60
2	101	1989-09-21 00:00:00.0	1993-10-27 00:00:00.0	AC_ACCOUNT	110
3	101	1993-10-28 00:00:00.0	1997-03-15 00:00:00.0	AC_MGR	110
4	201	1996-02-17 00:00:00.0	1999-12-19 00:00:00.0	MK_REP	20
5	114	1998-03-24 00:00:00.0	1999-12-31 00:00:00.0	ST_CLERK	50
6	122	1999-01-01 00:00:00.0	1999-12-31 00:00:00.0	ST_CLERK	50
7	200	1987-09-17 00:00:00.0	1993-06-17 00:00:00.0	AD_ASST	90
8	176	1998-03-24 00:00:00.0	1998-12-31 00:00:00.0	SA_REP	80

Figure: The output view

If an error occurs during execution, the SQL Commander automatically switches to the Log view so that you can further analyze the problem.

Log

The log keeps an entry for each SQL statement that has been executed. It provides generic information, such as how many rows were affected and the execution time. The important piece of information is the execution message which shows how the execution of that specific statement ended. If an error occurred, the complete log entry will be in red, indicating that something went wrong.




Figure: The Log with one failed statement

The detail level in an error message is dependent on the driver and database that is being used. Some databases are very good at telling what went wrong and why, while others are very quiet. The icon to the left of each log entry is used to pass the SQL for the entry into the current SQL editor when clicked.

Log controls

The Editor Control area contains a Log tab where you can control the log content. Use the Show controls to define which information you want to appear in the log. The Filter controls are used to specify which entries should be displayed.

Auto clear log

If you enable the Auto Clear Log control, the SQL Commander automatically clears the log between executions.

Result Set

A result set grid is created for every SQL that returns one or more result sets. These grids can be displayed in a tab or window style view, similar to how the SQL editors are displayed. Each grid shares the common layout and features as described in the [Getting Started and General Overview](#) document. The format of the result can be one:

- Grid
The result is presented in a grid.
- Text
The result is presented in a tabular format.
- Chart
Read more in [Monitor and Charts](#).

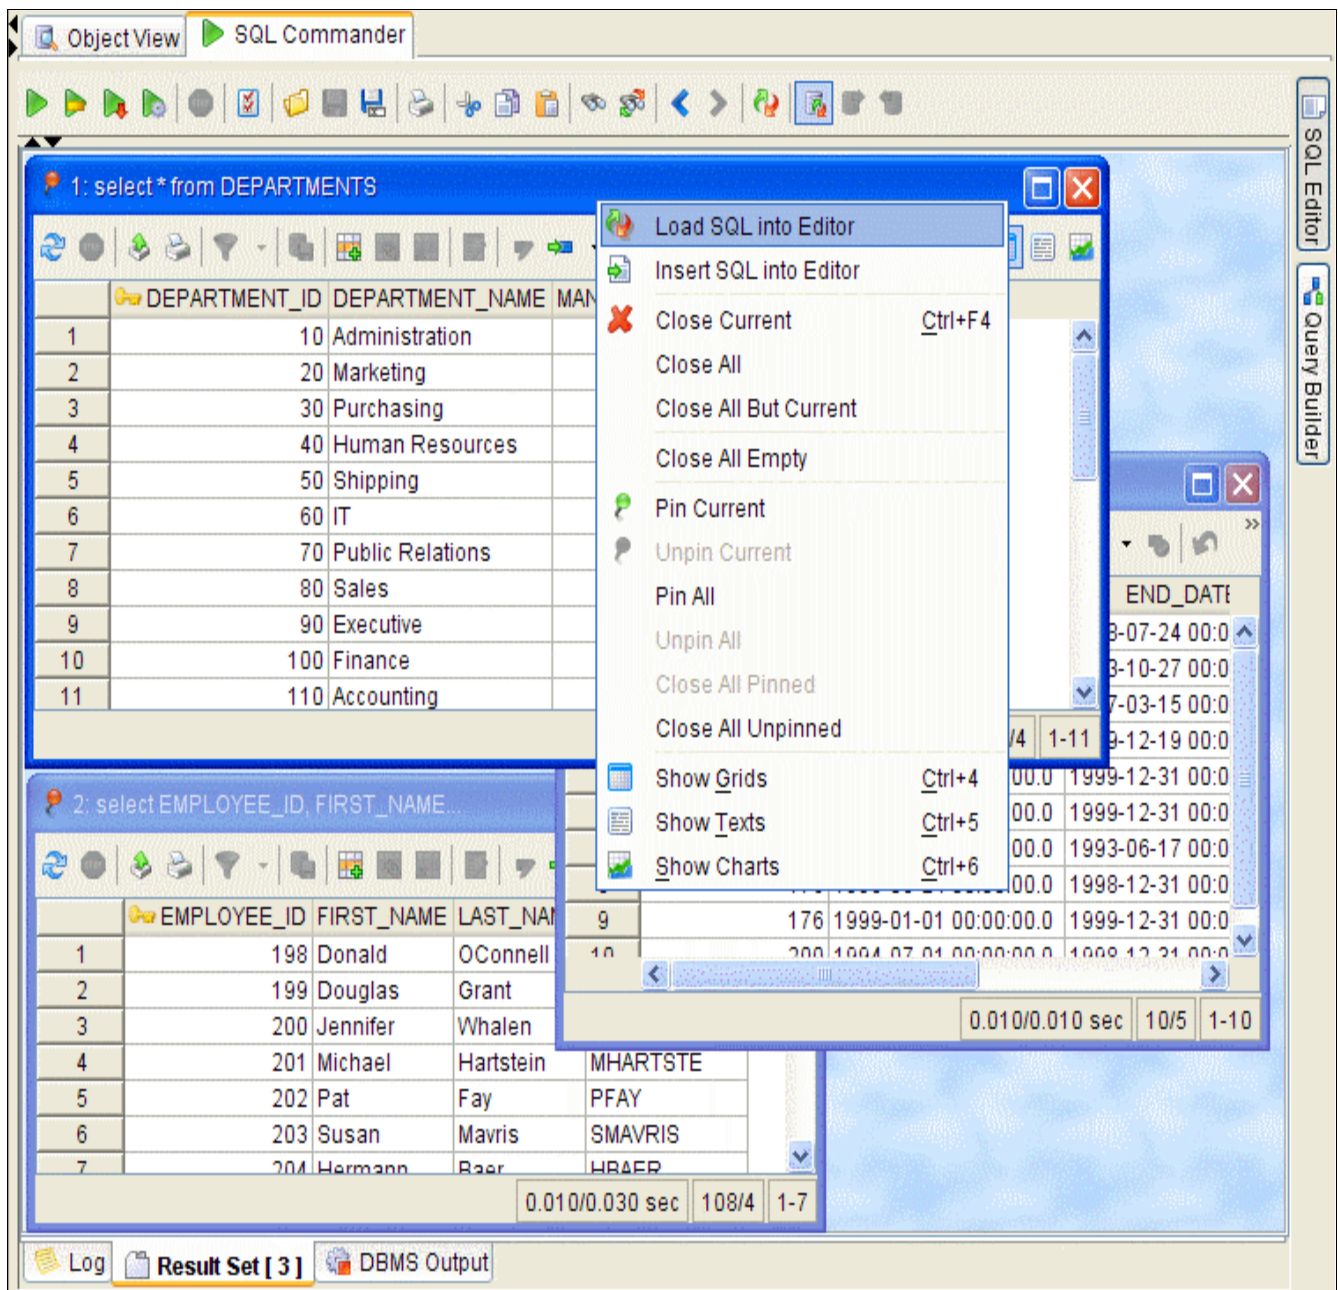


Figure: The windows output view

The figure above shows the windows output style with three result set grids. A result set grid can be closed using the red cross in the window frame header.

With the tabs style, you use the Close right click menu choice when the mouse pointer is in the tab header to close a result set:

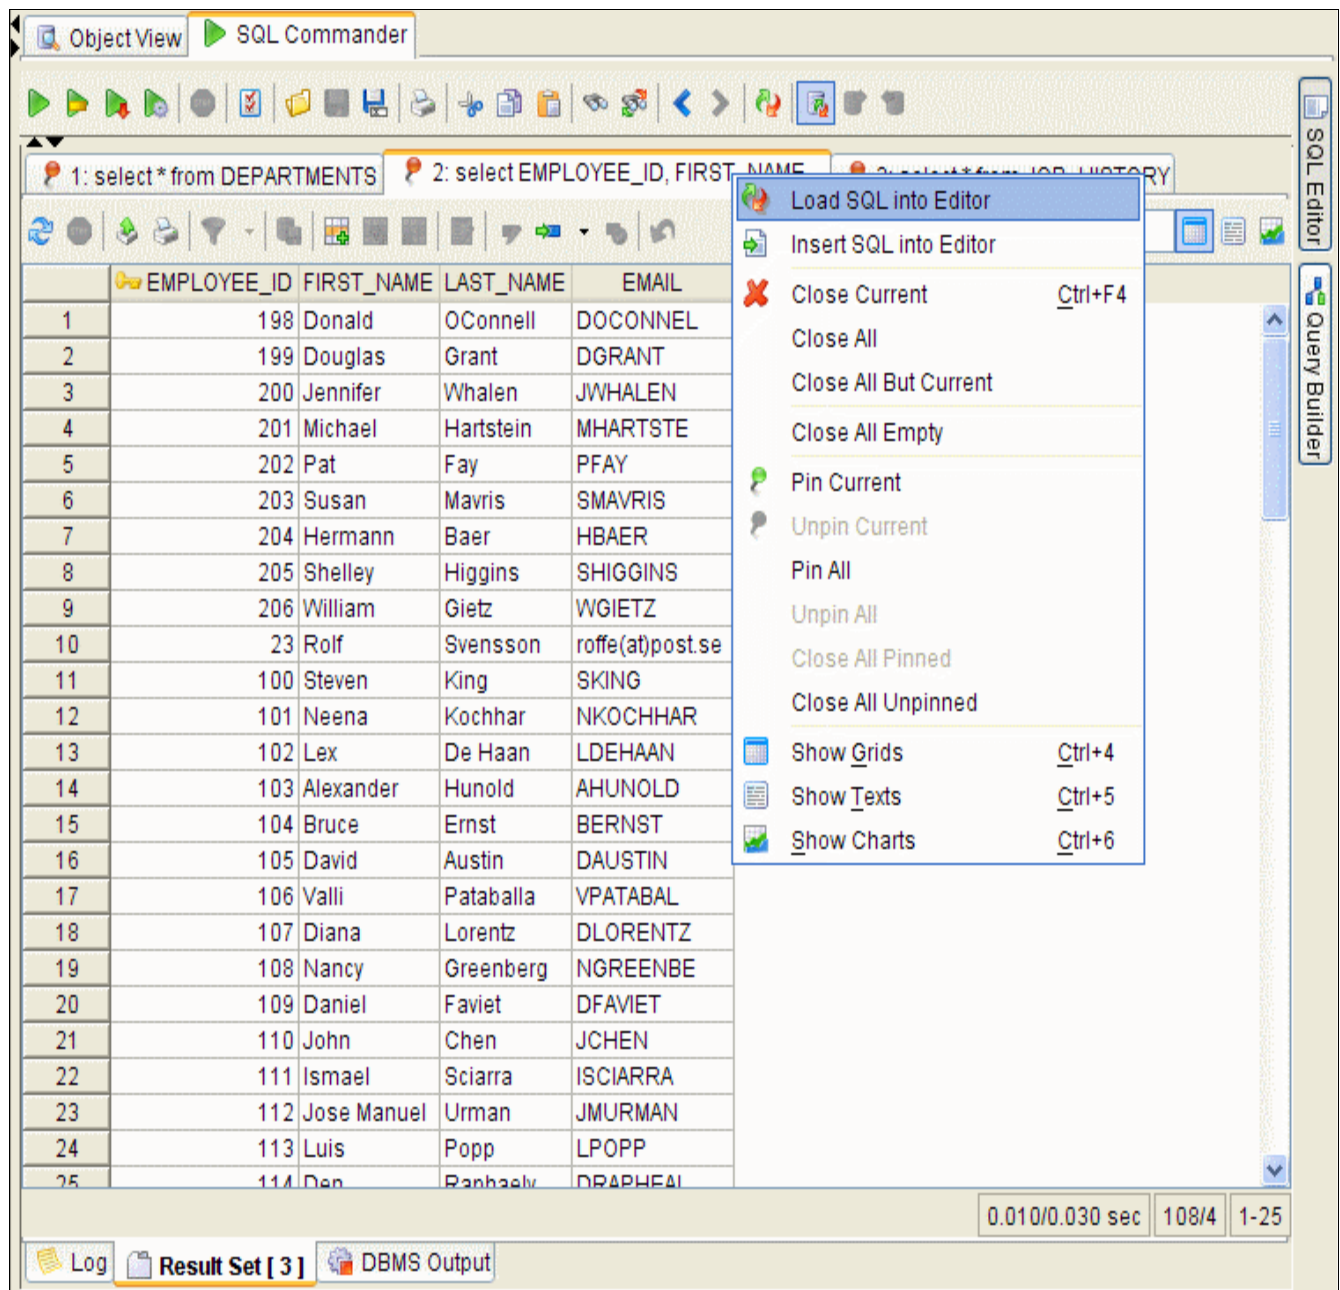


Figure: The right click menu for tabs

Result set menu

The result set menu is available by right-clicking on a tab or on the result set desktop (window style). It contains options to control the current result set and all result sets. The following actions are available:

Menu Choice	Description
Load SQL into Editor	Loads the SQL for the selected result set tab or window into the current editor.
Insert SQL into Editor	Inserts the SQL for the selected result set tab or window into the current editor at the cursor position.
Close Current	Closes the current result set
Close All	Closes all result sets
Close All But Current	Closes all but the current result set
Close All Empty	Closes all result sets that are empty (no data)
Pin Current	Pins the current result set (preventing it from being removed at the next execution).
Unpin Current	Unpins the current result set
Pin All	Pins all result sets. Pinning a result set prevents it from being removed at the next execution.
Unpin All	Unpins any pinned result sets, making them candidates for removal during the next execution.
Close All Pinned	Removes all pinned result sets directly.
Close All Unpinned	Removes all unpinned result sets directly.
Show Grids	Changes the display mode to show the grid tab for all result sets
Show Texts	Changes the display mode to show the text tab for all result sets
Show Charts	Changes the display mode to show the chart tab for all result sets

Editing

A result set grid may be enabled for editing based on the following criteria:

1. The result really is a result set
2. The SQL is a SELECT command
3. Only one table is referenced in the FROM clause
4. All columns in the result set exist in the table with exactly the same names

If all of the above is true, the standard editing tool bar appears just above the grid. Read more about editing in the [Edit Table Data](#) chapter.

Multiple result sets produced by a single SQL statement

Some SQL statements may produce multiple result sets. Examples of this are stored procedures in Sybase ASE and SQL Server. The SQL Commander checks the results as returned by the JDBC driver and add grids to the output view accordingly. The following shows the `sp_help Emps` command which returns several result sets with various information about the Emps table.

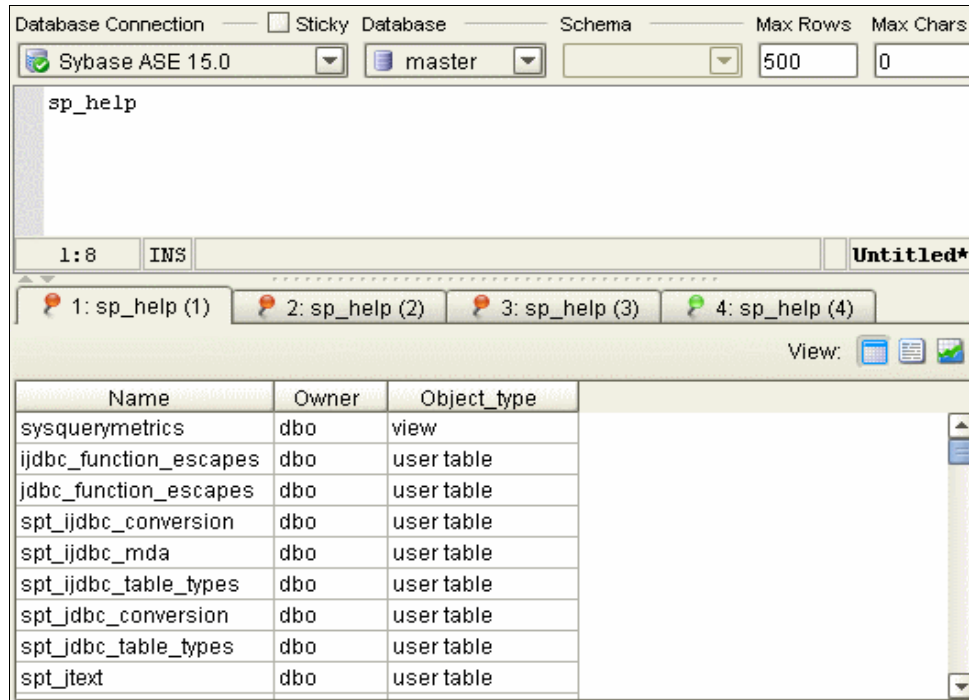


Figure: Multiple result set grids produced by a single SQL statement

The result set grids in the figure above all share the same label, `sp_help Emps`. The number after the label represents the order number for the actual result. A stored procedure can return different results, not all being result sets. The number helps you identify which entry matches which result set grid in the log. Here is the Log output view for the previous example.

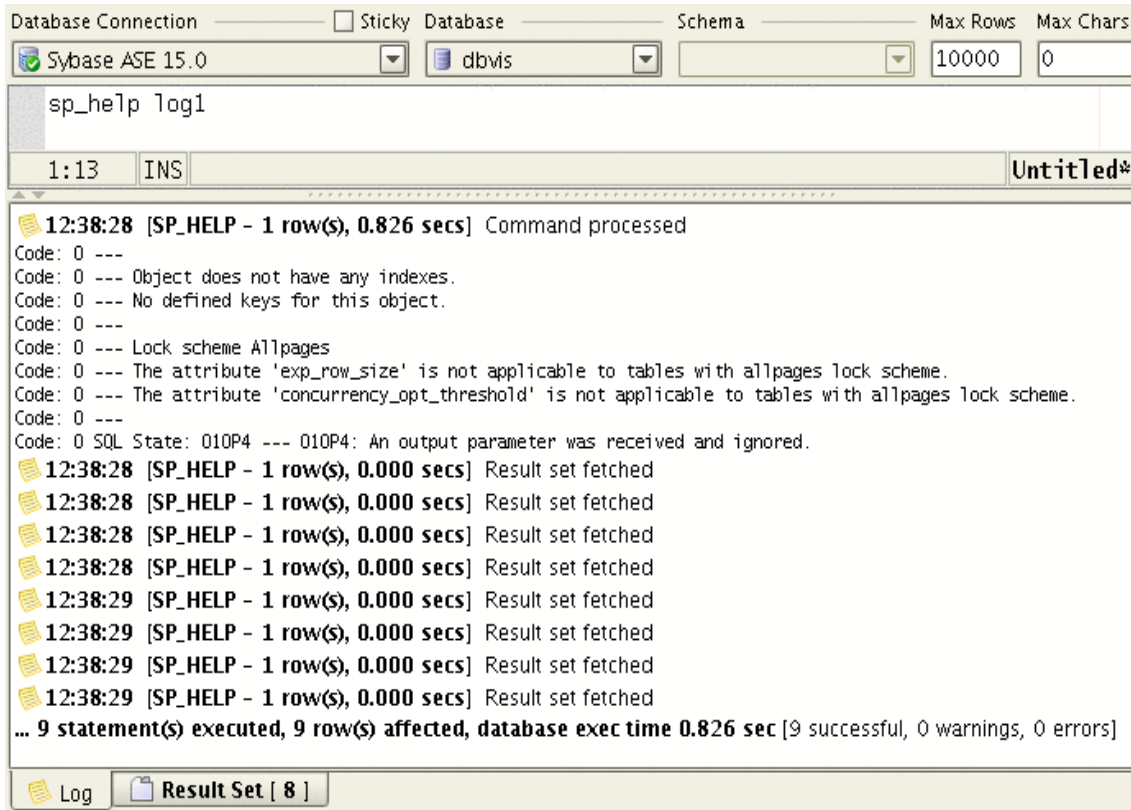



Figure: The Log after executing an SQL statement that returns multiple results

All entries with the log message "Result set fetched" are represented in the previous figure.

Text

The Text format for a result set presents the data in a tabular style. The column widths are calculated based on the length of each value and the length of the column label.

Note: The column widths may vary between executions of the SQL.



5: select EMPLOYEE_ID, FIRST_NAME...

View: [Grid] [Text] [Chart]

EMPLOYEE_ID	FIRST_NAME	LAST_NAME	EMAIL	PHONE_NUMBER	HIRE_DATE
198	Donald	OConnell	DOCONNEL	650.507.9833	1999-0
199	Douglas	Grant	DGRANT	650.507.9844	2000-0
200	Jennifer	Whalen	JWHALEN	515.123.4444	1987-0
201	Michael	Hartstein	MHARTSTE	515.123.5555	1996-0
202	Pat	Fay	PFAY	603.123.6666	1997-0
203	Susan	Mavris	SMAVRIS	515.123.7777	1994-0
204	Hermann	Baer	HBAER	515.123.8888	1994-0
205	Shelley	Higgins	SHIGGINS	515.123.8080	1994-0
206	William	Gietz	WGIEZT	515.123.8181	1994-0
23	Rolf	Svensson	roffe(at)post.se	87	2006-0
100	Steven	King	SKING	515.123.4567	1987-0

0.040 sec/0.060 sec | 108 / 11

Figure: The Text result set format

Chart

A result set can be charted using the Chart view in a grid. Please read more about it in the [Monitor and Charts](#) document.

DBMS Output (Oracle)

The DBMS Output tab for Oracle is used to enable and disable capturing of messages produced by stored procedures, packages, and triggers. These messages are typically inserted in the code for debugging purposes. For SQL*Plus users, the corresponding feature is enabled via the `set serveroutput on` command. To enable display of DBMS messages in DbVisualizer, select the DBMS Output tab and press the Enable button.

Once DBMS output is enabled, the icon in the tab header is changed. Invoking a stored procedure in the SQL editor will result in the following being displayed in the output tab. (Each block of output is separated with a timestamp).

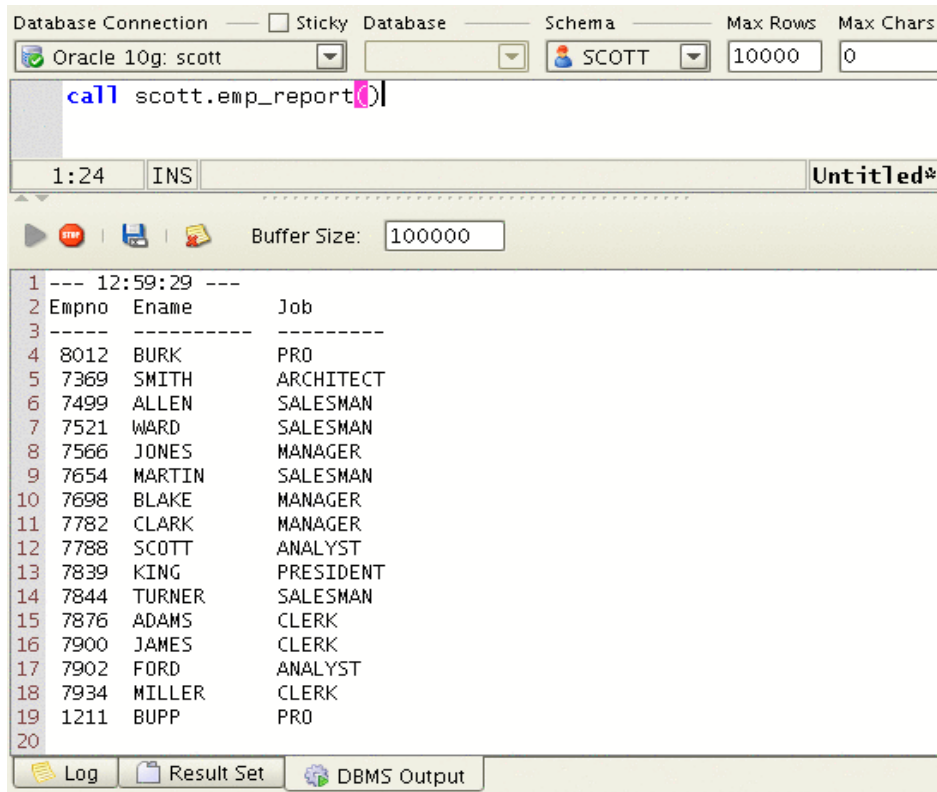


Figure: DBMS Output tab

Query Builder

Introduction

The Query Builder provides an easy way to develop database queries. The Query Builder provides a point and click interface and does not require in-depth knowledge about the SQL syntax.

The Query Builder is part of the SQL Commander, alongside the SQL Editor. To open the Query Builder, make sure the SQL Commander tab is selected and then choose either the SQL->Show Query Builder menu choice or click the vertical Query Builder button to the right in the SQL Commander. When you are ready to test a query built with the Query Builder, you just load it to the SQL Editor for execution.

Note: This document talks only about Tables even though the Query Builder supports both table and view objects.

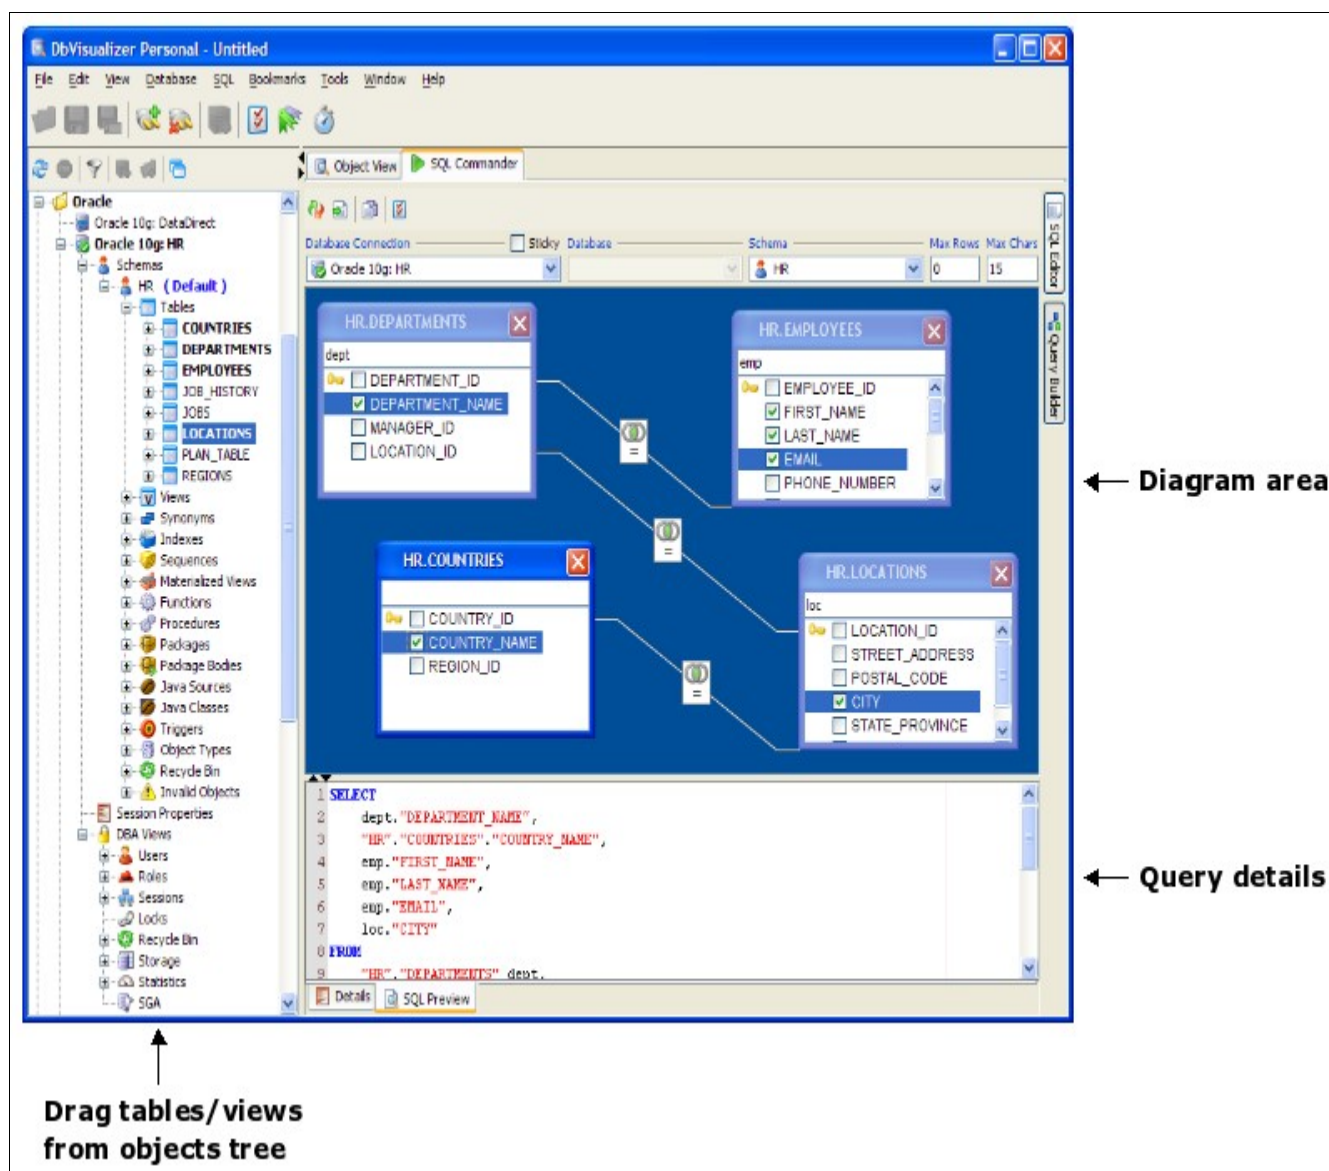


Figure: The query builder

Current Limitations

These are the current limitations in the Query Builder:

- Unions and sub queries are not supported.
- Not all join types are supported when joins are expressed as WHERE clause conditions. The Inner join type is supported for all databases, but the Left and Right types are only supported for databases with proprietary syntax to express these types, e.g., Oracle, SQL Server and Sybase. The Full type is not supported for any database. If a join type is not supported, the setting in the Join Properties dialog is silently ignored.
- When importing an SQL query from the SQL Editor, unsupported keywords and statement clauses are ignored. A dialog tells you which parts of the query are being ignored when unsupported parts are found in the imported statement.

Creating a Query

To create a query, open the query builder using the SQL->Show Query Builder menu choice or click the Query Builder button in the SQL Commander as described earlier. Make sure that the controls in the top section of the Query Builder are set correctly, as described in [Database Connection, Catalog and Schema](#).

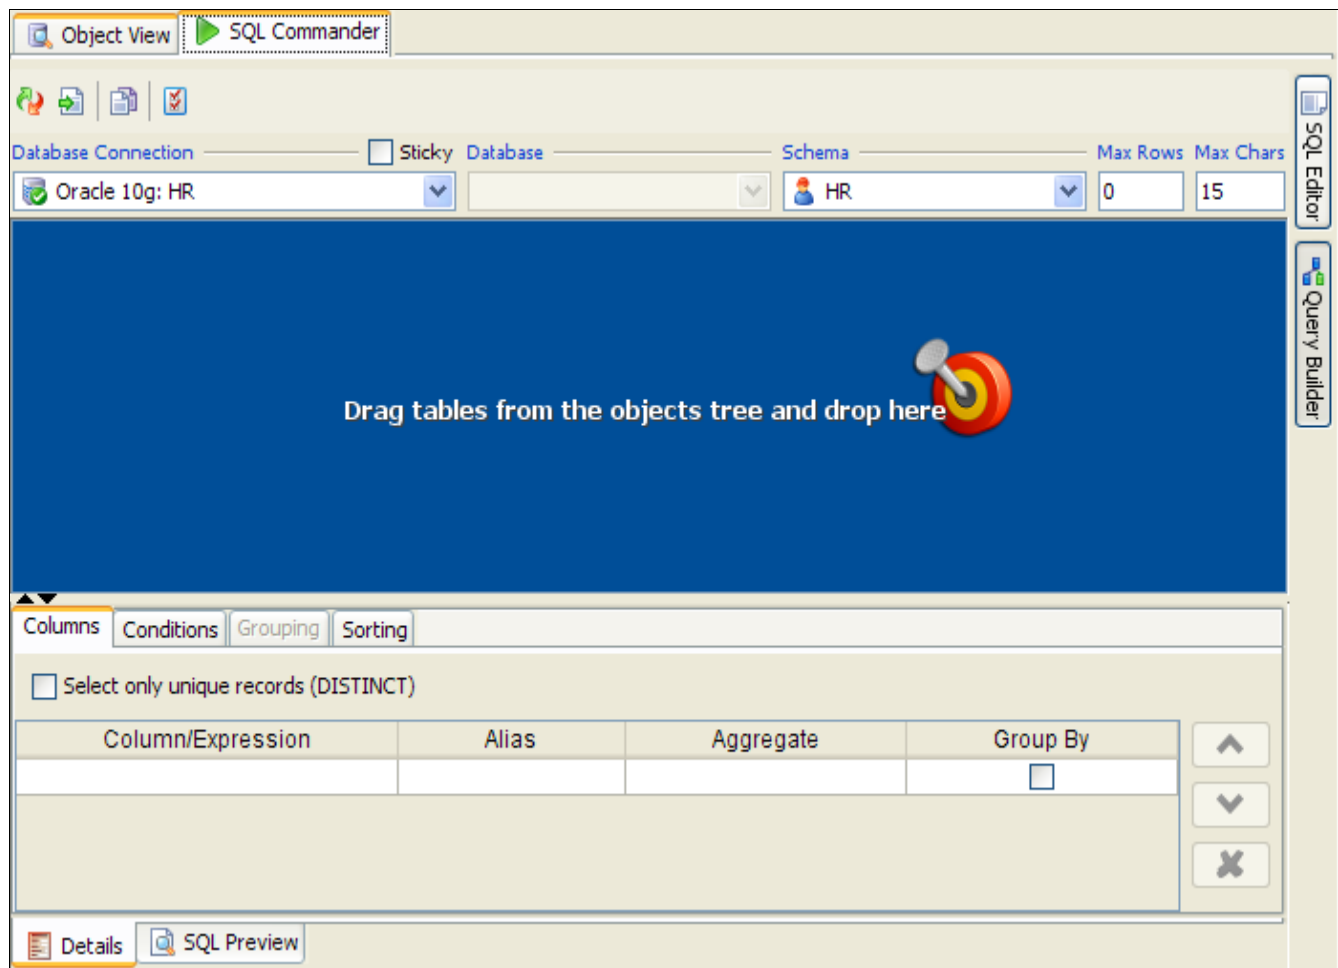


Figure: The initial appearance of the query builder

The easiest way to jump between the Query Builder and the SQL Editor is by clicking the vertical control buttons to the right in the SQL Commander. Clicking these buttons changes the display, but does not copy the query from one display to the other. To copy the current query from the Query Builder to the SQL Editor, use the toolbar buttons at the top of the Query Builder:




Figure: Query builder toolbar

1. The first button (from left) replaces the content of the SQL Editor with the query SQL
2. The second button adds the query last in the SQL Editor
3. The third button copies the query to the system clipboard
4. The fourth button opens the [editor properties](#)

The two first buttons automatically change the display to the SQL Editor.

You can also load a query from the SQL Editor into the Query Builder, as described in detail [below](#).

Adding Tables

To add tables, make sure the database objects tree and the table and/or view objects are visible. Then select and drag nodes from the tree into the diagram area.

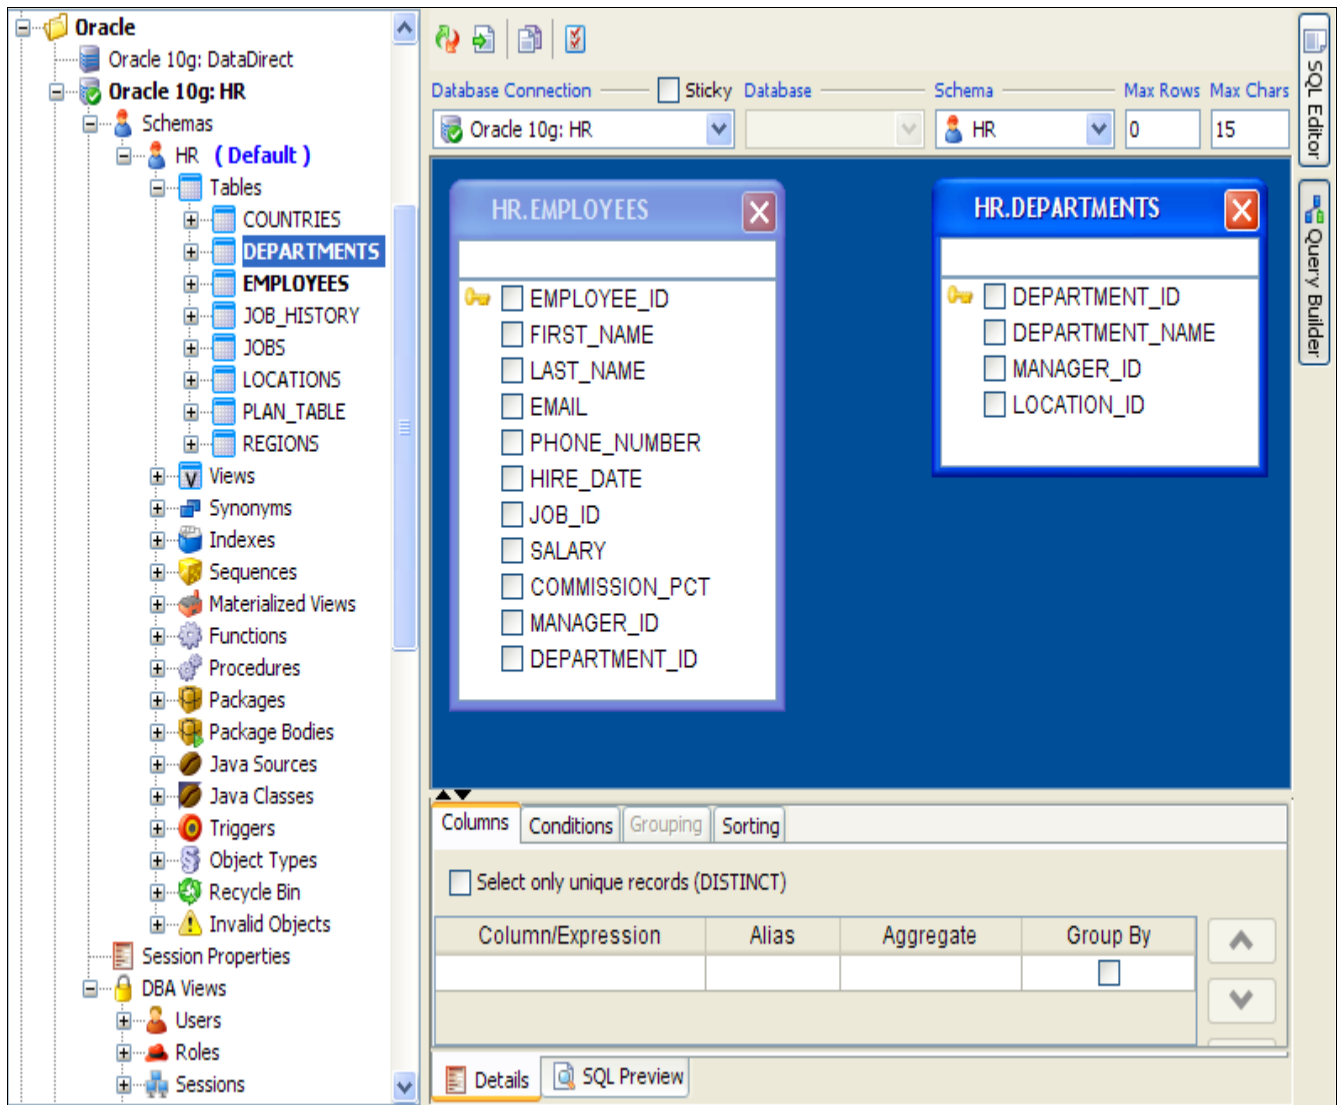


Figure: Adding tables to the query builder

To add a table, drag it from the object tree to the diagram area of the Query Builder. When the table is dropped in the diagram area, it is shown as a window with the table name as the window title.

Below the title is a text field where an optional table alias can be entered. If a table alias is specified, it is used in the Query Builder and the generated SQL statement to refer this table.

Under the table alias field is a list of all table columns. A check box in front of each name is used to select whether the column should be included in the query result set. Columns selected for the query result set also appear in the Columns and Sorting details tabs.

Joining Tables

Manually Joining Tables

To join two tables, select the column in the source table window with the mouse, drag it to the target table

column, and drop it.

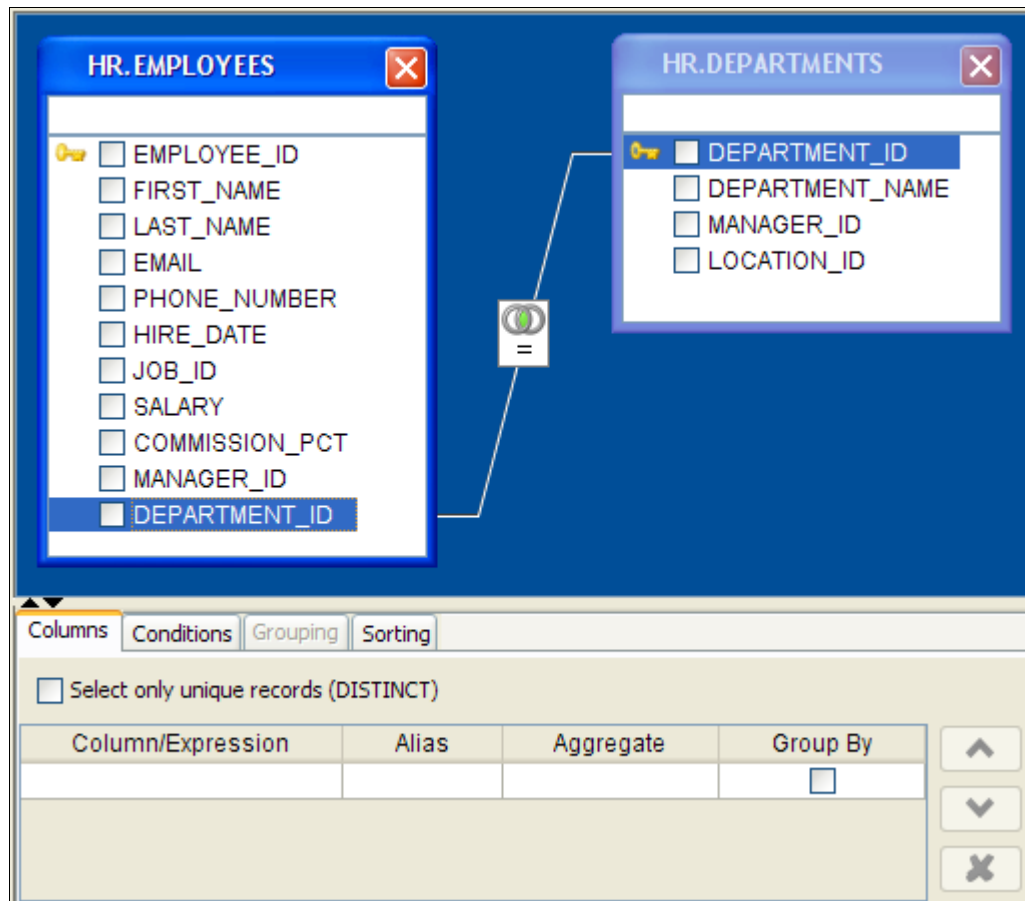


Figure: Joining two tables

The two columns now represents a join condition, represented by a link between the columns. If more than one join condition is needed, link additional columns in the two tables by dragging and dropping the columns in the same way as for the first join condition. The default join type is an Inner join and the default condition is "equal to" (=), represented as an icon with overlapping circles with the shared area shaded and an equal sign below them.

Joining Tables Automatically

Some database schemas declare how tables are related using primary and foreign keys. Other schemas use column names to indicate these relationships. For instance, in the figure above, the EMPLOYEES table has a column named DEPARTMENT_ID, which refers to the column with the same name in the DEPARTMENTS table. The Query Builder can be configured to use both kinds of rules to automatically join the tables you add to the query builder.

The auto-join feature is disabled by default. You can enable it in the tool properties for the database type (Tools->Tool Properties, under the Database tab) or for a specific connection (the Properties tab at the bottom of the Object View window when the connection is selected in the object tree).

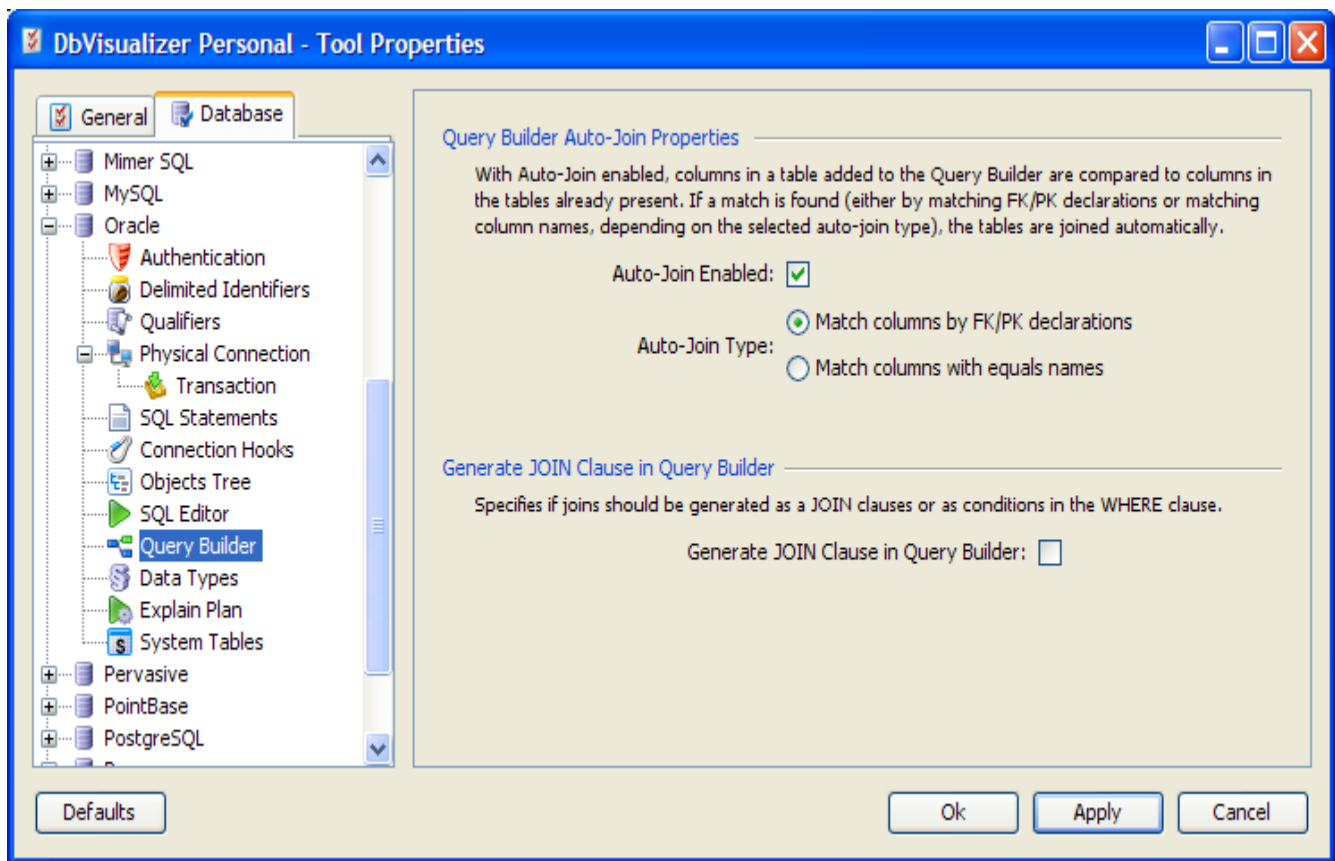


Figure: Query Builder Properties

The Query Builder node lets you enable the auto-join feature and select whether to use key declarations (FK/PK) or column names to find out how the tables are related.

When you add a new table with auto-join enabled, the Query Builder automatically joins it to the tables already in the builder if table columns match the selected matching rule.

If columns in the table you add are related to other columns in the same table, the Query Builder creates two windows for the table and joins them based on the matching rule. In this case, a table alias is also added for one of the windows so that you can tell the two windows for the same table apart.

Join Properties

A Join Properties dialog can be opened by double-clicking the icon or selecting Join Properties from the right click menu while the mouse pointer is over the join icon. The Join Properties dialog shows the source and target table columns and the conditional operator.

You can change the join type and the conditional operator in the Join Properties dialog. The join type defines how the records from the tables should be combined:

- Inner
This is the most common join type as it finds the results in the intersection between the tables.
- Left
This join type limits the results to those in the left table leaving 0 matching records in the right table as NULL.
- Right
This is the same as left join but reversed

- Full
A full join combines the results of both left and right joins.

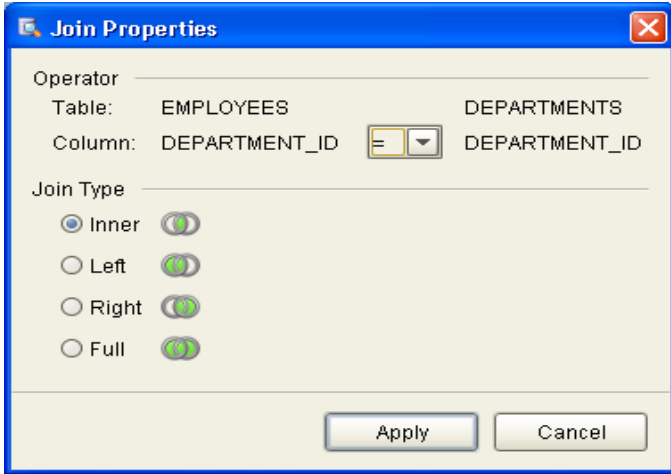


Figure: Join Properties dialog

Note: If you have multiple join conditions (linked columns) between two tables, you can specify different conditional operators for each join condition, but the join type is shared between all join conditions; if you change it for one join condition, it is changed for all the other join conditions linking the two tables. This is not a restriction in the Query Builder but rather how SQL is defined.

Here is the sample SQL generated from the previous join definition:

```

SELECT *
FROM HR.EMPLOYEES
INNER JOIN HR.DEPARTMENTS
ON (HR.EMPLOYEES.DEPARTMENT_ID = HR.DEPARTMENTS.DEPARTMENT_ID)

```

Remove Tables and Joins

A table window is removed by clicking the close icon in the window header. A join is removed by selecting Remove Join in the right click menu while the mouse pointer is over the join icon.




Figure: Diagram right click menu

All tables and joins may be removed via Remove All Joins and Remove All Tables.

Query Details

The Details tabs below the diagram area are used to define the various parts of the query. The tabs briefly represents the following parts of the final SQL:


```

SELECT <Columns>
FROM <tables>
WHERE <Conditions>
GROUP BY <Columns>
HAVING <Grouping>
ORDER BY <Sorting>

```

(The <tables> clause is defined in the diagram).

Columns

Use the Columns tab to specify characteristics of the columns that are included in the query. The list is initially empty until a column is checked in a table window or if manually adding a column expression. Columns will appear in the list in the same order as they are checked but may be manually moved at any time with the up and down buttons. To include all columns from a table, right click in the column list in the table window and choose Select All.

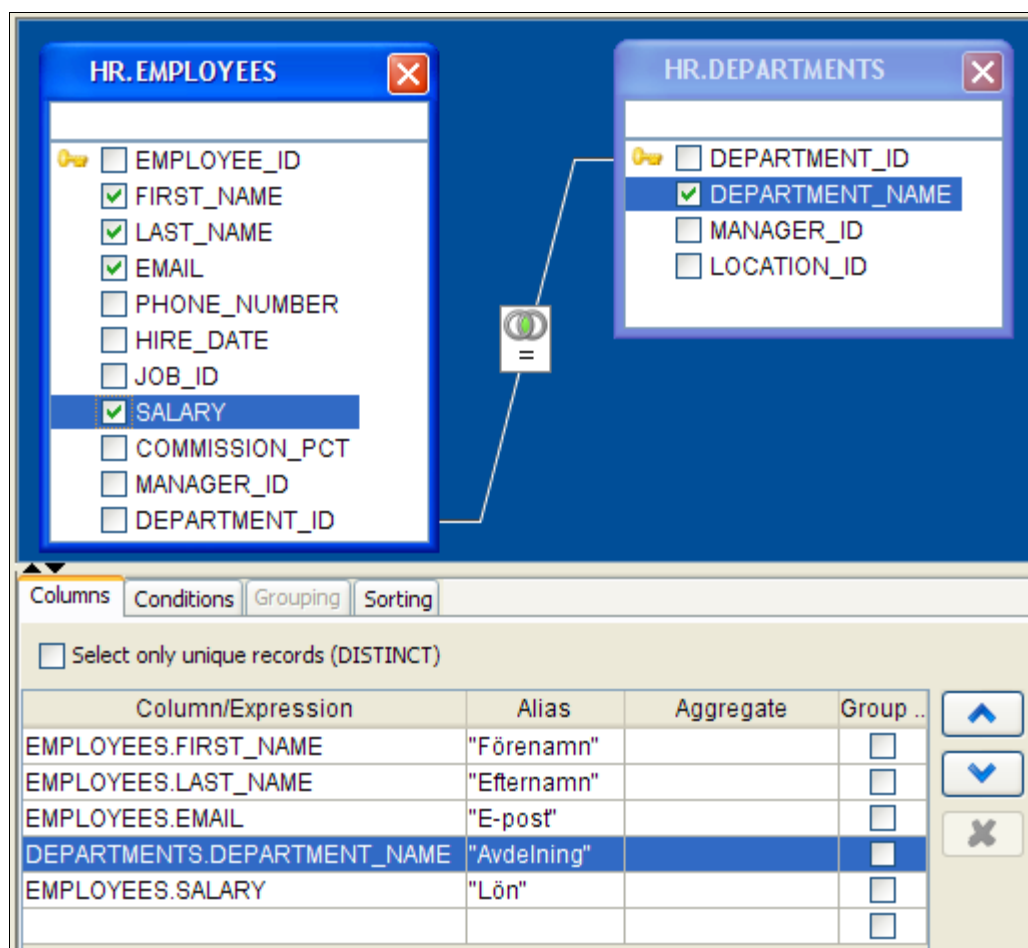


Figure: The Columns tab

The previous screenshot shows a total of 5 checked columns in the two tables. These are presented in the columns list by their full column identifier, qualified by either the table name or the table alias. To remove a column from the list, uncheck the corresponding column in the table window.

The alias field is used to specify an optional alias identifier for the column. The alias is used as the identifier for

the column in the final query and also appears as the column name in the result set produced by the query. Check the documentation for the actual database whether the alias must be quoted since the Query Builder does not do this for you.

The Aggregate and Group by fields are used in combination:

- The Aggregate field lists the available aggregation functions (AVG, COUNT, MAX, MIN, SUM) that may be used for columns
- The Group By field specifies whether the column should be included in the group for which aggregate columns are summarized

The Group By field is disabled unless an aggregate function is selected for at least one column, and once you select an aggregate function for one column, you must set Group By for at least one of the other columns to form a valid query. If you remove the aggregate function for all columns, Group By is automatically reset for all columns. Group By and aggregate are also mutually exclusive options for one column, so when you select one of them, the field for the other is disabled for that column.

A custom expression may be added by entering data in the empty row last in the list, e.g., "col1 + col2" or "TO_CHAR(ts_col, 'DD-MON-YYYY HH24:MI:SSxFF')". Once entered, press enter to insert a new empty row. You can remove a custom expression by selecting it and clicking the Remove button.

Conditions

The Conditions tab is used to manage the WHERE clause for the query. A WHERE clause may consist of several conditions connected by AND or OR. The evaluation order for each condition is defined by indentation in the condition list. Each level in the list will be enclosed by brackets in the final SQL.

Here is an example from the Conditions tab.

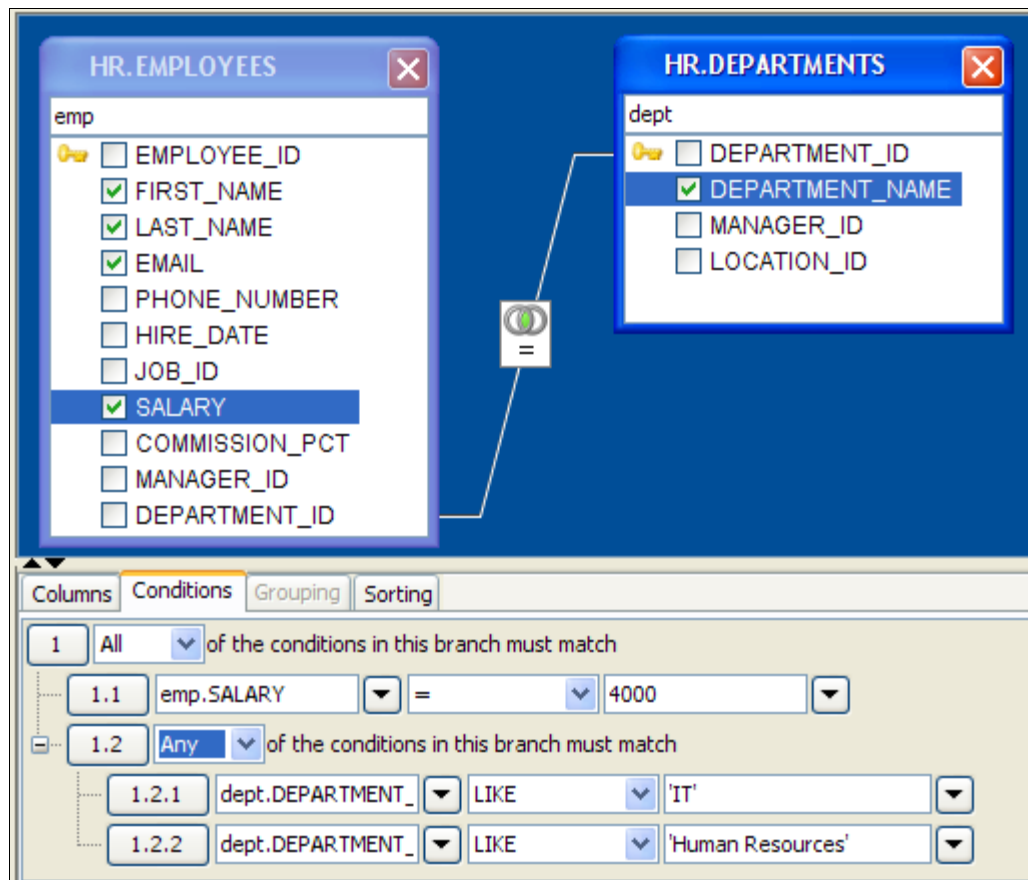


Figure: Condition settings

To create a new WHERE condition, press the indexed button in the list. In the menu that is displayed you may choose to create a new condition on the same level, a compound condition or delete the current condition. For compound conditions you may choose whether All (AND), Any (OR), None (NOT OR) or Not All (NOT AND) conditions must be met for its sub conditions. The SQL for the Conditions tab in the figure is:

```
WHERE emp.SALARY > 4000
AND (
    dept.DEPARTMENT_NAME = 'Human Resources'
    OR dept.DEPARTMENT_NAME = 'IT'
)
```

Next to the input field for each condition, there is a drop down button. When pressed it shows all columns that are available in the tables currently being in the Query Builder. You can pick columns from the list instead of typing these manually.

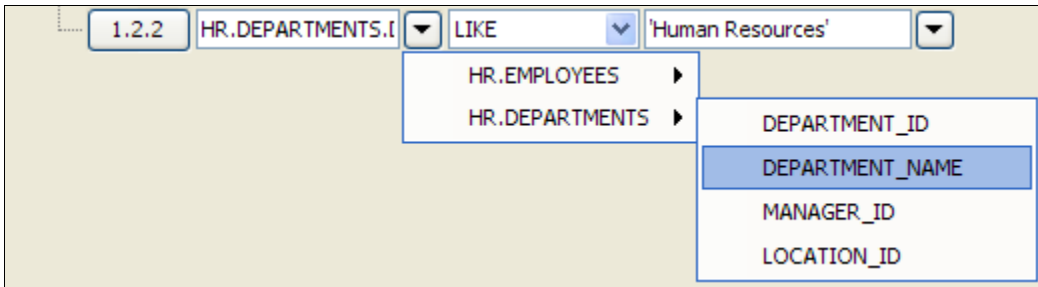


Figure: List of columns in the Conditions tab

Grouping

The Grouping tab is used to define the conditions for the HAVING clause that may follow a GROUP BY clause in an SQL query. This tab is only enabled when at least one of the columns in the Columns tab is marked as a Group By column.

The HAVING clause is similar to the WHERE clause, except that the HAVING clause limits what rows are included in the groups defined by the GROUP BY clause, after the WHERE clause has been used to limit the total number of rows to process.

You work with conditions in this tab in the same way as described in the [Conditions](#) section, with one exception regarding to the drop-down button for the fields in a condition. In the Grouping tab, the drop-down shows all columns listed in the Columns tab, with an aggregate function expression for columns that have an aggregate function defined. This is because (according to the SQL specification) the conditions in a HAVING clause must only refer to columns that are being returned by the query.

Sorting

The sorting tab is used to specify how the final result set will be sorted. All columns for the tables in the graph, plus any custom expressions created for the selection list in the Columns tab, are listed in the Sorting tab.

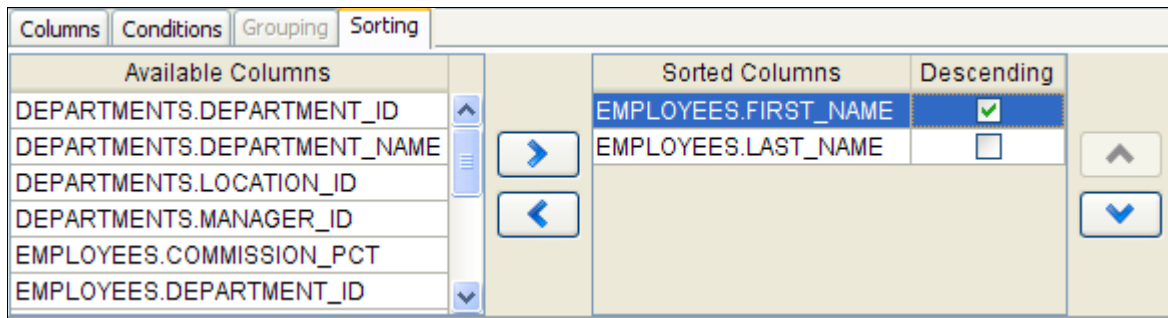


Figure: The Sorting tab

All columns listed in the Columns tab are initially listed in the Available Columns table. Select the ones you want to use in the sorting criteria and click the Move Left button to move them to the Sorted Columns table.

In the Sorted Columns table, you can change the default sort order (ascending) by clicking the check box in the Descending Order column. You can remove columns from the sorting criteria by selecting them in the Sorted Columns table and clicking the Move Right button.

SQL Preview

The SQL Preview tab at the bottom of the query builder is used to show a preview of the final SQL. This is a read-only view and cannot be modified.

Testing the Query

To test the query, simply press the appropriate toolbar buttons in the Query Builder to copy the SQL to the SQL Editor. Then execute the SQL as usual in the SQL Editor.

The screenshot shows the Oracle SQL Developer interface. At the top, there's a toolbar with various icons. Below it, the 'Database Connection' section shows 'Oracle 10g: HR' selected. The 'Schema' dropdown is set to 'HR'. The 'Max Rows' and 'Max Chars' fields are both set to 0. The main editor contains the following SQL query:

```

1 SELECT
2   "HR".EMPLOYEES."FIRST_NAME"      AS "Förenamn",
3   "HR".EMPLOYEES."LAST_NAME"      AS "Efternamn",
4   "HR".EMPLOYEES."EMAIL"          AS "E-post",
5   "HR".DEPARTMENTS."DEPARTMENT_NAME" AS "Avdelning",
6   "HR".EMPLOYEES."SALARY"         AS "Lön"
7 FROM
8   "HR"."EMPLOYEES"
9 INNER JOIN
10  "HR"."DEPARTMENTS"

```

Below the editor, the status bar shows '9:12', 'INS', 'Auto Commit: ON', and 'Untitled*'. A preview window below the editor shows the results of the query:


	Förenamn	Efternamn	E-post	Avdelning	Lön
1	Winston	Taylor	WTAYLOR	Shipping	3200
2	William	Gietz	WGIETZ	Accounting	8300
3	William	Smith	WSMITH	Sales	7400
4	Vance	Jones	VJONES	Shipping	2800
5	Valli	Pataballa	VPATABAL	IT	4800
6	Trenna	Rajs	TRAJS	Shipping	3500
7	TJ	Olson	TJOLSON	Shipping	2100
8	Timothy	Gates	TGATES	Shipping	2900

At the bottom right of the preview window, the execution statistics are shown: '0.030/0.030 sec', '106/5', and '1-8'.

Figure: Testing the SQL

To further refine the SQL press the Query Builder button and apply the necessary changes.

Loading a Query from the SQL Editor

If you have an existing SQL query that you want to modify using the Query Builder, you can load it from the SQL Editor into the Query Builder by clicking the rightmost button in the SQL Editor toolbar: 

It's important to be aware that the Query Builder does not support all features of the SQL SELECT statement, such as comments, UNION, and database-specific keywords. If you load a query into the Query Builder that contains unsupported constructs or keywords, they are ignored and a dialog pops up with a warning about this fact. You can then use the SQL Preview tab in the Query Builder to compare the SQL as it is represented in the Query Builder with the original SQL that you loaded to understand what was ignored.

Properties controlling Query Builder

There are a few properties that control how the Query Builder works and the SQL it generates. You can set these properties for the database type (Tools->Tool Properties, under the Database tab) or for a specific connection (the Properties tab at the bottom of the Object View window when the connection is selected in the object tree). Check the following sections for details.

Express joins as JOIN clause or WHERE condition

The Generate JOIN Clause in SQL Builder property is available in the [Database Type]->Query Builder node. Joins can be expressed either via the standardized SQL JOIN clause or a WHERE clause, using database-specific syntax for the Left and Right join types. The database-specific WHERE clause syntax is somewhat different between the supported databases and the Full outer join type is not supported. The default for this property is to use a JOIN clause.

A simple inner join expressed as a JOIN clause:

```
FROM HR.EMPLOYEES
INNER JOIN HR.DEPARTMENTS
ON (HR.EMPLOYEES.DEPARTMENT_ID = HR.DEPARTMENTS.DEPARTMENT_ID
```

Here is the same join expressed as a WHERE condition:

```
FROM HR.EMPLOYEES, HR.DEPARTMENTS
WHERE HR.EMPLOYEES.DEPARTMENT_ID = HR.DEPARTMENTS.DEPARTMENT_ID
```

The syntax for expressing Inner and Outer joins in WHERE conditions is different between databases. Oracle, for example, uses the "(+)" sequence to the left or right of the conditional operator to express left or right joins. SQL Server and Sybase use "*"=" or "=*" for the same purpose.

DbVisualizer automatically sets the correct join notation when generating joins as WHERE conditions for databases that support left and right joins using WHERE conditions. For databases that do not provide syntax for left and right joins, the join type is ignored and the WHERE condition that is generated produces an inner join result.

Table and Column Name qualifiers

Whether to qualify table names with the schema or database name and whether to qualify column names with table name are defined in the [Database Type]->Qualifiers node.

Delimited Identifiers

Identifiers that contain mixed case characters or include special characters need to be delimited. Define this in the [Database Type]->Delimited Identifiers node.

Drag style and Diagram Size

If you enable the editor controls from the Query Builder or SQL Editor toolbar, you can also set the style table windows in the Query Builder diagram should have when moving them, as well as the default size for newly added table windows.

Monitor and Charts

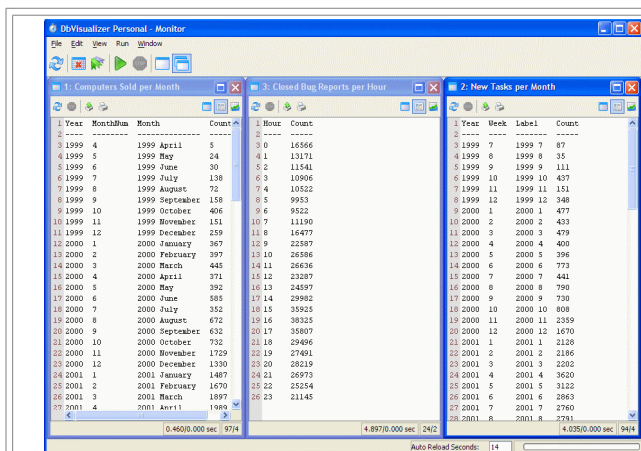
Introduction

With the monitor feature, you can track changes in data over time, viewing the results of one or many SQL statements either as grids or graphs. Typically, you configure the monitor to run the statements automatically at certain intervals.

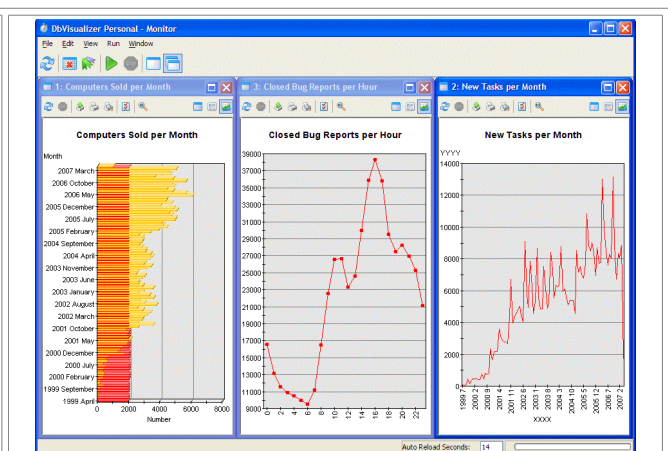
The monitoring feature combined with the charting capability in DbVisualizer Personal is really powerful, delivering real time charts of many result sets simultaneously. For example, you can use monitoring to spot trends in a production database, surveillance, statistics, database metrics, and so on. At Minq Software, we have a dedicated workstation that uses the monitoring feature to automatically present live chart information from our Internet servers and customer database.

Any SQL statement defined as an SQL Bookmark that produces a result set can be monitored, and when you monitor multiple statements, different statements may use different database connections concurrently.

Note: The chart customization covered in this document is also applicable to the charts for result sets in the SQL Commander (DbVisualizer Personal).



The Monitor window with four monitored SQL Bookmarks. The results can be viewed as windows or tabs. This example shows the grid data as returned from each SQL statement.



The same monitored SQL Bookmarks as in the left figure but here presented as charts. (DbVisualizer Personal)

Monitor an SQL statement

To monitor an SQL statement, you first define it as an SQL Bookmark in DbVisualizer. A bookmark is an SQL statement with associated information about the target database connection and (optionally) the catalog (the JDBC term which translates to a *database* for some databases, like Sybase, MySQL, SQL Server, etc). You use the Bookmark Editor to create and organize SQL bookmarks in a tree structured folder view. This is also where you mark which SQL Bookmarks should be available for monitoring. All information about the SQL Bookmarks is saved in an XML file between invocations of DbVisualizer.

You can read more about SQL Bookmarks in general in the [SQL Bookmarks](#) chapter; here we focus on the parts of interest for monitoring. Starting and stopping the monitoring takes place in the Monitor window, described in the next section. But first, let's see how you enable monitoring of an SQL Bookmark.

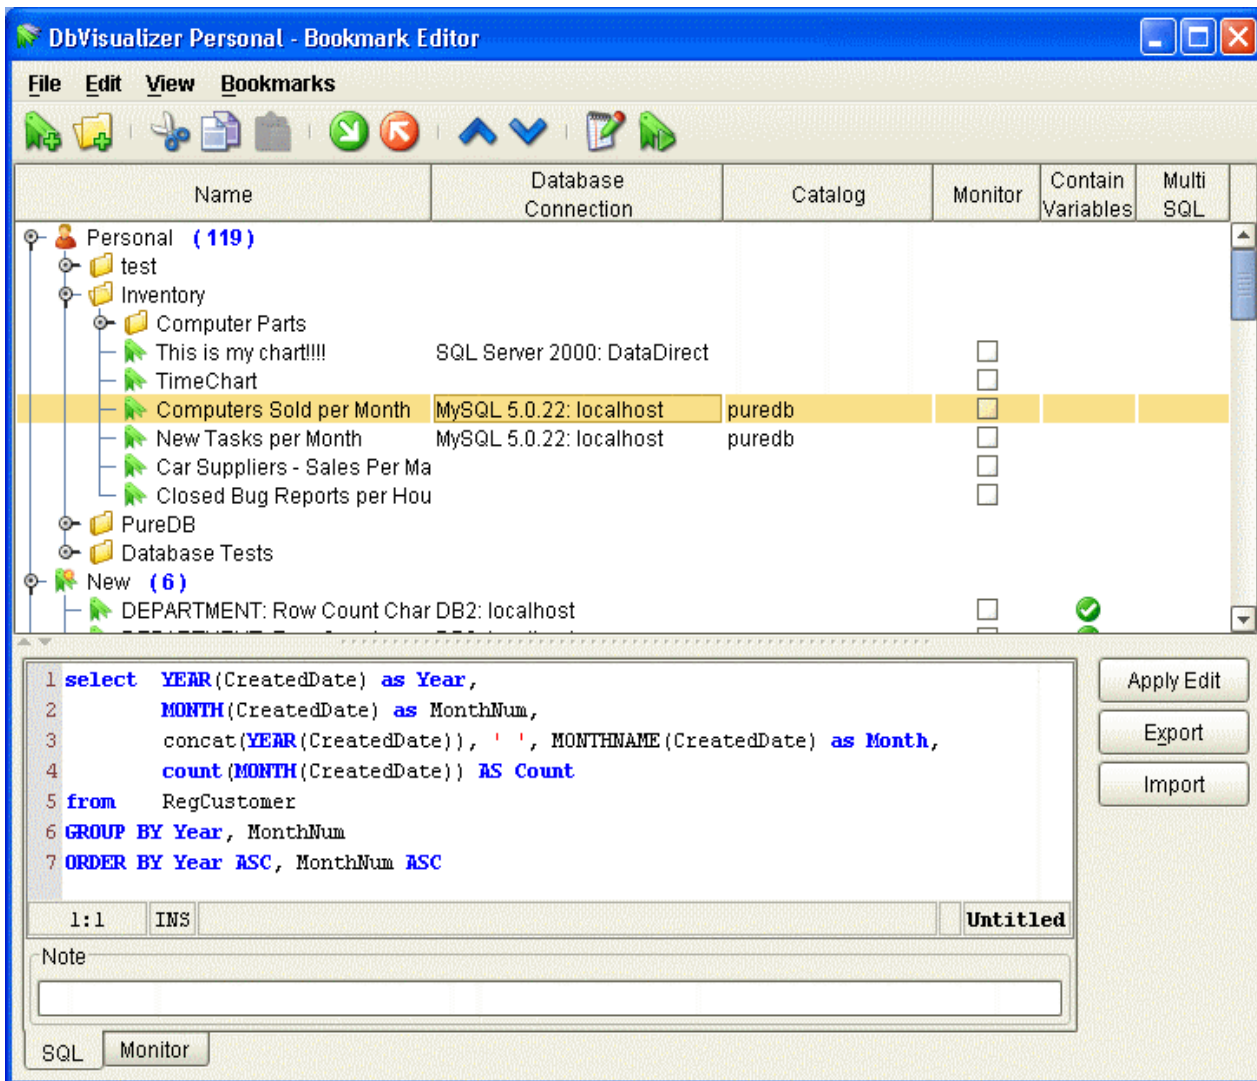


Figure: Bookmark Editor

The figure above shows the Computers Sold per Month bookmark and the SQL that is associated with it. The Monitor field in the tree is used to determine whether the SQL Bookmark is a monitor. Click on the checkmark to make the bookmark appear in the Monitor window, i.e., make it possible to monitor. Uncheck it to remove it from the Monitor window.

The following is an example of the result set the above SQL Bookmark produces:

Year	MonthNum	Month	Count
1999	2	1999 February	3
1999	3	1999 March	10
1999	4	1999 April	10
1999	5	1999 May	25
1999	6	1999 June	31
1999	7	1999 July	138
1999	8	1999 August	72
1999	9	1999 September	159
1999	10	1999 October	406
1999	11	1999 November	151
1999	12	1999 December	259
2000	1	2000 January	367
2000	2	2000 February	397
2000	3	2000 March	445

3.585 sec/0.000 sec 90 / 4 1-14

Figure: Monitor showing the result in Grid format

The interesting columns in the result are the Month and Count. The Year and MonthNum are there just to get the correct ascending order of the result.

You can read more about how to manually create and edit bookmarks in the [SQL Bookmarks](#) chapter. The following sections describe how you can get help creating the bookmarks for a couple of cases that are commonly used for monitoring.

Monitor table row count

It is very common to want to keep track of how the number of rows in a table varies over time. The right-click menu in the Data tab grid therefore has a Create Row Count Monitor operation that creates a monitoring enabled SQL Bookmark for you automatically.

It creates a monitor with SQL for returning a single row with the timestamp for when the monitor was executed and the total number of rows in the table at that time. Every time the monitor is executed, a new row is added to the grid, up to a specified maximum number of rows. When the maximum row limit is reached, the oldest row is removed when a new row is added. Example:

Computers: Row Count	
PollTime	RowCount
2003-01-23 12:19:10	43123
2003-01-23 12:11:40	43139
2003-01-23 12:21:10	43143
2003-01-23 12:22:40	43184
...	...

Figure: Example of the result from a Table Row Count monitor

The SQL for this monitor uses two variables, DbVis-Date and DbVis-Time. These variables are substituted with the current date and time, formatted according to the corresponding Tool Properties settings. The reason for using these variables instead of using SQL functions to retrieve the values is simply that it is almost impossible to get the values in a database independent way. Another reason is that we want to see the client machine time rather than the database server time. You can, of course, modify the SQL any way you see fit, as long as the PollTime and RowCount labels are not changed.

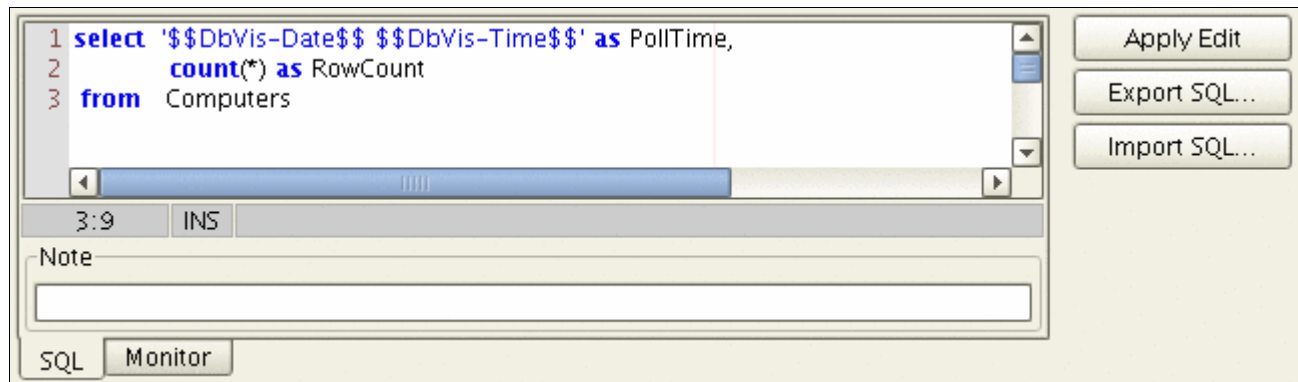


Figure: Sample of the SQL for the Table Row Count monitor

DbVisualizer keeps the result for previous executions of the SQL Bookmark, up to the specified maximum number of rows, so that you can see how the result changes over time. You define the maximum number of rows with the Allowed Row Count property in the Bookmark Editor, under the Monitor tab. This property is automatically set to 100 when you use Create Row Count Monitor to create the monitor.

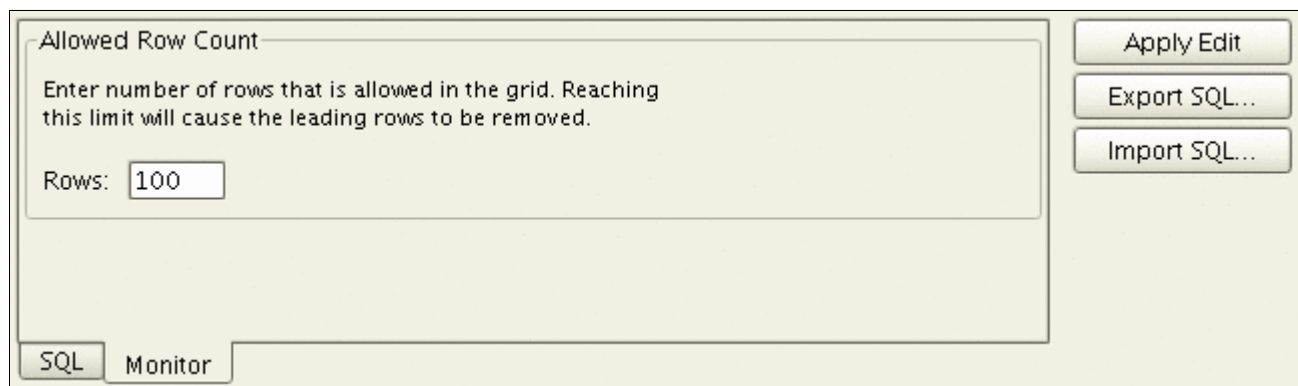


Figure: Allowed Row Count property pane

You can change the value in the Bookmark Editor, to limit or extend the number of rows that DbVisualizer should keep. Setting it to 0 or a negative number tells the DbVisualizer to always clear the grid between executions of monitors. Click the Apply Edit button to save the new value.

Monitor table row count difference

In addition to tracking the number of rows in a table over time, you may want to see by how many rows the value changes. You can create a monitor for this purpose with the Create Row Count Diff Monitor operation, available in the right-click menu for the grid in the Data tab.

In addition to the Row Count Monitor, the Row Count Diff Monitor reports the difference between the number of rows in the last two executions:

Computers: Row Count Change		
PollTime	RowCount	RowCountChange
2003-01-23 12:19:10	43123	0
2003-01-23 12:11:40	43139	16
2003-01-23 12:21:10	43143	4
2003-01-23 12:22:40	43184	41
...

Figure: Example of the result from a Table Row Count Difference monitor

The SQL for this monitor adds a third column, named RowCountChange. It utilizes the fact that DbVisualizer automatically creates variables for the columns in a monitor result set, holding the values from the latest execution. The RowCountChange column is set to the value returned by the count(*) aggregate function for the current execution minus the value from the previous execution, held by the RowCount variable. All columns in a monitor result set can be used like this to reference values from the previous execution of the monitor.

```

1 select '$$DbVis-Date$$ $$DbVis-Time$$' as PollTime,
2        count(*) as RowCount,
3        count(*) - $$RowCount||count(*)$$ as RowCountChange
4 from   Computers

```

Figure: Sample of the SQL for the Table Row Count Difference monitor

Monitor window

The Monitor window, launched via the Tools->Monitor menu option, is where you browse the active monitors. The monitors can be organized either as tabs or internal windows. The monitor results can be viewed only as grids in DbVisualizer Free, while DbVisualizer Personal adds the capability to view them as charts. The following figure is a screenshot of the Monitor window:

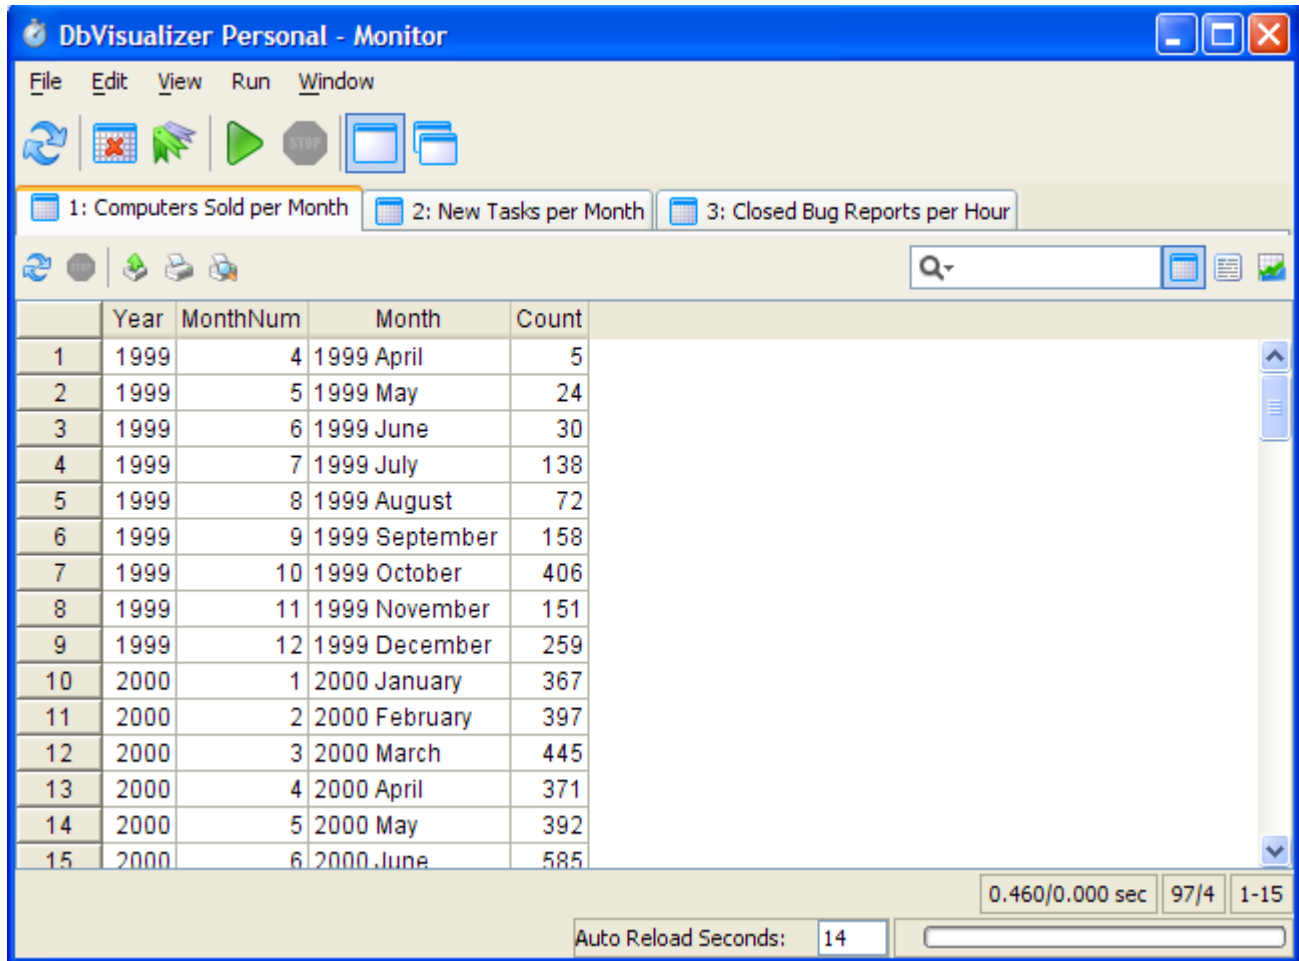


Figure: The Monitor window with all monitors organized as tabs

The Monitor window has a menu and a toolbar at the top, an area for the monitors (i.e., monitoring enabled SQL Bookmarks) in the middle, and an Auto Reload Seconds field and status bar at the bottom. The screenshot is from DbVisualizer Personal, with View buttons in a toolbar for the selected monitor; these buttons are not included in DbVisualizer Free.

The main toolbar buttons have the following functions:

Toolbar Button	Description
Reload	Reloads all grids (i.e., executes all monitors and updates the result set grids)
Clear Current	Clears the currently selected grid. To clear all grids at the same time, use the Edit->Clear All Monitors m
Show Bookmarks	Launches the Bookmark Editor
Start Monitors	Starts auto-update of all monitors, repeatedly executing all monitors at the intervals specified by the Auto

Stop Monitors	Stops the auto-update
Show as Tabs	Shows the monitors as tabs
Show as Windows	Shows the monitors as internal windows

The Auto Reload Seconds field at the bottom of the main window is used to control how often to execute the monitors when auto update is running. Use the field to specify how many seconds to wait between auto-reloads. The status bar shows the time until the next auto-update.

Note: The specified number of seconds may be increased automatically by DbVisualizer if the total execution time for all monitors is longer than the specified value.

The Window menu contains choices to control the appearance in the Monitor:


 Show <u>G</u> rids	Ctrl+4
 Show <u>T</u> exts	Ctrl+5
 Show <u>C</u> harts	Ctrl+6
Select <u>N</u> ext Tab/Window	Alt+Right
Select <u>P</u> revious Tab/Window	Alt+Left
 <u>C</u> ascade Windows	Alt+C
 Tile Windows <u>V</u> ertically	Alt+V
 Tile Windows <u>H</u> orizontally	Alt+H

Figure: Window menu operations

Show Grids, Show Texts and Show Charts toggle the monitor display between the selected views. Cascade and Tile are used to automatically arrange the windows in the Windows view.

Charts

This section is only applicable to DbVisualizer Personal.

Charts in conjunction with the Monitor feature is really powerful, since monitored data is very often a good candidate to be charted. The charting capability in DbVisualizer Personal is also available in the SQL Commander; everything described here also applies to the grids for the result sets in the SQL Commander.

The basic setup of a chart is really easy. It is just a matter of selecting one or more columns that should appear as series in the chart. The basic requirement is that the monitor has been executed, so that there are columns to choose the series from. The appearance of the charts can be thoroughly customized using the advanced customization editor.

The chart view is controlled by the buttons in toolbar shown in the monitor area when the graph mode is selected:




Figure: Chart control buttons

The controls are from the left:

1. Show/Hide chart controls pane
2. Reset any zoom


The following sections explain the features and how to setup the chart.

Chart Controls

The chart controls are used to customize the Data that shall be displayed in the chart, optional axis labels, titles, etc. It is also used to control the Layout of the chart in terms of chart type, legend type, etc.

Data

Use the controls in the Data tab to customize which data shall appear in the chart.



The screenshot shows the 'Data' tab of a chart customizer. It features a table with columns 'Serie', 'Column', 'Label', and 'Type'. The 'Serie' column has checkboxes for 'Year', 'Week', 'Label', and 'Count'. The 'Week' and 'Count' series are selected. The 'Label' field for 'Week' is 'Vecka'. Below the table are fields for 'Chart Title' (Red Alers per month), 'X-Axis Title' (X), 'Y-Axis Title' (Y), and 'Rotation' settings (None).

Serie	Column	Label	Type
<input type="checkbox"/>	Year		Long
<input checked="" type="checkbox"/>	Week	Vecka	Long
<input type="checkbox"/>	Label		String
<input checked="" type="checkbox"/>	Count		Long

Figure: Data customizer

Select at least one Series from the list of columns. As soon as you select the a series, it is immediately added to the graph. The Label field can be used to specify an optional label for the series as it will appear in a legend. The column name is used if no label is specified.

The X-Axis Label box is used to specify the column in the result that should be used to render the labels of the X-axis. Chart Title specifies the main title of the chart. This SQL Bookmark name, as defined in the Bookmark Editor, is used as the default title. X-Axis Title and Y-Axis Title specifies the titles for the X and Y axis. You can use the Rotation settings if you want the X and Y axis text rotated.

Layout

The layout tab is used to configure the appearance of the chart, primarily the type of chart want to use. Note that all settings are per monitor. The following screen shots show some of the most commonly used chart types.

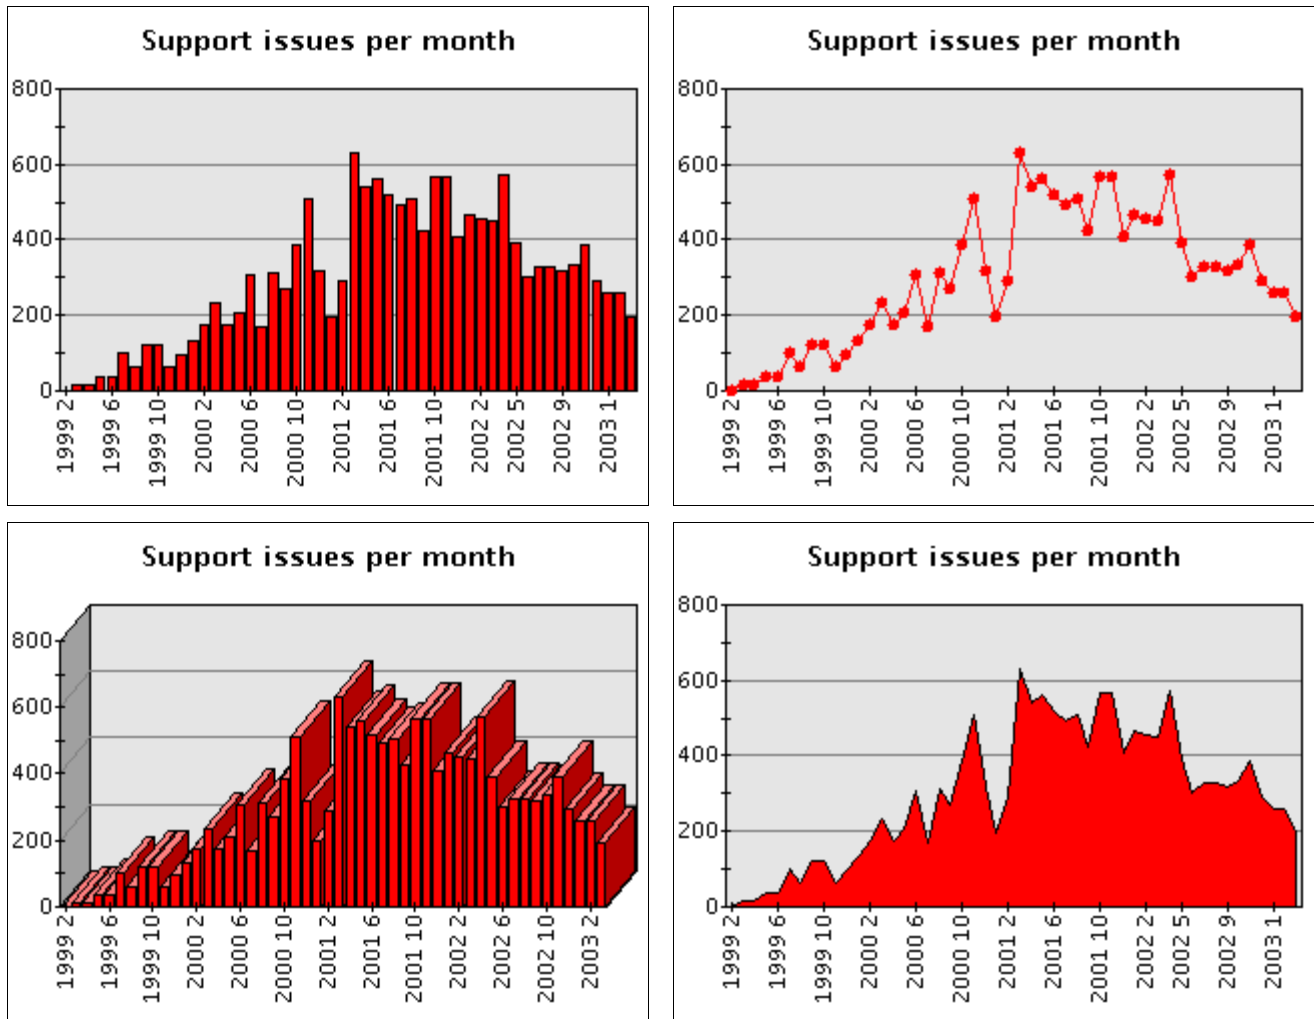


Figure: Chart type examples

The advanced layout editor can be used to customize every aspect of the layout. The basic layout settings, however, are the following:

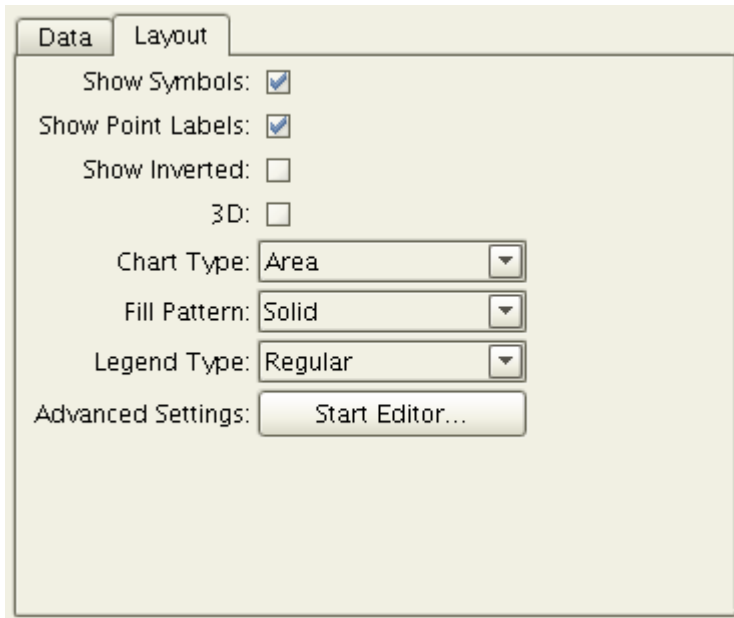


Figure: Layout customizer

Show symbols specifies whether each value in a line chart should be represented by a symbol. Show Inverted defines whether the X and Y axis will be switched. 3D specifies if a bar chart will be displayed in 3D. The Chart Type lists all the available chart types. Fill Pattern defines how a bar, area and pie chart shall be filled. Legend Type specifies whether a legend will be displayed or not.

You can use the Advanced Settings editor to customize all the bits and pieces of the chart. This document does not explain all the configurations that can be done using this editor since that would result in a 100 page book. Play around with the different settings and see how the graph changes.

Note: Settings that are made in the Advanced Editor are not saved between invocations of DbVisualizer.

Chart View

Zooming

Charts support zooming by selecting a rectangle in the chart area. Selecting another rectangle in that zoomed area will zoom the chart even further, and so on. To reset the zoom, click the Reset Zoom button or press the "r" keyboard button while the mouse pointer is in the chart area.

Rotating

All 3D chart types support rotating and changing the depth of the chart. Use the following to change the appearance:

- Shift+Left Mouse button
Changes the depth of the chart
- Ctrl+Left Mouse button
Changes the rotation of the chart

Examples:


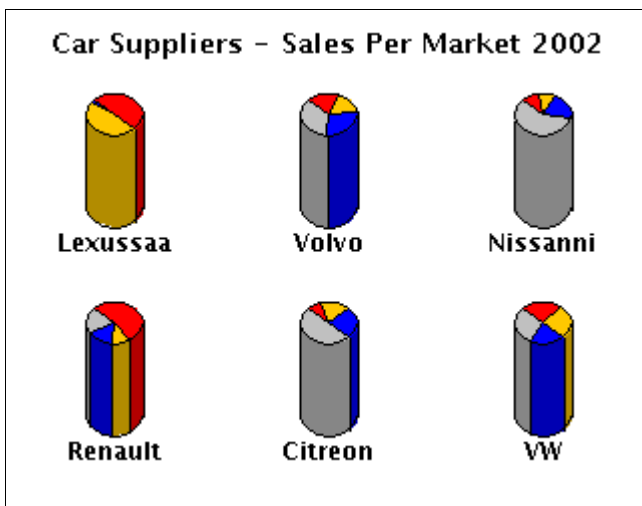

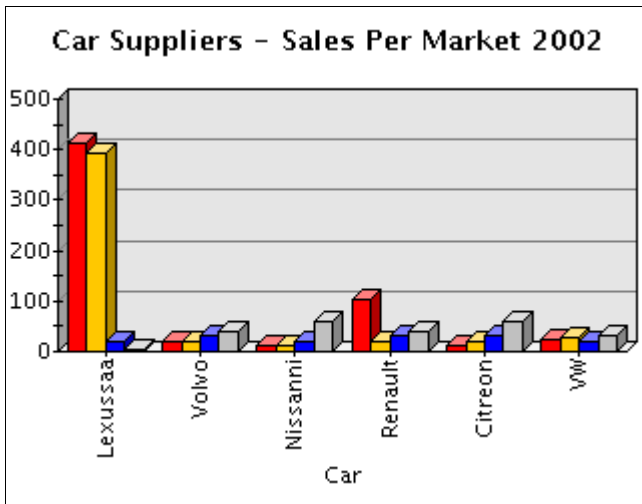


Figure: Example of 3D charts

The above screen shots are just a few examples of the 3D chart types and how depth and rotation settings are used to change the appearance.

Export

The export operation is context sensitive and works on the currently selected chart, graph or grid. The controls in the export dialog also adapt to the currently selected object. If a chart is the current object the following export dialog will appear:

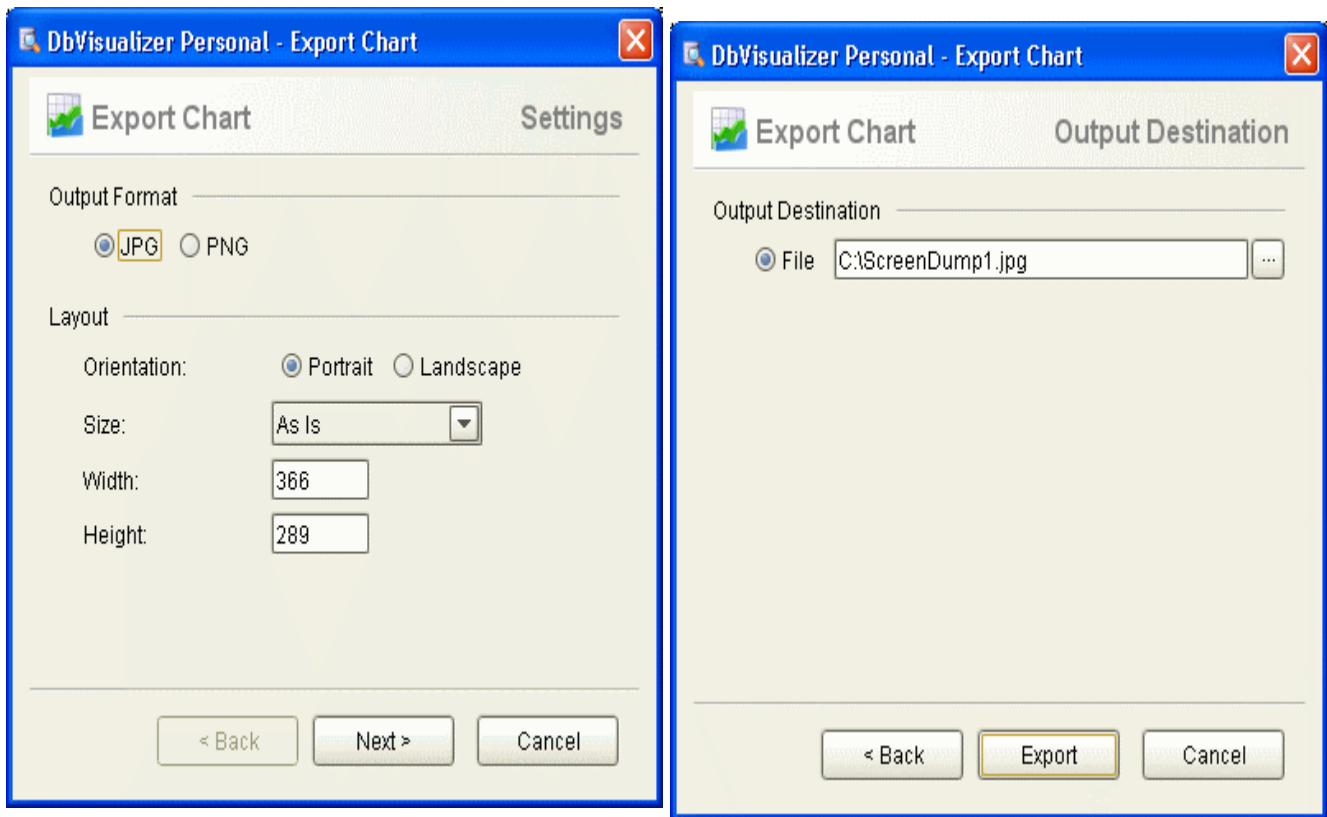


Figure: Export dialog for charts

The default size of the image that is about to be exported is the same as it appears on the screen. To change the size then either select a pre-defined paper size in the Size list or enter a size in pixels.

Create and Alter Table

Introduction

The Create Table, Create Index and Alter Table Assistants are used to create new tables and indexes and to alter existing tables. The assistants are quite simple to use since they examine various metadata in the database (depending on which assistant is used) and then let you point and click to define the table or index.

The assistants are launched from the Database main menu, from in the Database Objects tree right-click menu, or from the Actions menu button in the object view. The menu choices are enabled only if a table or index can be created for the selected node in the Database Objects tree.

Create Table

To create a table, select an appropriate node in the objects tree, typically a Tables node, and launch the Create Table assistant from one of the menus as described earlier.

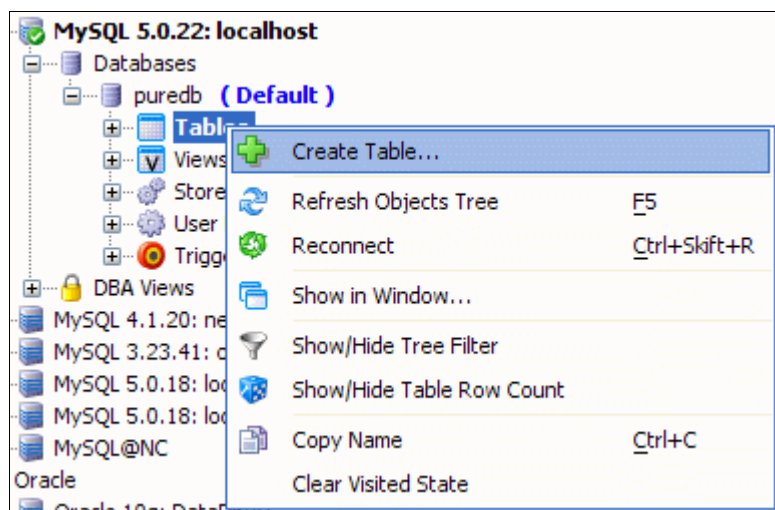


Figure: The right-click menu in the Database Objects tree

The Create Table assistant is organized in three areas from the top:

- General Table Info
Specifies the owning database connection, database and/or schema. These are picked up from the selection in the tree when the assistant is started. Table name is set to a default name that you can change to the real table name.
- Table Details
A number of tabs where you specify information about the columns and, optionally, various constraints. The Columns, Primary Key and Foreign Key tabs are available for all databases. The remaining tabs are database-specific and depends on the features supported by the database engine.
- SQL Preview
The SQL previewer instantly shows the SQL statement for creating the table.

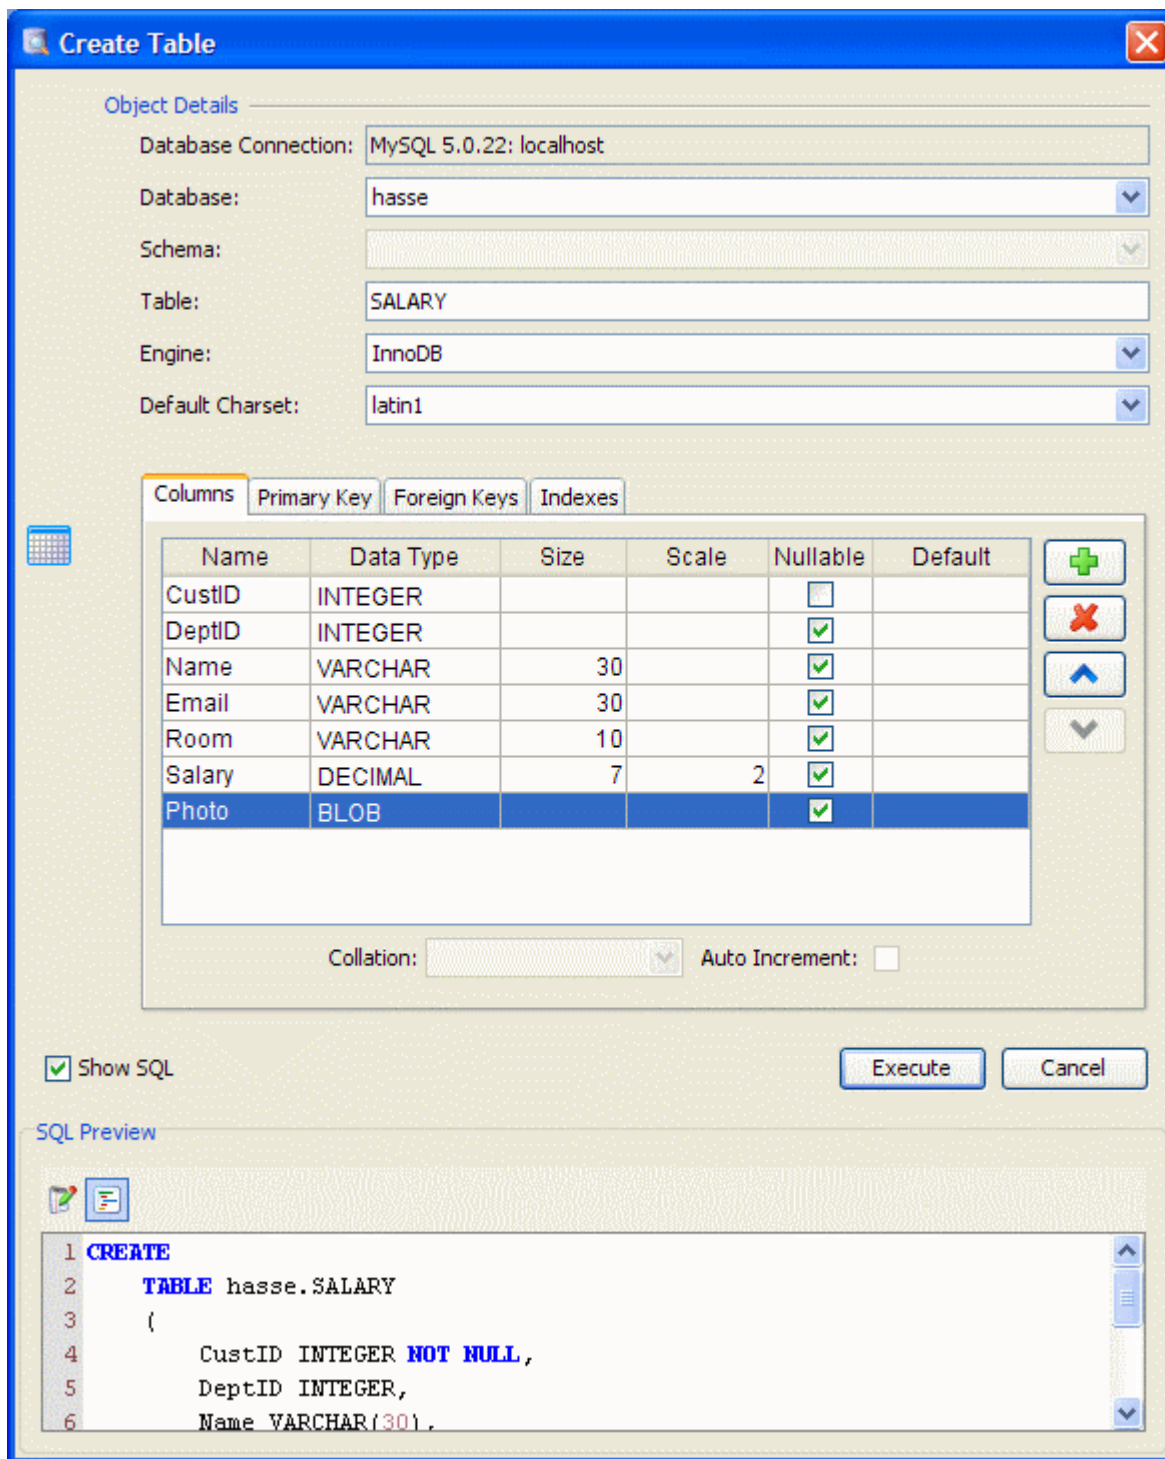


Figure: The Create Table assistant

Just enter as much information as is needed to describe the table and click Execute to create the table.

Columns tab

The Columns tab lists all table columns along with their attributes.

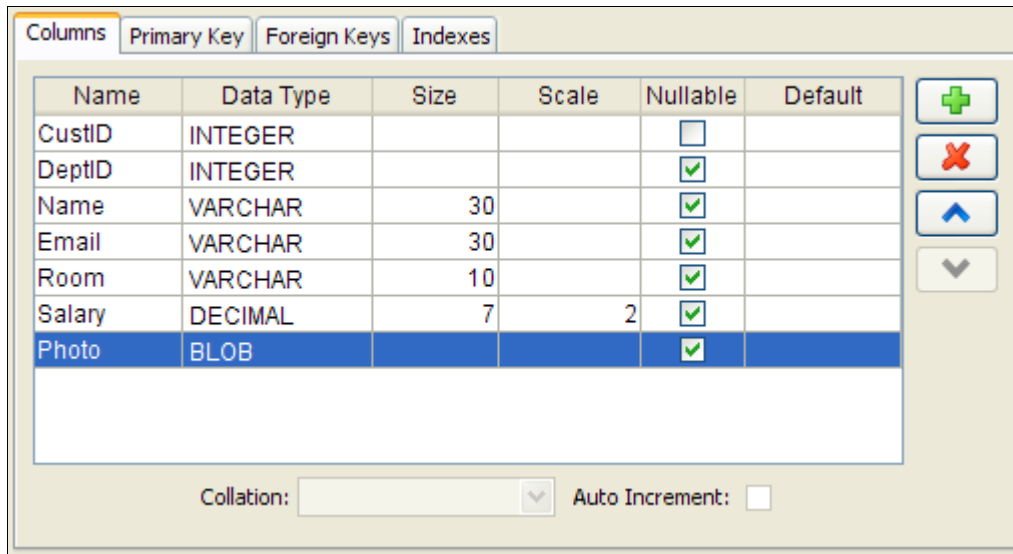


Figure: The Columns tab (for MySQL)

Add columns by clicking on the Add button, and remove the currently selected column by clicking on the Remove button. You can reorganize the columns using the Up and Down buttons.

Enter the name of the column in the first field and select a data type from a drop down list in the second field. The list contains the names of all data types the database supports.

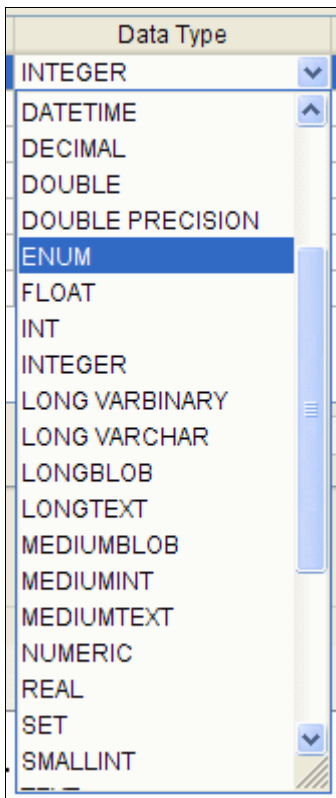


Figure: Data Type list (for MySQL)

For some data types, such as character types, you may also specify a size, i.e., the maximal length of the value. For others, like the decimal types, you can may specify both a size and a scale (the maximal number of decimals).

Room	VARCHAR				<input type="checkbox"/>
Salary	DECIMAL	7	2	<input checked="" type="checkbox"/>	
Phone	BLOB			<input type="checkbox"/>	

Figure: Size and scale for a DECIMAL data type

The above example will allow a total length (including the decimal places) of 7. Examples:

```

1.02
9871.1
8172.0
18291.22
    12.112 <- Error
1921211.11 <- Error

```

The last two fields let you specify if the table is nullable and a default value to use for rows inserted into the table without specifying a value for the column.


Below the column list, you may see one or two additional fields, depending on the features supported by the database you create the table for. The fields are enabled when you select a column that they apply to. The Collation field is enabled for character columns if the database supports the declaration of a collation for textual data. The Auto Increment field is enabled for numeric fields if the database supports automatically inserting the next available sequence number in a numeric column.

The Create Table assistant uses database metadata to try to enable only the fields that apply to the selected data type, but please note that it is not always possible. For instance, there is no metadata available to tell if a data

type requires, or allows, a size. If you don't enter a required attribute or enter an attribute that is unsupported for a data type, you will get an error message when you click Execute to create the table.

Primary Key tab

The Primary Key tab contains information about an optional primary key for the table. A primary key is a column, or a combination of columns, that uniquely identifies a row in a table.



The screenshot shows a dialog box with four tabs: 'Columns', 'Primary Key', 'Foreign Keys', and 'Indexes'. The 'Primary Key' tab is selected. At the top, there is a text field labeled 'Constraint Name:'. Below it is a table with two columns: 'Column' and 'Include'. The 'CustID' row is highlighted in blue and has a checked checkbox in the 'Include' column. Other rows include 'DeptID', 'Name', 'Email', 'Room', 'Salary', and 'Photo', all with unchecked checkboxes. To the right of the table are three buttons: an up arrow, a down arrow, and a close button (X).

Column	Include
CustID	<input checked="" type="checkbox"/>
DeptID	<input type="checkbox"/>
Name	<input type="checkbox"/>
Email	<input type="checkbox"/>
Room	<input type="checkbox"/>
Salary	<input type="checkbox"/>
Photo	<input type="checkbox"/>

Figure: Primary Key tab

You can, optionally, enter a constraint name for the primary key constraint in the Constraint Name field. Select columns to be part of the primary key by clicking the checkboxes in the Include field in the columns list.

You can change the order of the columns in the key by selecting a column and move it using the Up and Down buttons.

Foreign Keys tab

In the Foreign Keys tab, you can declare one or more foreign keys for the table. A foreign key is a column, or a combination of columns, that refer to the primary key of another table. Foreign keys are used by the database to enforce integrity, i.e., that there is a row in the referenced table with a primary key that matches the foreign key value when a new row is inserted or updated, and can optionally declare rules for what to do when a referenced primary key is removed or updated in the referenced table.

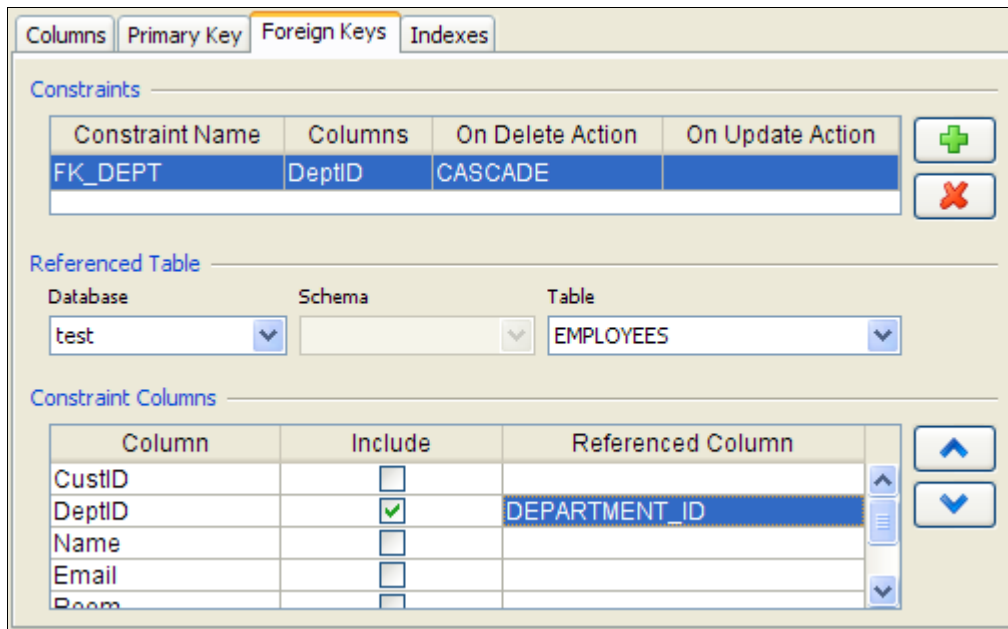


Figure: Foreign Keys tab

The tab has the following sections:

- A list of foreign keys.
- Controls for selecting the table the currently selected foreign key refers to, including the database (catalog) and/or schema for the table.
- A list of all columns for the table being created.

To declare a new foreign key constraint, click the Add button next to the list of foreign keys. You can then enter a name for the foreign key in the first field in the list, and select On Delete and On Update actions from the pull-down menus. The pull-down lists include all actions that the database support, typically CASCADE, RESTRICT, NO ACTION and SET NULL. The Columns field is read-only and gets its value automatically when you select which columns to include in the key later.

Next, use the Referenced Table controls to select the table that the foreign key refers to.

Finally, check the Include checkbox for all columns in the column list that should be part of the foreign key and then select the corresponding column in the referenced table from the pull-down menu in the Referenced Column field. You can change the column order for the key with the Up and Down buttons.

To remove an existing foreign key, select it in the list in the top section and click the Remove button.

Unique Constraints tab (database-specific)

The Unique Constraints tab is only available for databases that support this constraint type. A unique constraint declares that the columns in the constraint must have unique values in the table.

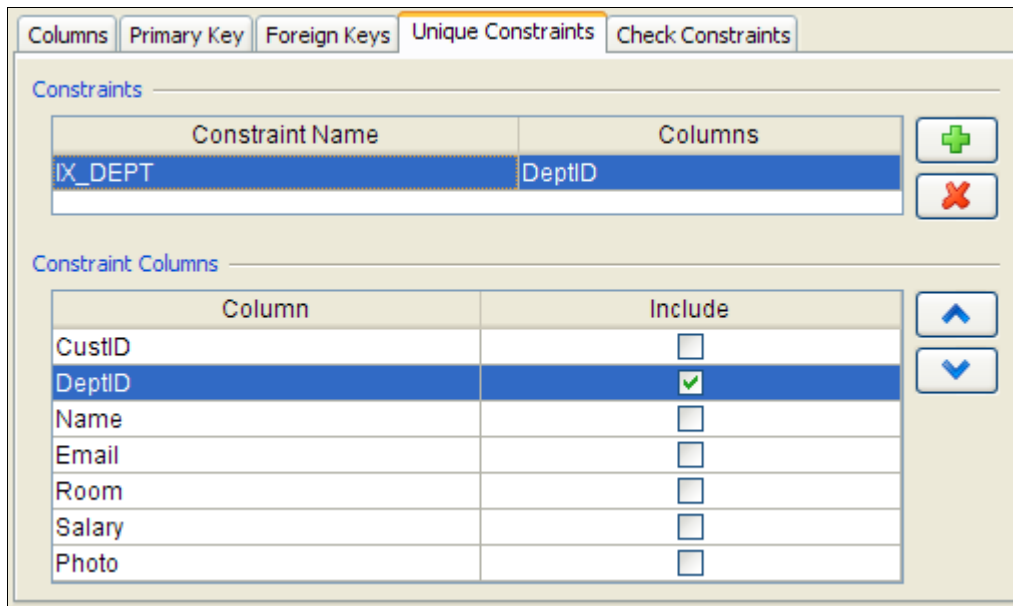


Figure: Unique Constrains tab

The top portion of the tab holds a list of all unique constraints, and the lower portion holds a list of all table columns.

To create a constraint, click the Add button and optionally enter a constraint name in the Constraint Name field. The Columns field in the constraints list is read-only, filled automatically as you include columns in the constrain. Select the columns to be part of the constraint by clicking the checkboxes in the Include field in the columns list.

You can change the order of the columns in the constraint by selecting a column and move it using the Up and Down buttons.

To remove an existing constraint, select it in the list in the top section and click the Remove button.

Check Constraints tab (database-specific)

The Check Constraints tab is only available for databases that support this constraint type. Check constraints declare that a column value fulfills a certain condition when a row is inserted or updated. Some databases uses check constraints to enforce nullability rules, so when you alter a table (as described later), you may see auto-generated check constraints for columns that you marked as not allowing null values in the Columns tab.

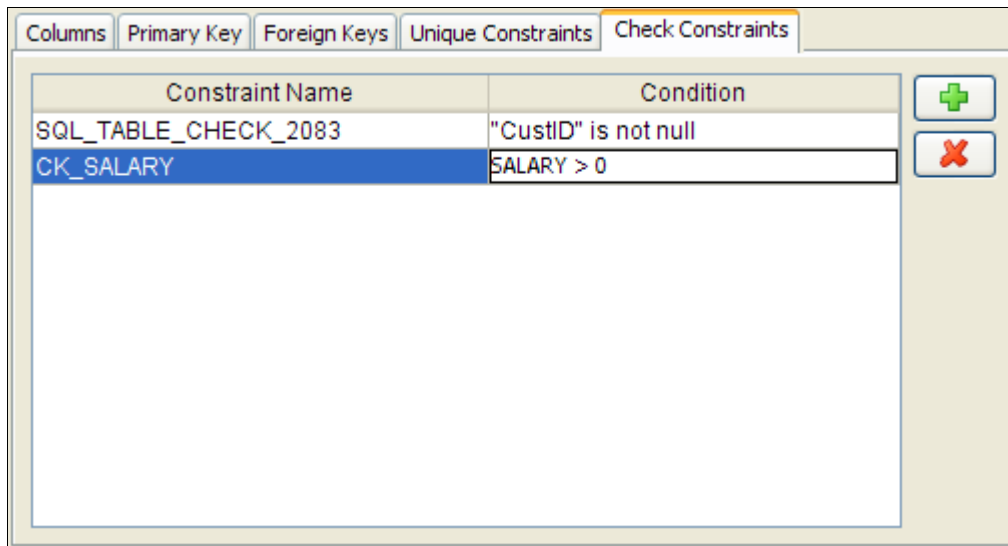


Figure: Check Constrains tab

To create a check constraint, click the Add button and optionally enter a constraint name in the Constraint Name field. Enter the condition for the column in the Condition field. You can use the same type of conditions as you use in a SELECT WHERE clause.

To remove an existing constraint, select it in the list and click the Remove button.

Indexes tab (MySQL only)

The Indexes tab is only used for the MySQL database, as a replacement for the Unique Constraints tab. The reason is that for MySQL, the CREATE TABLE statement can be used to declare both unique and non-unique indexes. MySQL also does not make a clear distinction between a unique constraint (a rule, most often enforced and implemented as an index by the database) and a unique index (primarily a database structure for speeding up queries, with the side-effect of ensuring unique column values), as most other databases do.

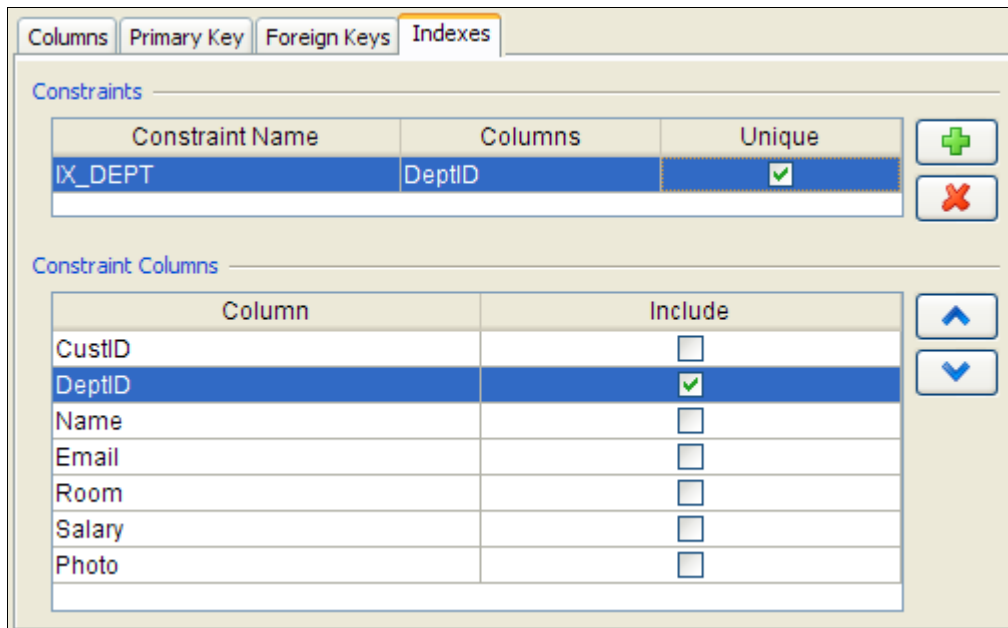


Figure: Indexes tab

The top portion of the tab holds a list of all indexes, and the lower portion holds a list of all table columns.

To create an index, click the Add button and optionally enter a name in the Constraint Name field. The Columns field in the constraints list is read-only, filled automatically as you include columns in the constraint. If you want the index columns to have unique values for all rows in the table, click the checkbox in the Unique field.

Select the columns to be part of the index by clicking the checkboxes in the Include field in the columns list. You can change the order of the columns in the constraint by selecting a column and move it using the Up and Down buttons.

To remove an existing constraint, select it in the list in the top section and click the Remove button.

SQL Preview

The SQL Preview area is updated automatically to match the edits made in the assistant. The preview is read only, but you can copy the SQL to the SQL Commander and flip between formatted and unformatted views using the two buttons in the toolbar above the preview area.

Execute

When you are satisfied with the table declaration, click the Execute button to create it.

Alter Table

To alter a table, select the table node in the objects tree and launch the Alter Table assistant from the Database main menu, the Database Objects tree right-click menu, or from the Actions menu button in the object view.

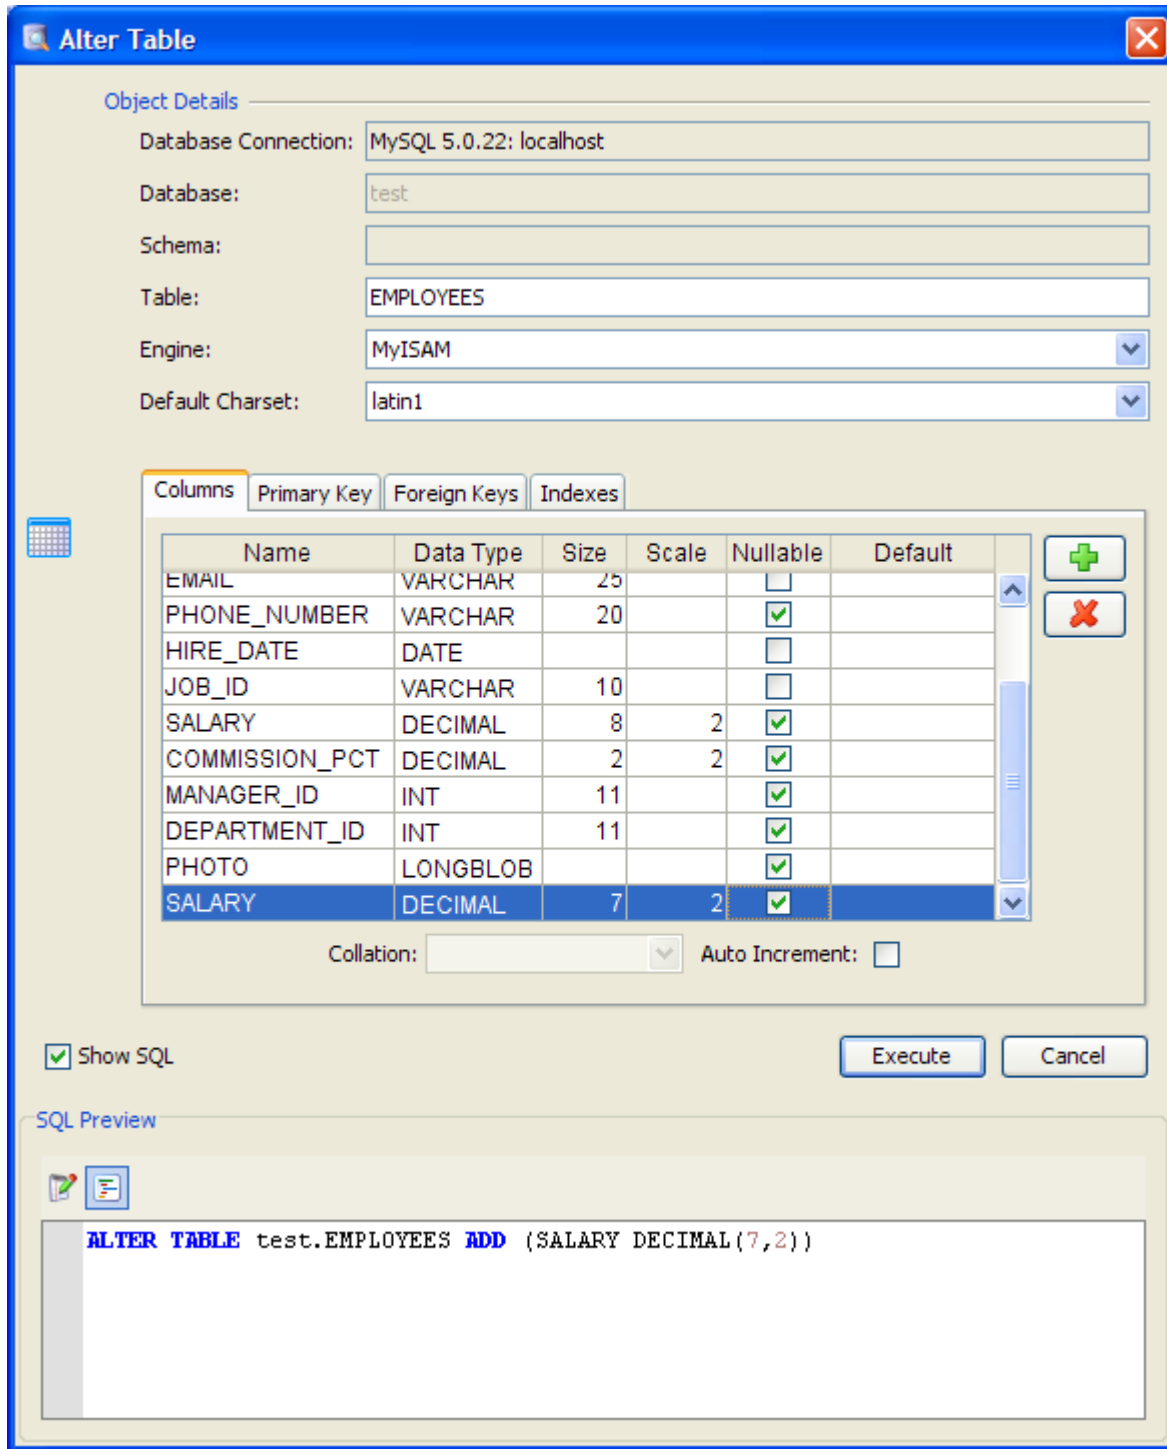


Figure: The Alter Table assistant

The Alter Table assistant has exactly the same layout as the Create Table assistant, with all information about the table you wish to alter shown when you launch it. As you make changes, such as adding a column, the SQL Preview area shows the corresponding ALTER TABLE statements. See the [Create Table section](#) for descriptions of all parts of the assistant.

The controls, such as the fields, pull-down menus and buttons, in the assistant are only enabled if the ALTER TABLE statement for the database holding the table provides a way to alter the corresponding table attribute. For instance, for a database only allows the size of a VARCHAR column to be altered, the Size field in the Columns tab is disabled for all columns with other data types. If you find that you can not make the change you want, it is because the ALTER TABLE statement does not allow that change to be made.

Create Index

The Create Index assistant helps you create an index for a table. Select the table node in the objects tree and launch the Create Index assistant from the Database main menu, the Database Objects tree right-click menu, or from the Actions menu button in the object view.

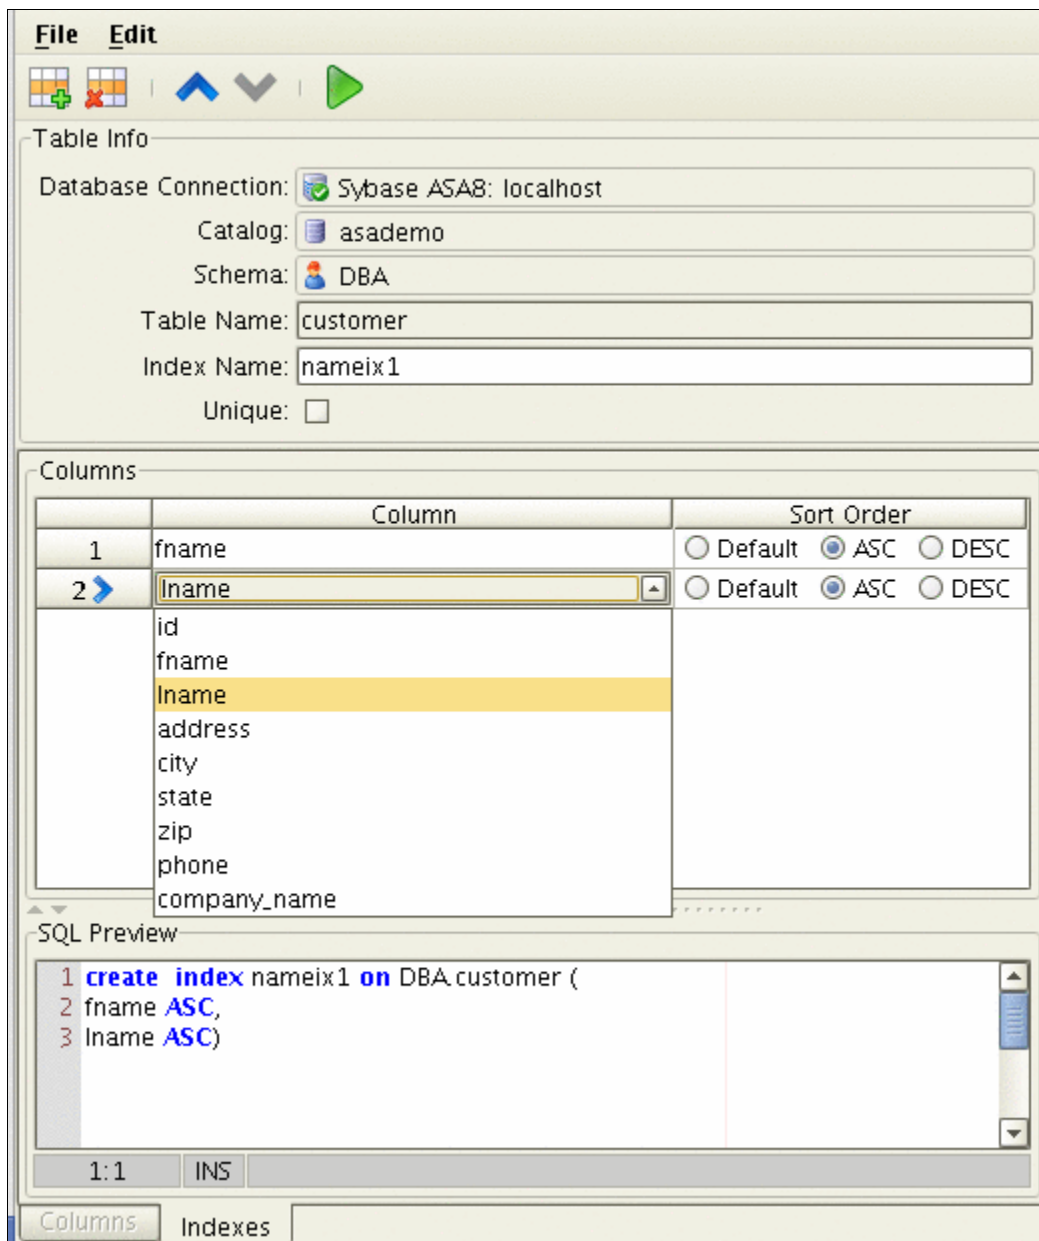


Figure: The Create Index assistant

The Create Index assistant is organized in three areas from the top:

- **Table Info**
Specifies the owning database connection, database and/or schema, and the table name. These are picked up from the selection in the tree when the assistant is started. The Index Name field is empty and must be specified. If you want the index columns to have unique values for all rows in the table, click the checkbox under the Index Name field.
- **Columns**
The list is used to select the columns include in the index.
- **SQL Preview**
The SQL previewer instantly shows the SQL that is used to create the index.


Columns

Add columns by using the Edit->Insert menu choice and Edit->Delete to remove the currently selected row. They can be re-organized using the Edit->Move Up and Edit->Move Down menu operations.

The Column field has a pull-down menu with the names all columns in the table. Columns can be re-ordered using the main controls. Use the radio buttons to select a sort order for the column in the index.

SQL Preview

The SQL Preview area is updated automatically to match the edits made in the assistant. The preview is read only.



```
SQL Preview
1 create table MyFirstTable (
2 CustId INT not null,
3 DeptId INT not null,
4 Name VARCHAR(20) not null,
5 Email VARCHAR(30) not null,
6 Room VARCHAR(10),
7 Salary DECIMAL(7, 2) not null,
8 Photo LONGBLOB ,
9 primary key (
10 CustId
11 ))
```

1:1 INS

Figure: The SQL Preview for an index

Execute

When you are satisfied with the index declaration, click the Execute button to create it.

Edit Table Data

Introduction

The table data editing feature mimics how editing is performed in standard spreadsheet applications; just click a cell value and edit. Edits are saved in a single database transaction which ensures that all or no changes are committed. The editing feature supports saving binary and large text data and it even presents common data formats in their respective viewers, such as image viewer, XML, HEX, etc.

A block of data can easily be interchanged via standard copy and paste operations between the grid editor and other applications, such as Microsoft Excel, StartOffice and OpenOffice.

Editing is primarily performed in the grid editor. For some data, such as binary or large formatted text data, editing in the grid editor is not optimal, so for these situations, we recommend that you to use the form or cell data editors. The form editor presents a single row of data in a separate window, organized as a form with the column name in the first column and the data in the second column. All editing capabilities in the grid editor are also available in the form editor. The cell editor is used to edit a single cell value in a separate window. This is useful when editing formatted text data or to browse binary data.

Most of the editing functions have a key binding assigned. Check the right click menu in the data grid to find out what they are.

Features that support editing

Editing of table data can be performed in the Database Objects->Data tab and in a result set generated by a SELECT statement in the SQL Commander.

There are a few conditions that must be fulfilled for editing to be enabled in the SQL Commander:

1. It is a result set
2. The SQL is a SELECT command
3. Only one table is referenced in the FROM clause
4. All current columns exist by name (case sensitive) in the identified table

The editing tool bar is hidden if these rules are not met.

Update and Delete must match one table row





The editing features in DbVisualizer ensure that only one row in the table will be affected by update and delete edits. This prevents the user from doing changes in one row that might also silently affect data in other rows. DbVisualizer uses the following strategies to determine the uniqueness of the edited row:

1. Primary Key
2. Unique Index
3. Manually Selected Columns

The Primary Key concept is widely used in databases to uniquely identify the key columns in tables. If the table has a primary key, DbVisualizer will use it. There are situations when primary keys are not supported by a database or when primary keys are supported but not used. If no primary key is defined, DbVisualizer will check if there is a unique index. If there are several unique indexes, DbVisualizer will pick one of them. If there is no primary key or any unique indexes defined for the table, you need to manually choose what columns to use. The key column chooser is automatically displayed if the key columns can't be determined automatically.

Edit Multiple Rows

The grid editor supports editing multiple rows and saving all changes in one database transaction. Edited rows are indicated with an icon in the row header:

-  Cell(s) in the row has been edited
-  Row is new
-  Row is duplicated from another row
-  Row is marked for deletion (edit is not allowed)

Edits are saved when explicitly saving changes via the Edit Table Data->Save Edit(s) right click menu choice or via the Save button in the tool bar.

Data Type checking

When leaving an edited cell the new value is validated with the data type for the column. If there is an error, the following dialog is displayed.

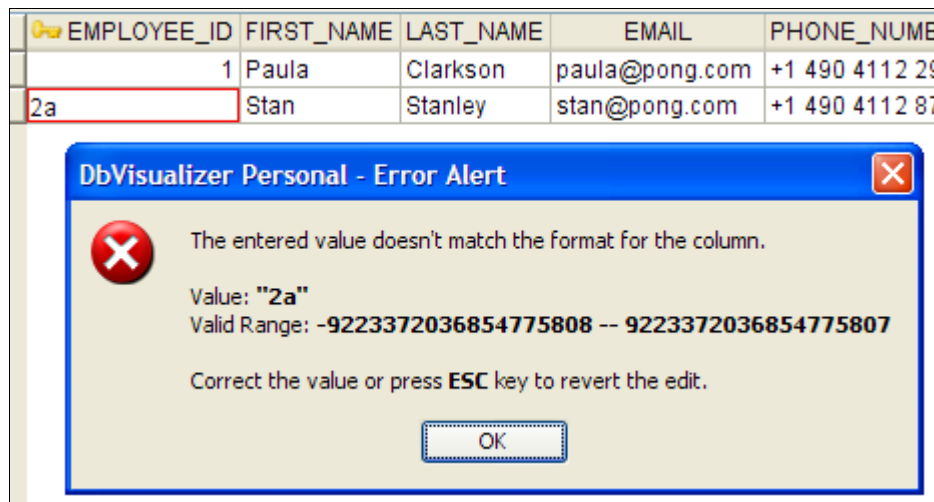


Figure: Data type error

New Line and Carriage Return

If a cell in the grid editor contains new line, carriage return or tab characters, these are not visually represented in the grid. These characters instead appear as `\n` (newline), `\r` (carriage return) and `\t` (tab) when you edit the cell value.

	Id	Name
1	1	PaulBurtAndersonis a good human being
2	2	Paul\r\nBurt\r\nAnderson\tis a good human being

Figure: Carriage Return, Line Feed and Tab characters

The figure show identical values in the Name column. The second is in edit mode and the special characters are

visible, while in the first non editing row they are not visible.

A better solution for editing multi-lined data is with the cell editor, as it handles these characters more naturally:

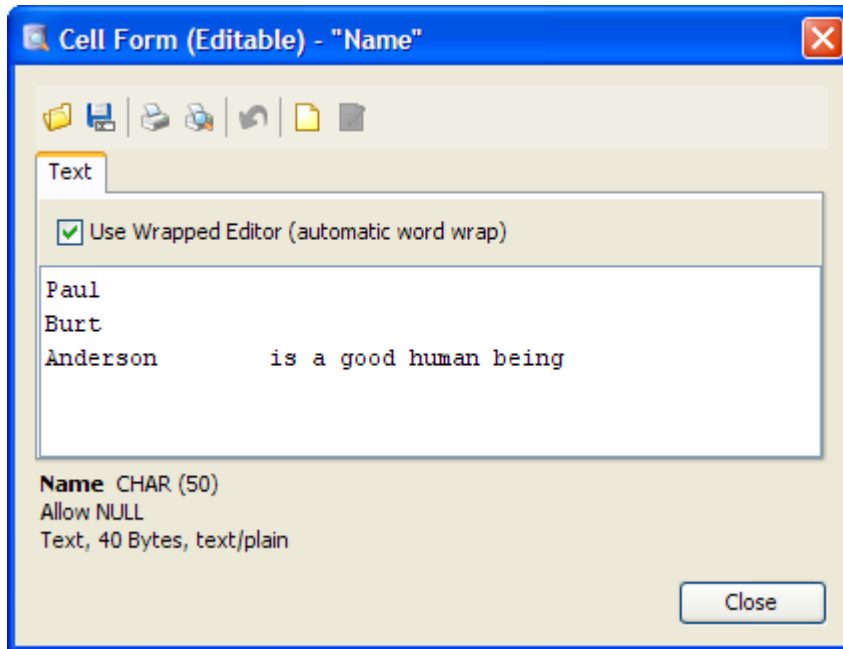


Figure: Editing multi lined text data in the cell editor

Grid Editor

The grid editor tool bar is decorated with buttons for editing and the right-click menu contains all related operations.

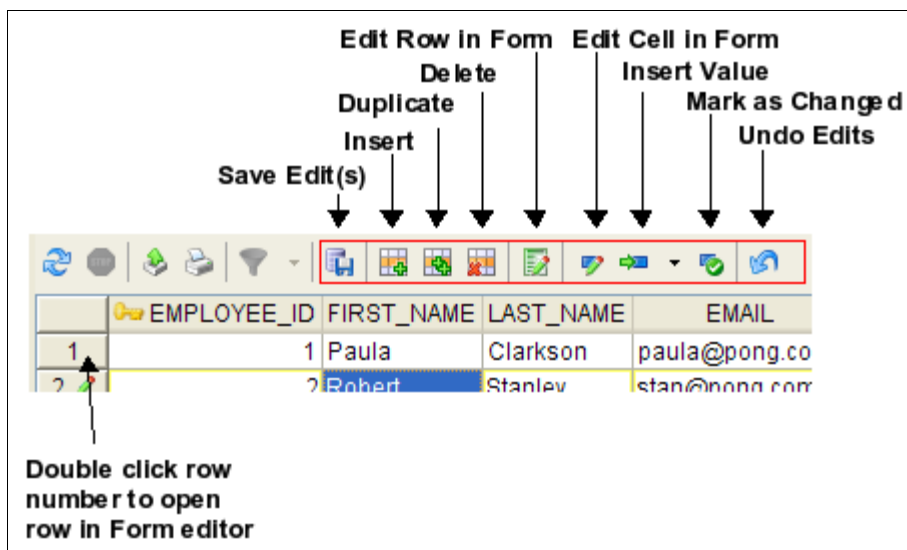


Figure: Toolbar buttons to control the grid editor

Cells that have been edited are indicated by a yellow background color. Only these columns will be updated

when the changes are saved in the database. All cells in the edited row are highlighted with a yellow border to indicate that some cell(s) in the row has been edited.

Insert row

To insert a new row, choose the Edit Table Data->Insert New Row right-click menu choice or press the Insert toolbar button. The new row will be inserted after the selected row or at the top if no row is selected. You can now start editing the cells in the grid or open the form editor to insert values.

Update row

To update a row, just double click in the target cell and modify the value.

Delete row(s)

One or multiple rows are marked for deletion via the Edit Table Data->Delete Row(s) right-click menu choice or by pressing the Delete toolbar button. Each deleted row will be highlighted with a red background color and no further editing of the content is allowed.

Deleting a row that has been updated will automatically undo all edits and show the original values. This is done so that it is obvious which data will be deleted. Deleting a row that has been inserted (or via duplication) will be removed from the grid.

Duplicate row(s)

Duplicate a row or several rows by selecting the cells in the rows that should be duplicated, then choose Edit Table Data->Duplicate Row(s) or press the Duplicate tool bar button. All cells in the new row will be marked as being edited (yellow background color). The exception is any Auto Increment field, which should be assigned a value by the database.

Copy/Paste

Copying selected cell values is accomplished via the Copy Selection right-click menu choice. The data on the clipboard may then be pasted either into DbVisualizer or any external application. The copy and paste operations in the grid editor are defined by the Grid->Copy category in Tool Properties. The default setting for column and newline delimiters are sufficient for most uses.

Paste data from Excel and OpenOffice

The grid editor supports pasting data from the major spreadsheet applications. The grid editor support pasting single data as well as block of data.

Copy from Excel	Paste into DbVisualizer Grid
A single cell is copied in Excel	The selected cell is pasted into one selected target cell

	A	B	C	D
1	98	Roger	Wruck	jsewell@hotmail.co
2	104	Steeve	Heer	gertveters@hotmail
3	105	Marc	Meller	mintegen@hotmail.
4	102	Bernhard	Kessler	brian@yahoo.com
5	193	Luci	Young	stephen@osp.edu
6	185	Robert	Austad	kathleen@bto.com
7	183	Pawan	Yucel	michael@ware.de
8	181	Hong	Graff	mail1235@yahoo.de
9	174	Rick	Stewart	jay@tury.com
10	171	Michael	Palm	mike@gmail.com
11	154	Peter	Pullabhotla	dsimmons@at123.i

==>

	EMPLOYEE_ID	FIRST_NAME	LAST_NAME
1	98	Roger	Wruck
2	104	Steeve	Heer
3	105	Marc	Meller
4	102	Bernhard	Kessler
5	193	Luci	Young
6	185	Robert	Austad
7	183	Pawan	Yucel
8	181	Hong	Kessler
9	174	Rick	Stewart
10	171	Michael	Palm

A single cell is copied in Excel

The selected cell is pasted into multi selected target cells

	A	B	C	D
1	98	Roger	Wruck	jsewell@hotmail.co
2	104	Steeve	Heer	gertveters@hotmail
3	105	Marc	Meller	mintegen@hotmail.
4	102	Bernhard	Kessler	brian@yahoo.com
5	193	Luci	Young	stephen@osp.edu
6	185	Robert	Austad	kathleen@bto.com
7	183	Pawan	Yucel	michael@ware.de
8	181	Hong	Graff	mail1235@yahoo.de
9	174	Rick	Stewart	jay@tury.com
10	171	Michael	Palm	mike@gmail.com
11	154	Peter	Pullabhotla	dsimmons@at123.i

==>

	EMPLOYEE_ID	FIRST_NAME	LAST_NAME
1	98	Roger	Wruck
2	104	Steeve	Heer
3	105	Marc	Kessler
4	102	Bernhard	Kessler
5	193	Luci	Kessler
6	185	Robert	Kessler
7	183	Pawan	Kessler
8	181	Hong	Graff
9	174	Rick	Stewart
10	171	Michael	Palm

A block of cells is copied from Excel

The block is pasted into the selected region

	A	B	C	D
1	98	Roger	Wruck	jsewell@hotmail.co
2	104	Steeve	Heer	gertveters@hotmail
3	105	Marc	Meller	mintegen@hotmail.
4	102	Bernhard	Kessler	brian@yahoo.com
5	193	Luci	Young	stephen@osp.edu
6	185	Robert	Austad	kathleen@bto.com
7	183	Pawan	Yucel	michael@ware.de
8	181	Hong	Graff	mail1235@yahoo.de
9	174	Rick	Stewart	jay@tury.com
10	171	Michael	Palm	mike@gmail.com
11	154	Peter	Pullabhotla	dsimmons@at123.i

==>

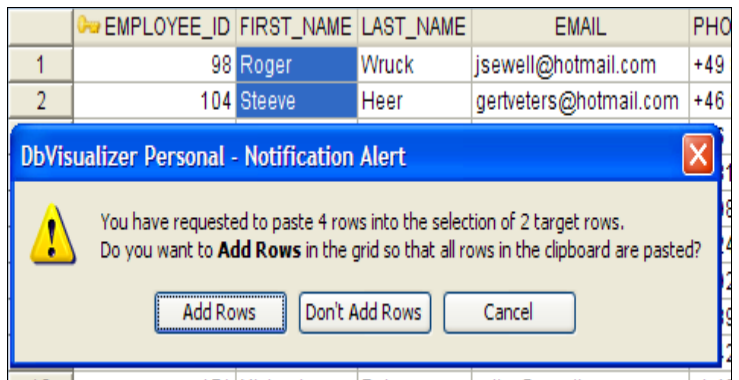
	EMPLOYEE_ID	FIRST_NAME	LAST_NAME
1	98	Roger	Wruck
2	104	Marc	Meller
3	105	Bernhard	Kessler
4	102	Luci	Young
5	193	Robert	Austad
6	185	Robert	Austad
7	183	Pawan	Yucel
8	181	Hong	Graff
9	174	Rick	Stewart
10	171	Michael	Palm

A block of cells copied from Excel

The block is pasted into a non equal number of target cells

==>

	A	B	C	D
1	98	Roger	Wruck	jsewell@hotmail.co
2	104	Steeve	Heer	gertveters@hotmail
3	105	Marc	Meller	mintegen@hotmail.
4	102	Bernhard	Kessler	brian@yahoo.com
5	193	Luci	Young	stephen@osp.edu
6	185	Robert	Austad	kathleen@bto.com
7	183	Pawan	Yucel	michael@ware.de
8	181	Hong	Graff	mail1235@yahoo.de
9	174	Rick	Stewart	jay@tury.com
10	171	Michael	Palm	mike@gmail.com
11	154	Peter	Pullabhotla	dsimmons@at123.i



Insert pre-defined values (Set Selected Cells)

The Edit Table Data->Set Selected Cells right-click menu choice or the Set Selected Cells tool bar button lists a few pre-defined functions that will fill the selected cells with data.

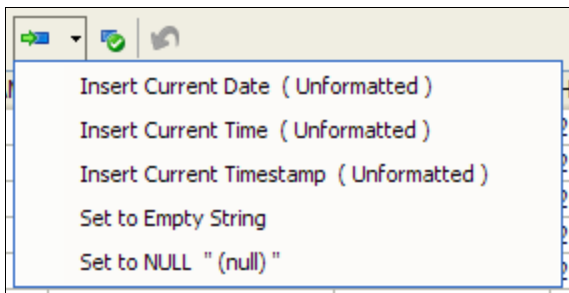


Figure: Set Selected cells functions

Use these to insert data into the selected cells. Note that the target column type must accept the selected function. Nothing will happen if choosing for example "Insert Current Time" into a DATE data type.

Undo Edit(s)

The Edit Table Data->Undo Edit(s) operation is used to revert all edits in the selected cell(s). Reverting all cells in a row that are marked as Insert or Duplicate will remove the complete row from the grid while a Delete marked row is cleared from its delete state. Undoing updated cells simply reverts the changes to the original values.

Key Column(s) Chooser

Normally database tables have a primary key or at least one unique index. If this is the case, editing is no problem. If there is no way to uniquely identify rows in the table, you need to manually define what columns DbVisualizer should use. While saving the changes DbVisualizer will check that there is a way to identify unique rows in the table. If it cannot be accomplished, the following window is displayed.

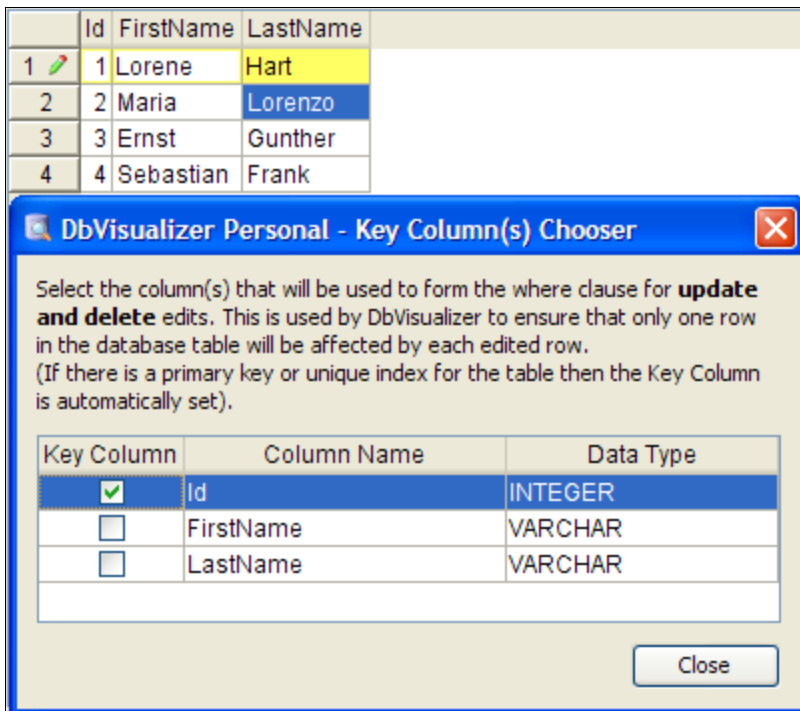


Figure: Key Column(s) Chooser

The key column chooser can be manually opened via the Edit Table Data->Key Column Chooser right-click menu choice.

If the database request to save the edits cannot uniquely identify the single row that should be changed, the error dialog is displayed and the editing state is kept for that row in the grid editor.

Preview Changes

You may preview the SQL statements that will be executed when choosing to Save the edits. It is displayed via the Edit Table Data->SQL Preview right-click menu choice.

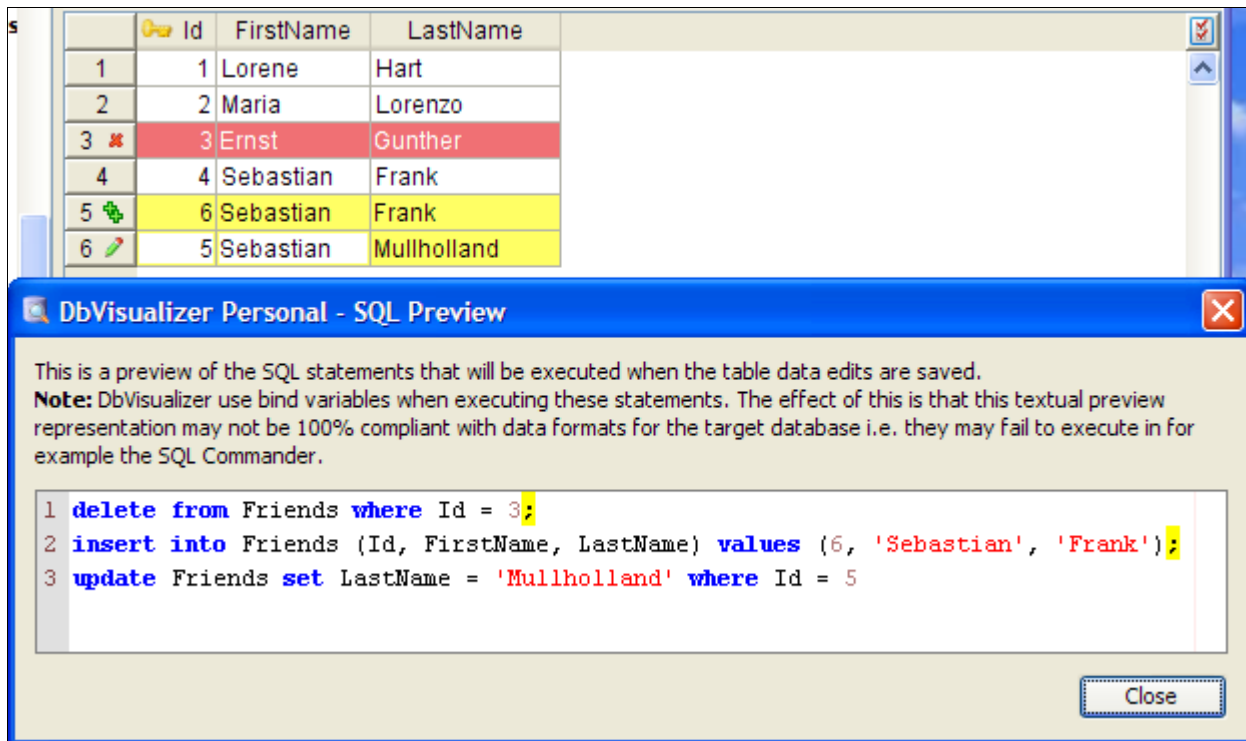


Figure: SQL Preview

(The listed SQL statements may not be 100% compliant as the save process use variable binding to pass values to the database).

Saving Changes

To save table data edits, select the Edit Table Data->Save Edit(s) right-click menu choice or the Save toolbar button. If there are rows that has been edited or deleted, these are first checked so that there is only one database table row affected by each edited row. If this pass is successful DbVisualizer will save the changes in the table. The progress is displayed in the status bar and Save may be interrupted by pressing the Stop button in the toolbar. While save is in progress, no other operation may be performed in DbVisualizer, i.e., the rest of the application is locked.

Transaction Control

DbVisualizer use the physical root connection for the actual database connection when saving table data. Once save is requested, DbVisualizer will implicitly set the auto commit state to off and reset it to what it was prior to requesting save when saving is finished. If the Use Single Shared Physical Database Connection is enabled in connection properties, DbVisualizer will check whether there are any uncommitted updates in the database when save is requested. If there are uncommitted changes you must first commit or rollback these changes before save is started.

Saving table data edits are batched in a single transaction. There is no restriction of number of edits that may be saved in a single save operation, but the database server may put either explicit or implicit restrictions. The connection property Physical Connection->Transaction->Commit Batch Size specifies how many edits should be performed in the database table until commit is automatically initiated. If for example saving 150 edited rows and an error occur while saving the 121:st row, then the first 100 rows will have been committed and the rest are left unchanged. The visual indication in the grid after a non complete save operation is that rows that weren't saved keep their original editing state indicator. Rows that were saved properly are indicated with a green checked

cylinder icon.

	PADATE	TDATE	DESCRIPTION	S
95	2005-05-10	2005-05-10 00:52:34.0	Program Registration	-
96	2005-05-10	2005-05-10 10:59:15.0	Program Registration	-
97	2005-05-10	2005-05-10 17:04:42.0	Program Registration	-
98	2005-05-10	2005-05-10 17:20:16.0	Program Registration	-
99	2005-05-10	2005-05-10 19:51:18.0	Program Registration	-
100	2005-05-10	2005-05-10 20:48:33.0	Program Registration	-
101	2005-05-11	2005-05-11 07:24:41.0	Program Registration	-
102	2005-05-11	2005-05-11 09:31:38.0	Program Registration	-
103	2005-05-11	2005-05-11 11:18:05.0	Program Registration	-
104	2005-05-12	2005-05-12 11:49:18.0	Program Registration	-
105	2005-05-12	2005-05-12 16:38:52.0	Program Registration	-

Figure: Saved rows state

The cylinder icon with the green check mark indicate that the row was saved in the database table. Normally the grid is reloaded after a successful save operation and there is no cylinder icon displayed. It only appear if the save operation was partly successful. Rows that weren't saved are still represented with the original editing state icon and you may request save one more time.

Note: Rows that has been properly saved (indicated with the cylinder icon) cannot be edited until all rows are saved properly or the grid is reloaded.

Permissions

All of the insert, update and delete requests performed by the grid editor may be confirmed before being executed by the database server. Specify in [Tool Properties->Permissions](#) the confirmation state. The default behavior is that delete operations must be confirmed while insert and update need no confirmation.

Errors

If a database error occur while saving changes, details about the errors are displayed in a window along with the actual SQL that was executed.

Form Editor/Viewer

The Form Viewer is available in the right click menu for all grids in DbVisualizer. It is used to browse information and to present binary data in viewers.

The Form Editor add editing capability in the form viewer. This editor is useful when inserting new rows and when it is important to get a more balanced overview of all the data.

The form editor "rotate" the data in one row and presents it in a form with the column name as label. All edits made in the form editor are reflected in the grid with the edited state icon being updated along with new values. Saving edits in the database is always done in the grid editor.

Open the form editor via the Edit Row in Form right-click menu choice, via the button in the toolbar or by double clicking the row number header.

	EMPLOYEE_ID	FIRST_NAME	LAST_NAME	EMAIL	PHOTO
1	101	Roger	Bjärevall	roger@minq.se	BINARY, 409 679 Bytes
2	104	Ulf	Andersson	uffe@minq.se	BINARY, 514 969 Bytes
3	105	Hans	Bergsten	hans@minq.se	BINARY, 4 358 Bytes
4	102	Jan	Andersson	ianne@minq.se	BINARY, 532 348 Bytes

Figure: A row in the grid

Here is the same row as selected in the previous screen shot displayed in the row form window.

Key	Name	Value
	EMPLOYEE_ID	101
	FIRST_NAME	Roger
	LAST_NAME	Bjärevall
	EMAIL	roger@minq.se
	PHONE_NUMBER	+46 8 55659250
	HIRE_DATE	1998-07-01
	JOB_ID	2
	SALARY	12419.00
	COMMISSION_PCT	0.91
	MANAGER_ID	0
	DEPARTMENT_ID	4
	PHOTO	 BINARY, 409 679 Bytes, image/png
	CV	(null)

JOB_ID VARCHAR (10)
Not NULL

Close

Figure: Form viewer

The Key optionally show with an icon what columns are primary key columns and the Name is the column name in the grid. None of Key or the Name can be edited. Edit the values in the form in the same way as these are edited in the grid editor.

The form viewer presents images as thumbnails. The size of these are controlled in Tool Properties->Form Viewer->Image Thumbnail Size. To see the original size of an image, open the cell in the cell viewer either by

selecting Edit in Cell Window in the right-click menu, the toolbar button or by double clicking on the image.

Cell Editor/Viewer

The Cell Viewer is available in the right-click menu for all grids in DbVisualizer. It present the data for a single cell (column in a row) in a window. If the data is recognized it will be presented in a corresponding viewer:

- Image viewer
- XML viewer
- Serialized Java object viewer
- Hex viewer
- Text viewer

The cell viewer allows saving data to file and print.

The Cell Editor adds editing capability in the cell viewer. You may import data from a file or manually change the text in a text editor.

Opening an image in the cell editor will display the following window.



Figure: Cell editor

The tabs at the top shows the available viewers for the current data. When you load a file into the cell editor, the tabs may change to reflect the newly loaded data. To nullify the cell value, press the Set Value to NULL toolbar button.

Read more about binary and formatted text data in the following chapters.

Binary/BLOB

Due to the nature of binary/BLOB and CLOB data, cells of these types can only be fully modified and viewed in the cell editor. (There is partial support in the form editor to view image data and to load from file). Editing binary data can be done by importing from a file or via the text editor.

Binary data in DbVisualizer is the generic term for several common binary database types:

- LONGVARBINARY
- BINARY
- VARBINARY
- BLOB

Image Viewers

The image viewer supports displaying the full size images for the following formats:

- GIF
- JPG
- PNG
- TIFF
- BMP

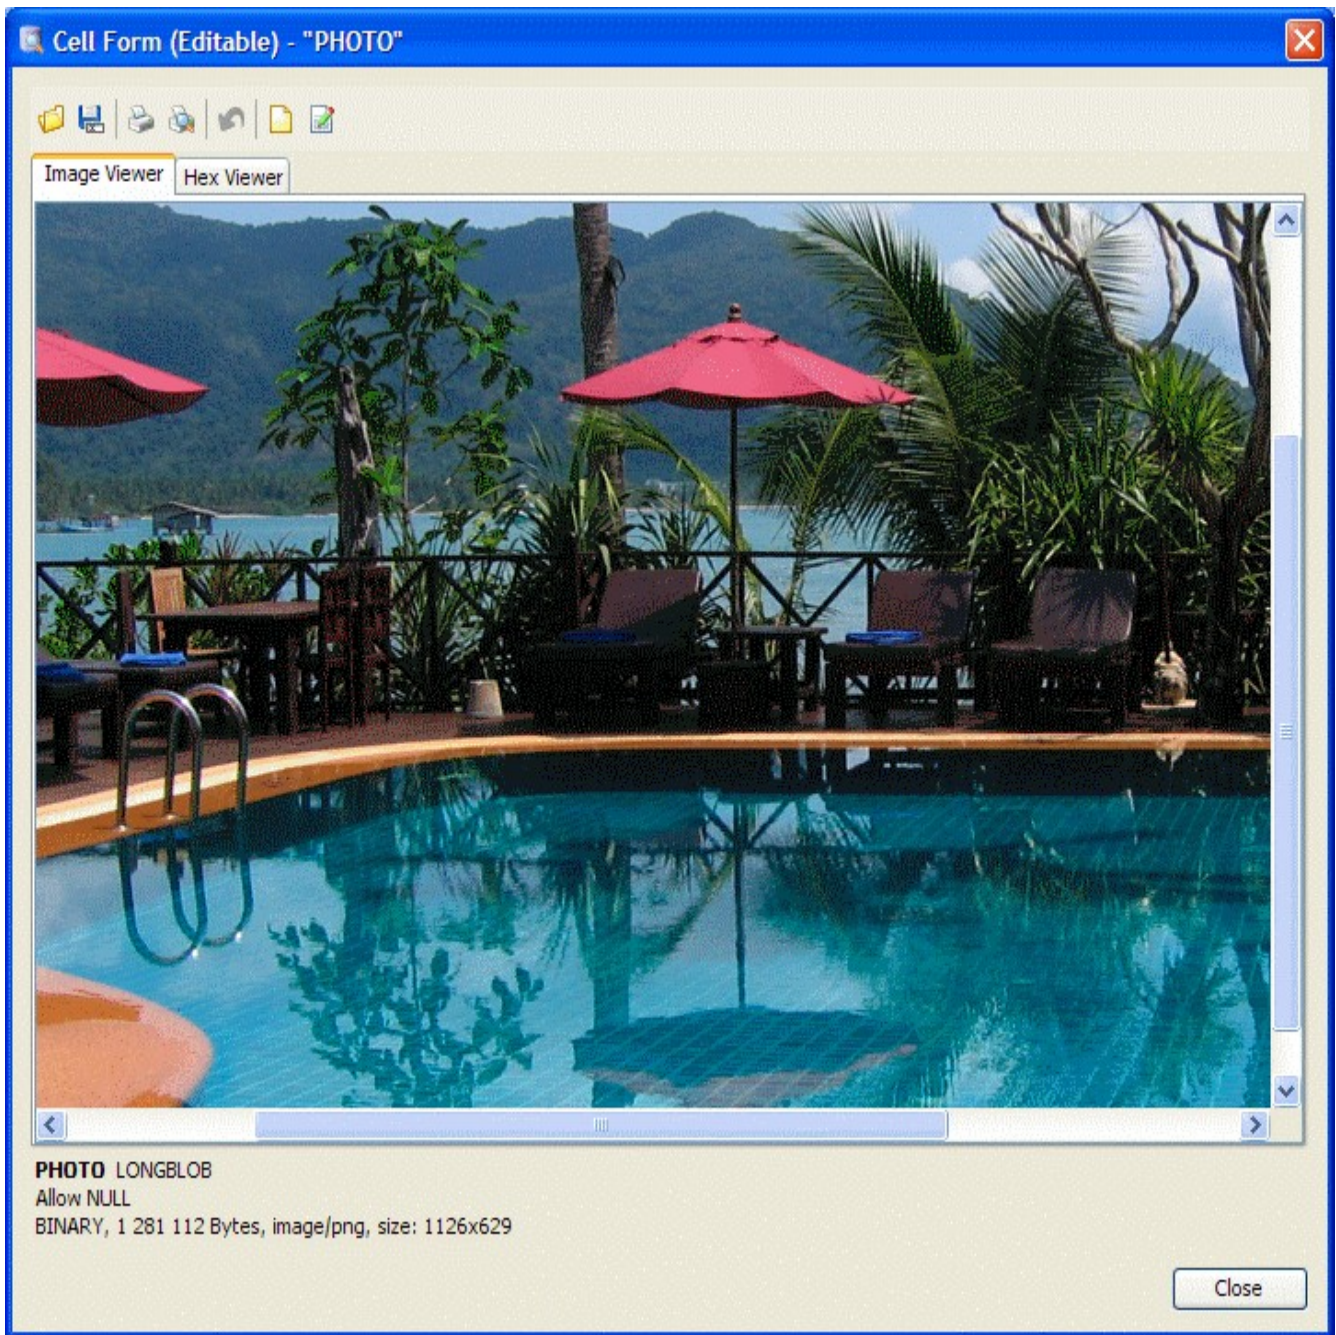


Figure: Image viewer

XML Viewer

The XML Viewer shows the content of an XML document in a tree with color highlighting. You can switch to an editable text view by pressing the Edit value in text editor tool bar button.

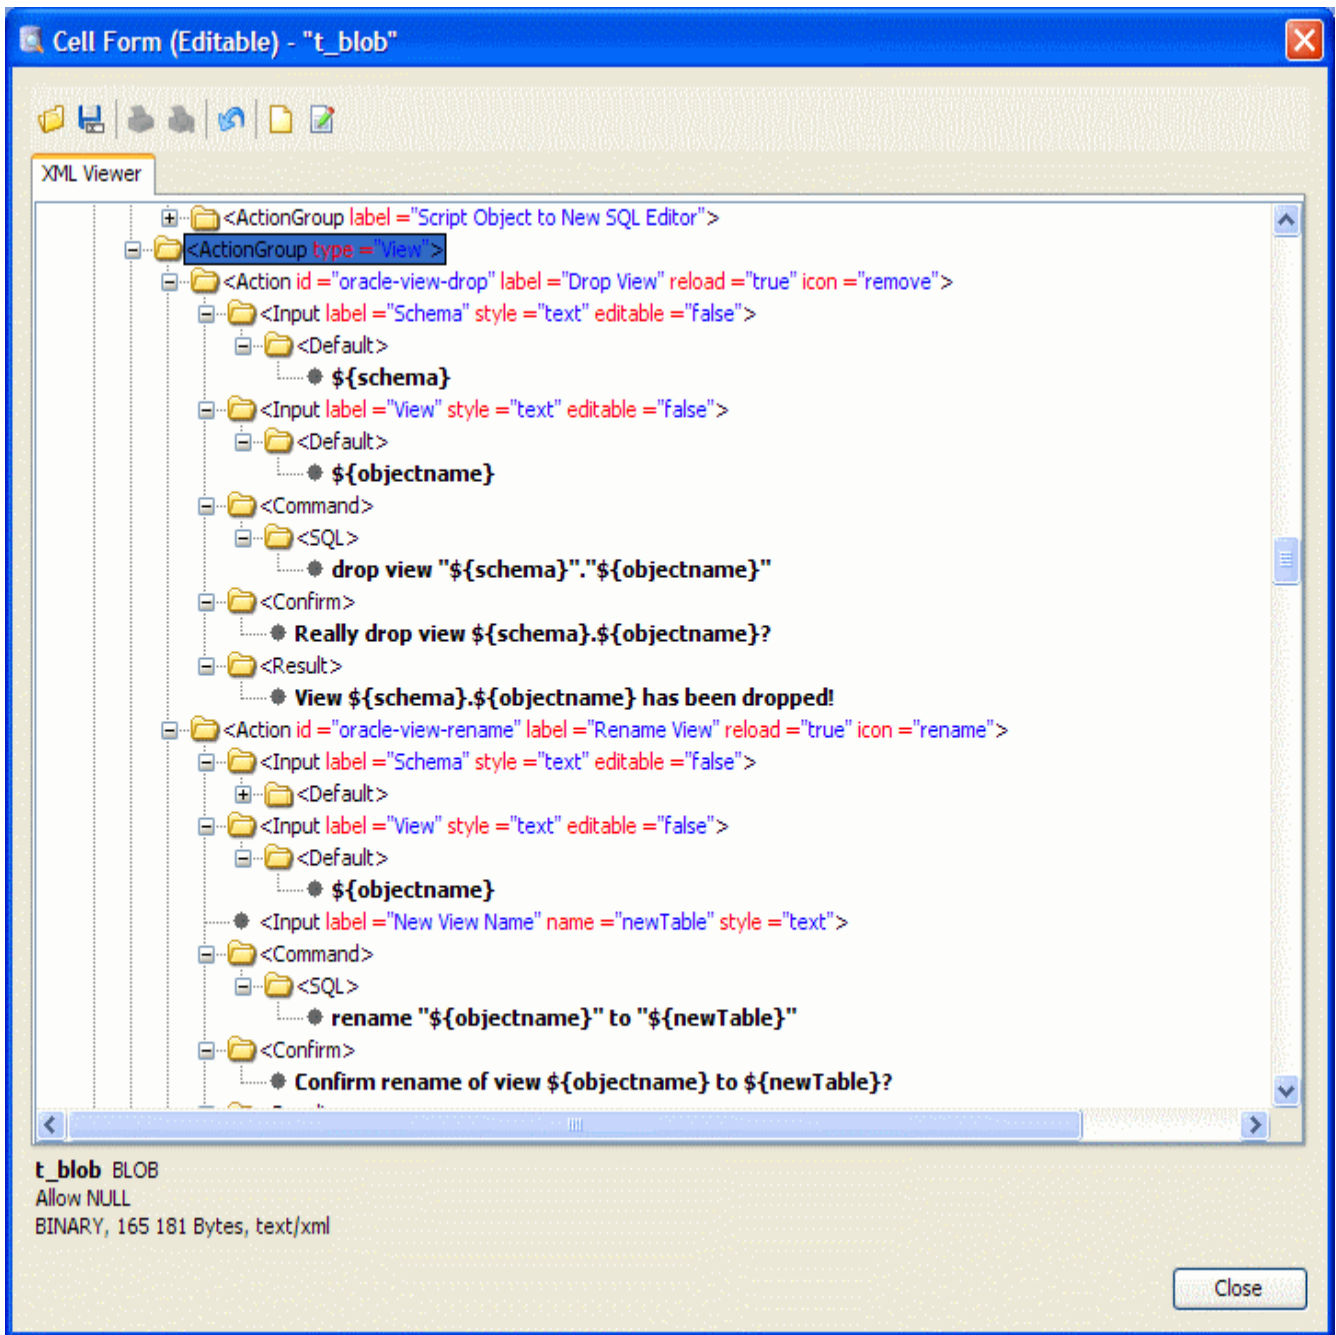


Figure: XML viewer

Serialized Java Objects Viewer

The serialized Java object viewer render a java object in a tree style. All aspects about the object may be browsed.

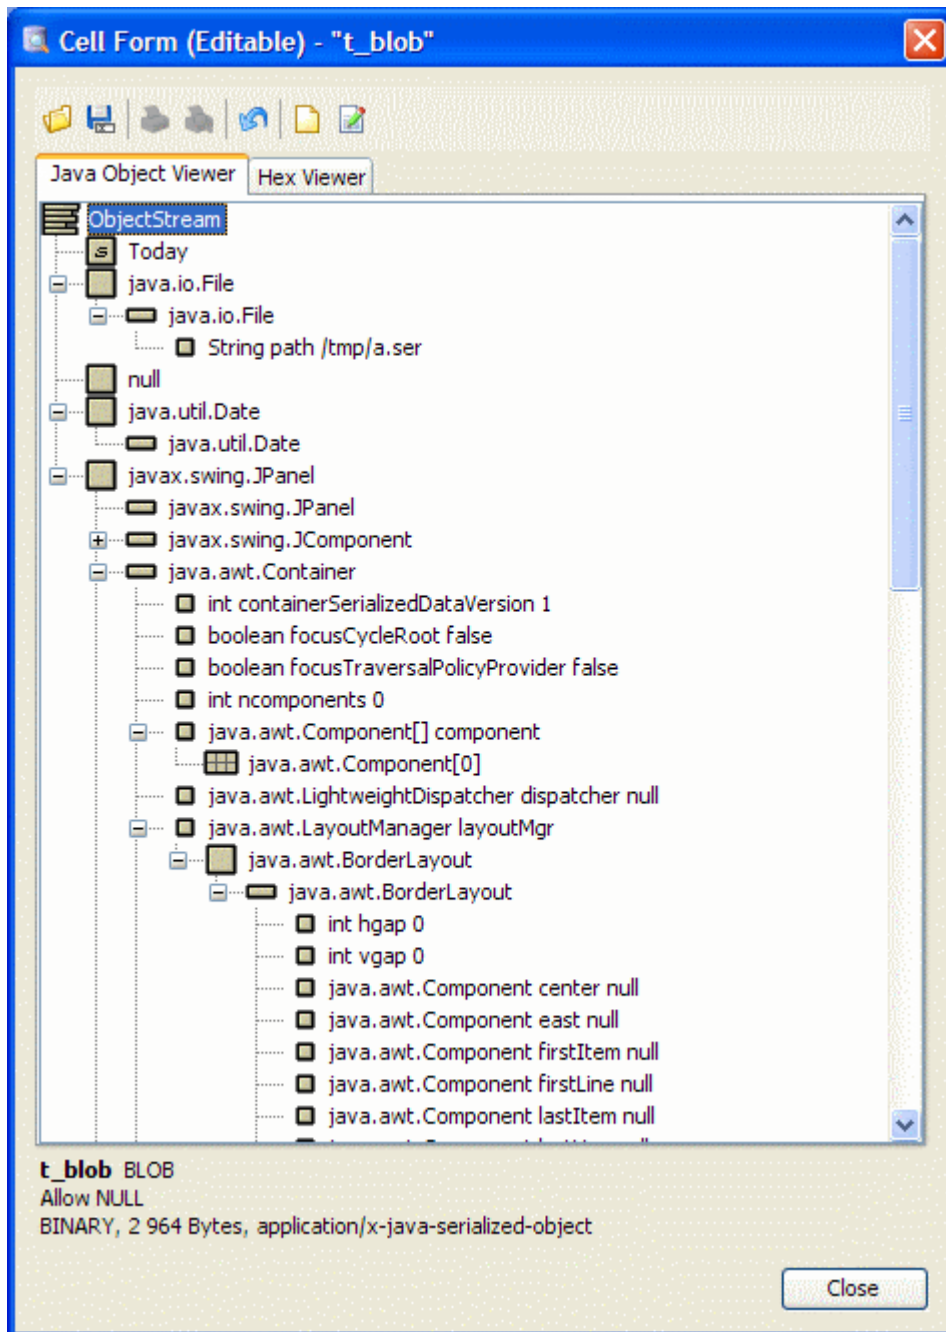


Figure: Binary data viewer for serialized Java objects

Hex Viewer

The generic Hex/Ascii viewer shows the hex representation of every byte in the data and its text representation. This is the default viewer for unknown data.

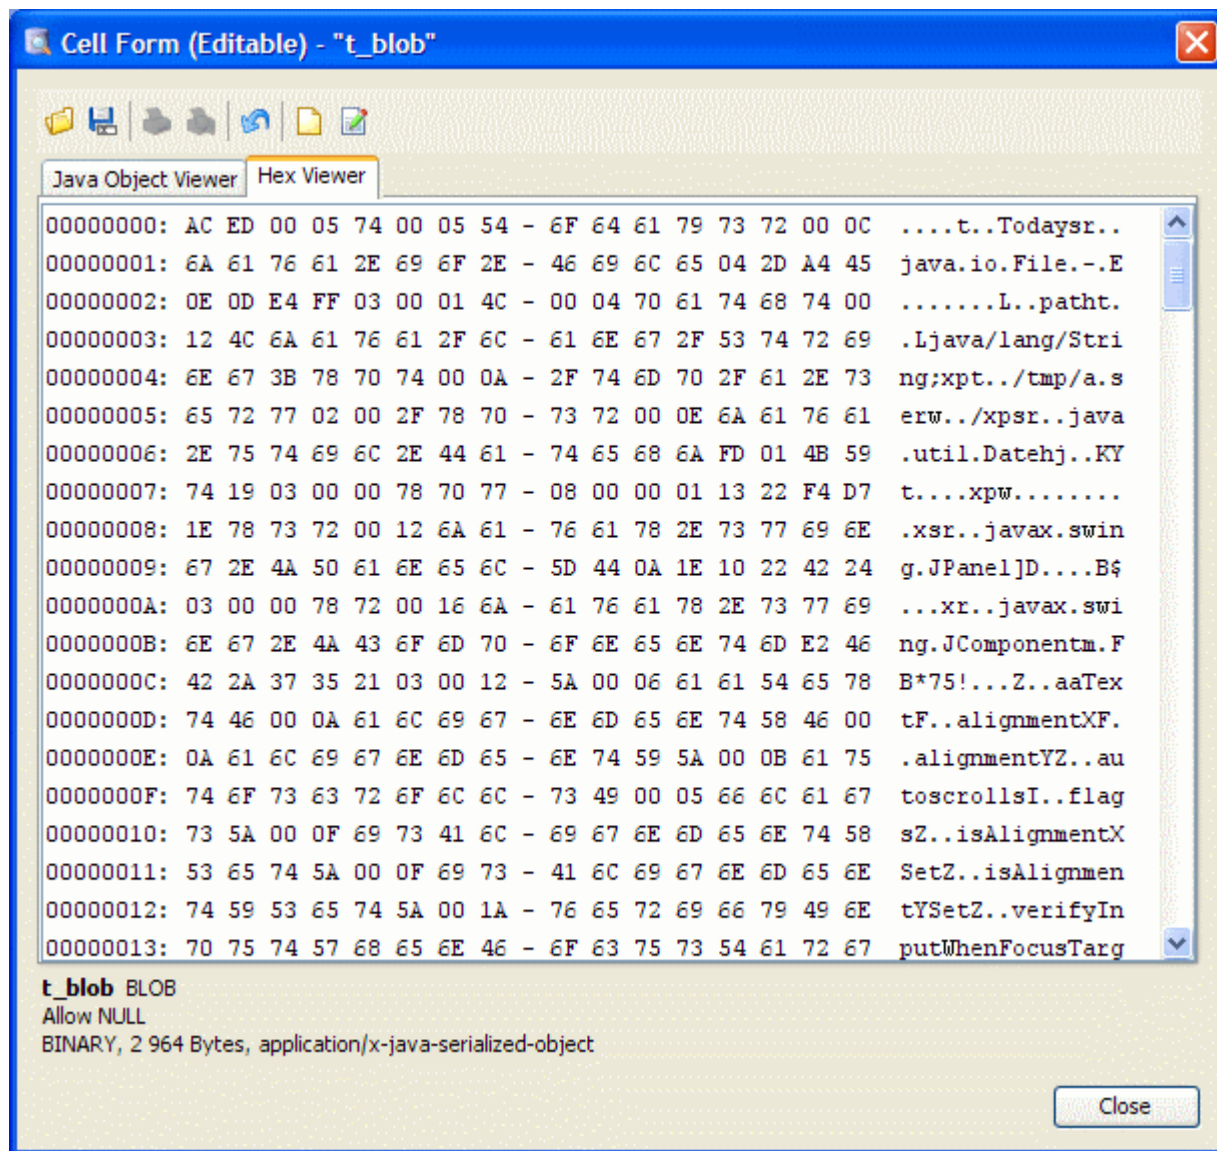


Figure: Hex/Ascii viewer

Large text data/CLOB

Large text data and CLOB data types are typically edited in the multi line text editor. For formatted data (that includes new lines), the default editor is useful. If editing a large chunk of non-formatted data, enable the Use Wrapped Editor setting and DbVisualizer will then automatically wrap the text for easy editing.

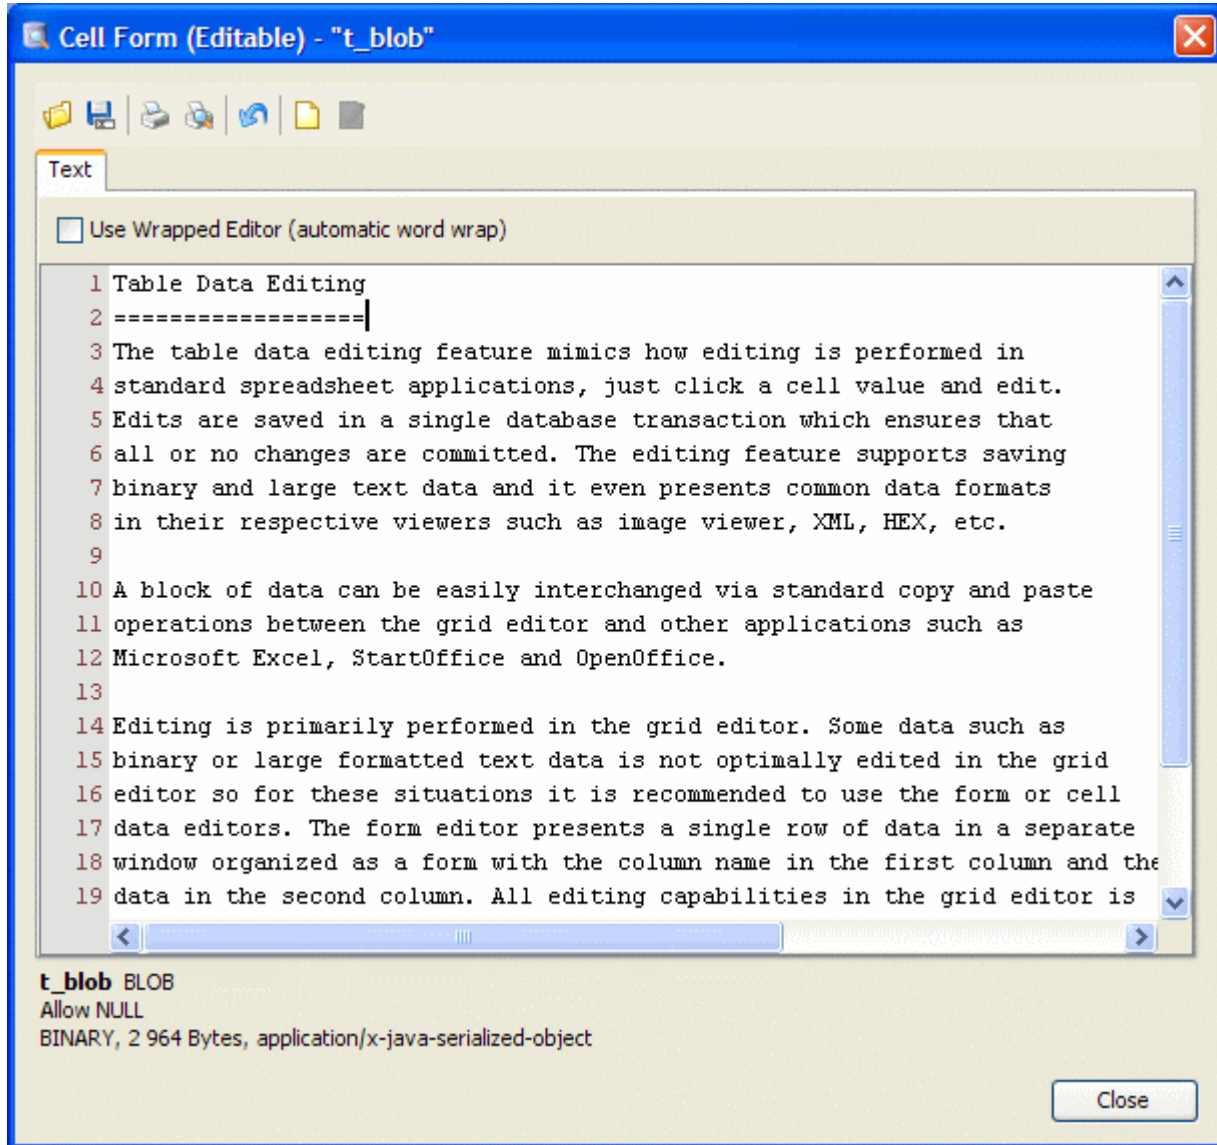


Figure: The text viewer and editor

Import from File

Importing data from a file can be done in the form and cell editors. Imported binary data with a recognized binary viewer will be displayed accordingly. Import is supported for binary and text data.

Export to File

Export can be made in the grid, form and cell editors for binary and text data.

Table Data Navigation

Introduction

A powerful way to study database data is to navigate between the tables in a schema by following table relationships declared by primary and foreign keys. DbVisualizer includes a Navigator feature for this purpose, visualizing the relationships graphically while making the data for each navigation case easily accessible in a data grid.

To launch the Navigator, select the table you want to start the navigation from in the Database Objects Tree, and then open the Navigator tab in the Object View.

Table: DEPARTMENTS Actions ▾

Derby on Mac | Schemas | HR | Tables | DEPARTMENTS

Info | Columns | Data | Row Count
 Primary Key | Indexes | Row Id | References | Navigator | DDL

Powered by yFiles

Related Table:

EMPLOYEE_ID	FIRST_NAME	LAST_NAME	EMAIL	PHONE_NUMBER	HIRE_DATE	JOB_ID	SALARY	COMMISSION
103	Alexander	Hunold	AHUNOLD	590.423.4567	1990-01-03	IT_PROG	9000.00	(null)
104	Bruce	Ernst	BERNST	590.423.4568	1991-05-21	IT_PROG	6000.00	(null)
105	David	Austin	DAUSTIN	590.423.4569	1997-06-25	IT_PROG	4800.00	(null)
106	Valli	Pataballa	VPATABAL	590.423.4560	1998-02-05	IT_PROG	4800.00	(null)
107	Diana	Lorentz	DLORENTZ	590.423.5567	1999-02-07	IT_PROG	4200.00	(null)

Max Rows: Max Chars: 0.002/0.002 sec | 5/11 | 1-5

Figure: The Navigator tab showing the initial table

The Navigator has two parts: a graphical view and a data grid. Initially, the graphical view shows just the selected start table, and the data grid shows the data for the start table.

The data grid is a read-only grid of the same type as you encounter in other parts of DbVisualizer, but extended with a Related Table list and a Tag button. You can learn more about the general data grid in the [Data Grid](#) section of the [Getting Started](#) chapter. The Navigator specific extensions are described in detail in the following section.

Data Navigation

Data navigation in DbVisualizer means following table relationships declared by primary and foreign keys, using a unique key value. In the example schema shown in the screenshots in this section, there is a table named DEPARTMENTS with a primary key named DEPARTMENT_ID. Another table named EMPLOYEES has a foreign key constraint, declaring that values in its DEPARTMENT_ID column refer to primary key values in the column with the same name in the DEPARTMENTS table.

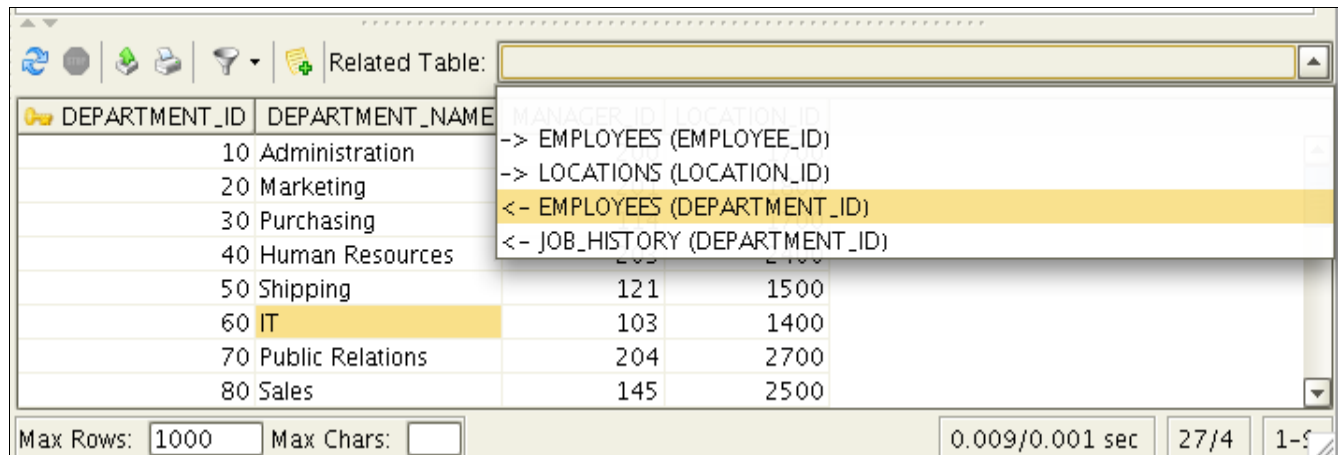


Figure: The Related Table list

If you use DEPARTMENTS as you start table, you can easily navigate to the EMPLOYEES table for different DEPARTMENT_ID values. In the data grid, select one or more columns in the row that holds the DEPARTMENT_ID you want to use for navigation. In the figure above, the DEPARTMENT_NAME column in the row for DEPARTMENT_ID = 60 is selected.

Next, bring up the Related Table list. It lists all tables the DEPARTMENTS table is related to through primary and foreign keys, with the key columns within parenthesis. A forward arrow (->) in front of the table name means that the DEPARTMENTS table has a foreign key relation to the named table. A backward arrow (<-) means that the named table has a foreign key relation to the DEPARTMENTS table.

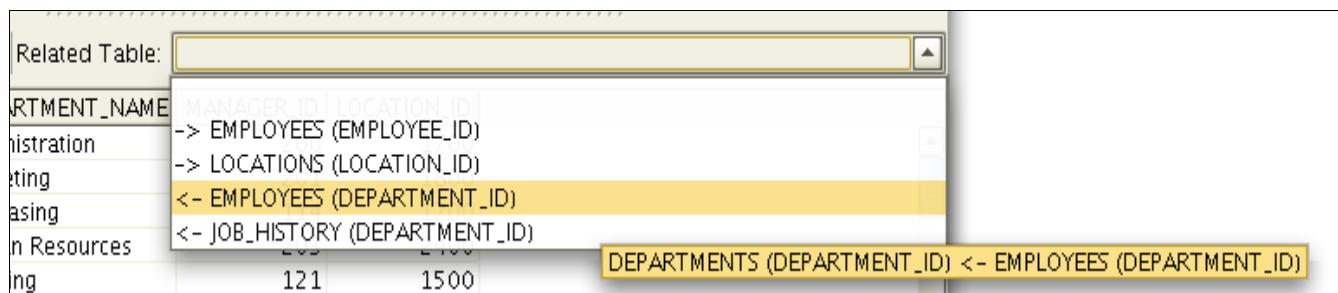


Figure: Tooltip for a Related Table list entry

The Related Table list shows only the table name and columns of the related table, because there is not room for more when a key contains many columns with long names. Sometimes this information is not enough to understand what the relation really means. To make it easier to figure out, you can let the mouse hover over a list entry. A tooltip then shows you the other end of the relation as well, e.g., in the figure above, the tooltip shows that "<- EMPLOYEES (DEPARTMENT_ID)" represents a foreign key from the EMPLOYEES DEPARTMENT_ID column to the DEPARTMENTS DEPARTMENT_ID column.

Table: DEPARTMENTS Actions ▾

Derby on Mac | Schemas | HR | Tables | DEPARTMENTS

Info Columns Data Row Count
 Primary Key Indexes Row Id References Navigator DDL

HR.DEPARTMENTS → DEPARTMENT_ID DEPARTMENT_NAME IT → HR.EMPLOYEES DEPARTMENT_ID 60

Powered by yFiles

Related Table:

EMPLOYEE_ID	FIRST_NAME	LAST_NAME	EMAIL	PHONE_NUMBER	HIRE_DATE	JOB_ID	SALARY	COMM
103	Alexander	Hunold	AHUNOLD	590.423.4567	1990-01-03	IT_PROG	9000.00	(null)
104	Bruce	Ernst	BERNST	590.423.4568	1991-05-21	IT_PROG	6000.00	(null)
105	David	Austin	DAUSTIN	590.423.4569	1997-06-25	IT_PROG	4800.00	(null)
106	Valli	Pataballa	VPATABAL	590.423.4560	1998-02-05	IT_PROG	4800.00	(null)
107	Diana	Lorentz	DLORENTZ	590.423.5567	1999-02-07	IT_PROG	4200.00	(null)

Max Rows: Max Chars: 0.003/0.002 sec 5/11 1-5

Figure: Navigation from DEPARTMENTS to EMPLOYEES for DEPARTMENT_ID = 60

When you select "<- EMPLOYEES (DEPARTMENT_ID)" in the Related Table list, a node is added to the graph for the EMPLOYEES table, with an arrow from the DEPARTMENTS table node to show the navigation direction. We call this a *navigation case*.

The EMPLOYEES node contains the key columns (just one in this example) and their values.

The arrow between the nodes is labeled with the key column name. In addition, the arrow label also shows the name and value of the column that you selected in the DEPARTMENTS table when you created this navigation case, i.e., the DEPARTMENT_NAME column. If you select multiple columns when you create a navigation case, all non-key column names and values are included in the arrow label. This can make it easier to see at a glance what a navigation case represents.

The grid is also updated when you create a navigation case, to show all rows in the table you navigated to that

has a key value corresponding to the selected key value in the table you navigated from. In this case, it shows all rows in the EMPLOYEES table with DEPARTMENT_ID equal to 60.

You can continue to create more navigation cases from any node in the graph. For instance, if the schema contains a table with job history for employees, you can navigate to the history for an employee from the EMPLOYEES node. Or, you can select the DEPARTMENTS node in the graph to navigate to the EMPLOYEES table for a different department. Just click on the DEPARTMENTS node, select another row in the data grid and then the same Related Table list entry.

The screenshot shows a database navigation interface for the 'DEPARTMENTS' table. The main area displays a graph with 'HR.DEPARTMENTS' on the left. Two arrows point to 'HR.EMPLOYEES' nodes on the right. The top arrow is labeled 'DEPARTMENT_ID Human Resources' and points to a node with 'DEPARTMENT_ID 40'. The bottom arrow is labeled 'DEPARTMENT_ID IT' and points to a node with 'DEPARTMENT_ID 60'. Below the graph is a 'Related Table:' dropdown menu and a data grid. The data grid shows the following records:

EMPLOYEE_ID	FIRST_NAME	LAST_NAME	EMAIL	PHONE_NUMBER	HIRE_DATE	JOB_ID	SALARY	COMM
103	Alexander	Hunold	AHUNOLD	590.423.4567	1990-01-03	IT_PROG	9000.00	(null)
104	Bruce	Ernst	BERNST	590.423.4568	1991-05-21	IT_PROG	6000.00	(null)
105	David	Austin	DAUSTIN	590.423.4569	1997-06-25	IT_PROG	4800.00	(null)
106	Valli	Pataballa	VPATABAL	590.423.4560	1998-02-05	IT_PROG	4800.00	(null)
107	Diana	Lorentz	DLORENTZ	590.423.5567	1999-02-07	IT_PROG	4200.00	(null)

At the bottom of the interface, there are controls for 'Max Rows: 1000', 'Max Chars: []', and performance metrics '0.002/0.002 sec', '5/11', and '1-5'.

Figure: Two navigation cases

Every time you select a node in the graph, the data grid is updated to show the corresponding data. The grid settings for one node are independent of the settings for another node. For instance, if you define a filter for one node, the filter is only associated with the grid for that node.

Adding Context Information to the Graph

The navigation node always show the key columns and their values, but sometimes you may want to add other columns to the node to better describe what it represents. This is called *tagging* the node. There are two ways to do so: drag and drop cells from the grid to any node, or use the Tag button in the grid toolbar to tag the currently selected node with the currently selected cells in the grid.

To drag and drop cells to a node, select one or more cells in the grid. With the left mouse button pressed and the mouse positioned over one of the selected cells, drag the cells over a node in the graph and release the mouse button. The cells are added to the node.

Table: DEPARTMENTS

Derby on Mac | Schemas | HR | Tables | DEPARTMENTS

Info Columns Data Row Count
Primary Key Indexes Row Id References Navigator DDL

HR.DEPARTMENTS

DEPARTMENT_ID
DEPARTMENT_NAME Human Resources

HR.EMPLOYEES

DEPARTMENT_ID 40
FIRST_NAME Susan
LAST_NAME Mavis

DEPARTMENT_ID
DEPARTMENT_NAME IT

HR.EMPLOYEES

DEPARTMENT_ID 60

Powered by yFiles

Related Table:

EMPLOYEE_ID	FIRST_NAME	LAST_NAME	EMAIL	PHONE_NUMBER	HIRE_DATE	JOB_ID	SALARY	COMMISSION_PCT
203	Susan	Mavis	SMAVRIS	515.123.7777	1994-06-07	HR_REP	6500.00	(null)

Max Rows: 1000 Max Chars: 0.032/0.002 sec 1/11 1-

Figure: A node tagged with additional column values

Alternatively, you can select the cells in the grid and click on the Tag button (📄➕) to add the cell values to the currently selected node.

Arranging the Graph

As you add navigation cases, you may find that you need to move nodes around, remove selected nodes, zoom and move around in the graph, etc.

You can rearrange the layout of the graph by selecting a node and, with the left mouse button pressed, drag it around. The arrow and its label moves with the node.

The toolbar for the graph offers a number of tools to help you with other tasks.












Figure: The graph toolbar

All these tasks can also be accessed through the graph popup menu.

-  Clicking the Reload button removes all navigation cases, leaving just the node for the table you started with.
-  You use the Show/Hide Controls button to control the display of an Overview control, see below.
-  The Zoom In button lets you zoom into the graph, one step per click.
-  The Zoom Out button zooms the graph out one step with each click.
-  Clicking the Zoom 100% button zooms the graph so that all items are shown with their standard size.
-  Use the Fit button to make all graph items fit in the graph display area.
-  The Relayout button lays out all graph item with standard positions, distances between items, etc. This can be useful after making manual changes, such as removing nodes or tagging nodes.
-  The Remove Node button removes the selected node. It is only enabled when a navigation case node is selected.

The Overview control is useful for large graphs that do not fit into the display area.

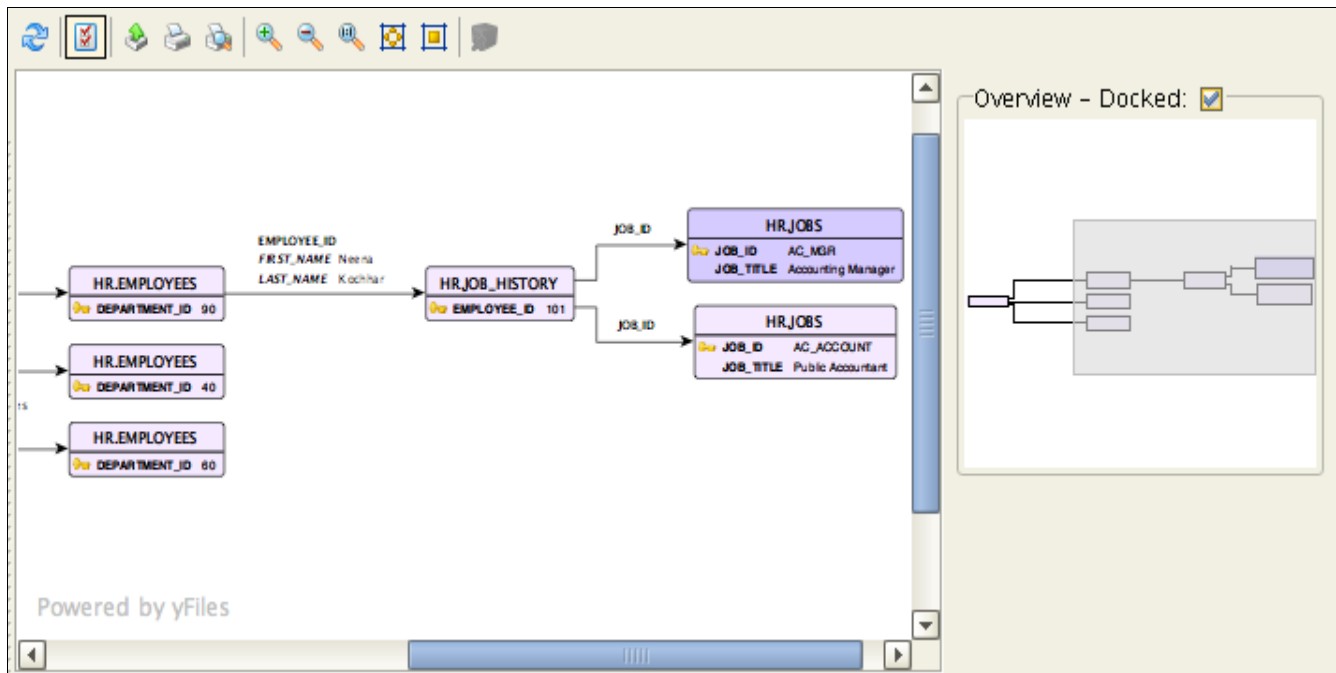





Figure: Graph with the Overview Window displayed

The gray area in the Overview control indicates the portion of the graph that is currently shown in the display area. You can drag the gray area around to study other portions of the graph.

To get a larger graph display area, you can put the Overview control in a separate window. Just uncheck the Docked checkbox.

Exporting and Printing the Graph

You can also export the graph to an image file or print it. Use the corresponding toolbar buttons to do this.

-  Export the graph to a JPG or GIF file.
-  Print the graph
-  Show a preview of how the graph will be printed

When you print the graph, you are prompted for information about what to print (the Graph or the View, i.e., just the portion visible in the display area) and how many rows and columns to split the printing over (one page is used for each row/column). See the Export and Print sections in the [Export, Import and Print](#) section for details.

Procedure Editor

Introduction

Many databases offers the capability to store programs in the database. These programs may be of two types: functions and procedures. The difference is that a function returns a value while a procedure does not. In addition, some databases offer the package concept, which means that a collection of functions and/or procedures are grouped together in one unit. A package is the interface describing the functions and procedures, while the package body contains the implementation. You can use DbVisualizer actions to create and drop procedural object of these types, and use the procedure editor to browse, edit and compile these object types.

The examples throughout this document refer to the procedure object type, but all described features can also be applied to function, package and package body objects.

Create Procedure

To create a new procedure, simply select the Procedures node in the objects tree and choose Create Procedure from its action menu.

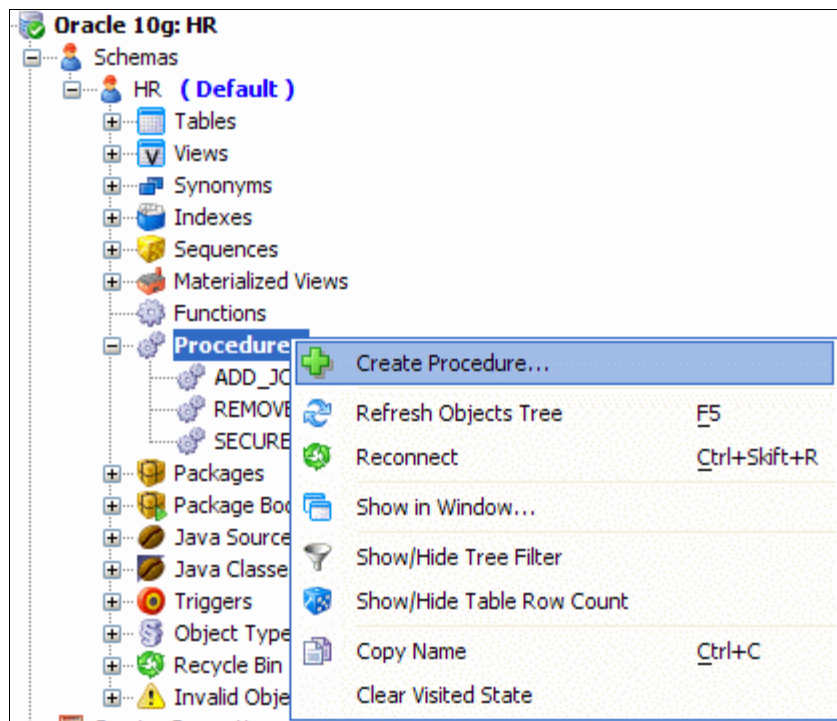


Figure: The actions menu for the Procedures node

Next, a dialog is displayed in which you enter the procedure name and the parameters for the new procedure. This data forms the interface for the procedure.

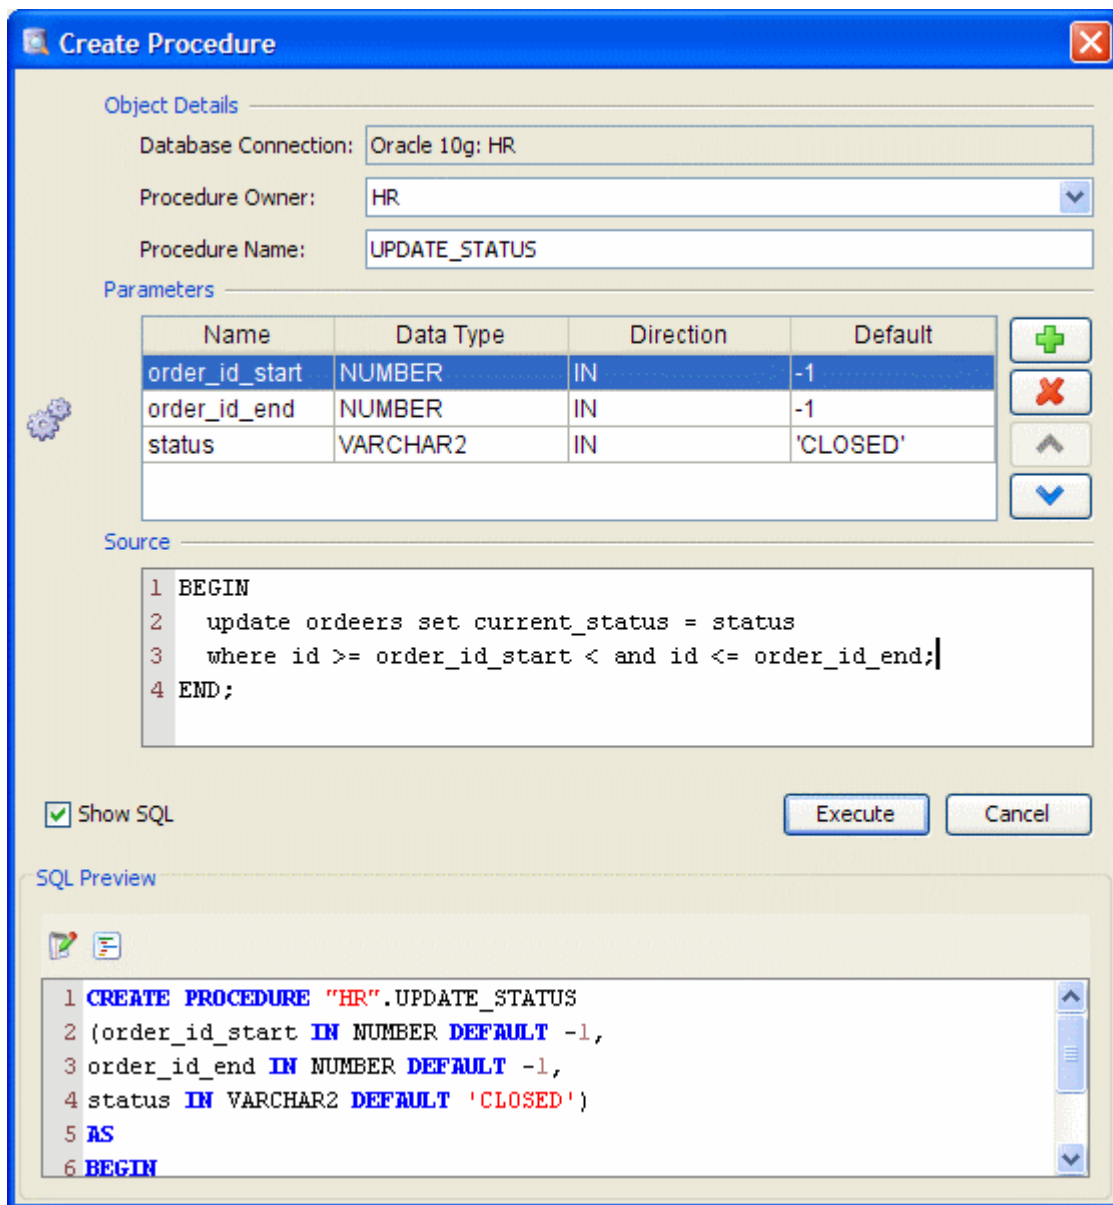


Figure: The create procedure dialog

Use the buttons to the right of the parameter list to insert, remove and move the parameters. For every parameter, you must supply its Name. Leaving out the Data Type results in the VARCHAR data type. The Direction is by default set to IN.

You can also enter the implementation source code in the Source area, but be aware that it is saved as-is for some databases; if there are any errors, you may not be notified when you save the procedure. You need to verify the status and fix errors, if any, as described in the next section. As an alternative, you may therefore like to leave just the default sample implementation in the Create Procedure dialog and enter the real the implementation later.

Click Execute in the dialog to create the new procedure.

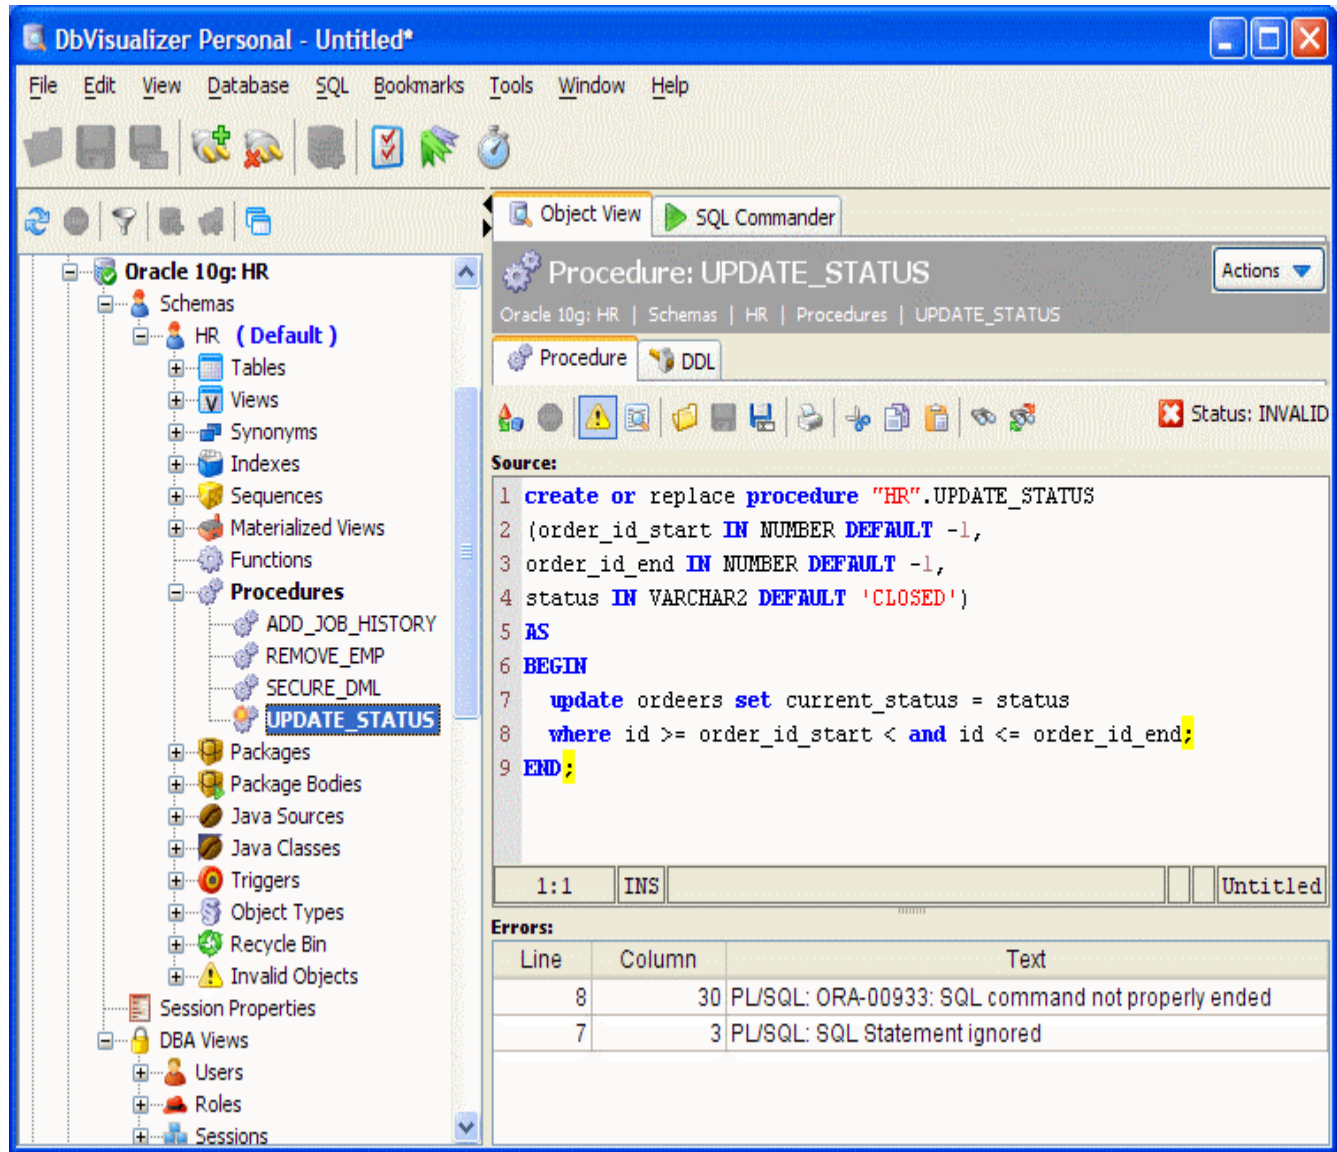


Figure: The newly created procedure

Selecting the newly created procedure in the tree will show the source for it in the procedure editor.

Edit and Compile

Note: This editor is currently supported only for Oracle databases! For other databases, you need to use the SQL Commander and commands like ALTER PROCEDURE to edit and compile the procedure. See your database documentation for details.

The editor has a toolbar with various actions to compile the procedure, save and load the source to/from file, show parameter information and common editing operations. The Status indicator shows whether the procedure is valid or invalid based on last compilation.

Edit the source code and compile the procedure when you are happy with the code, using the Compile toolbar button.

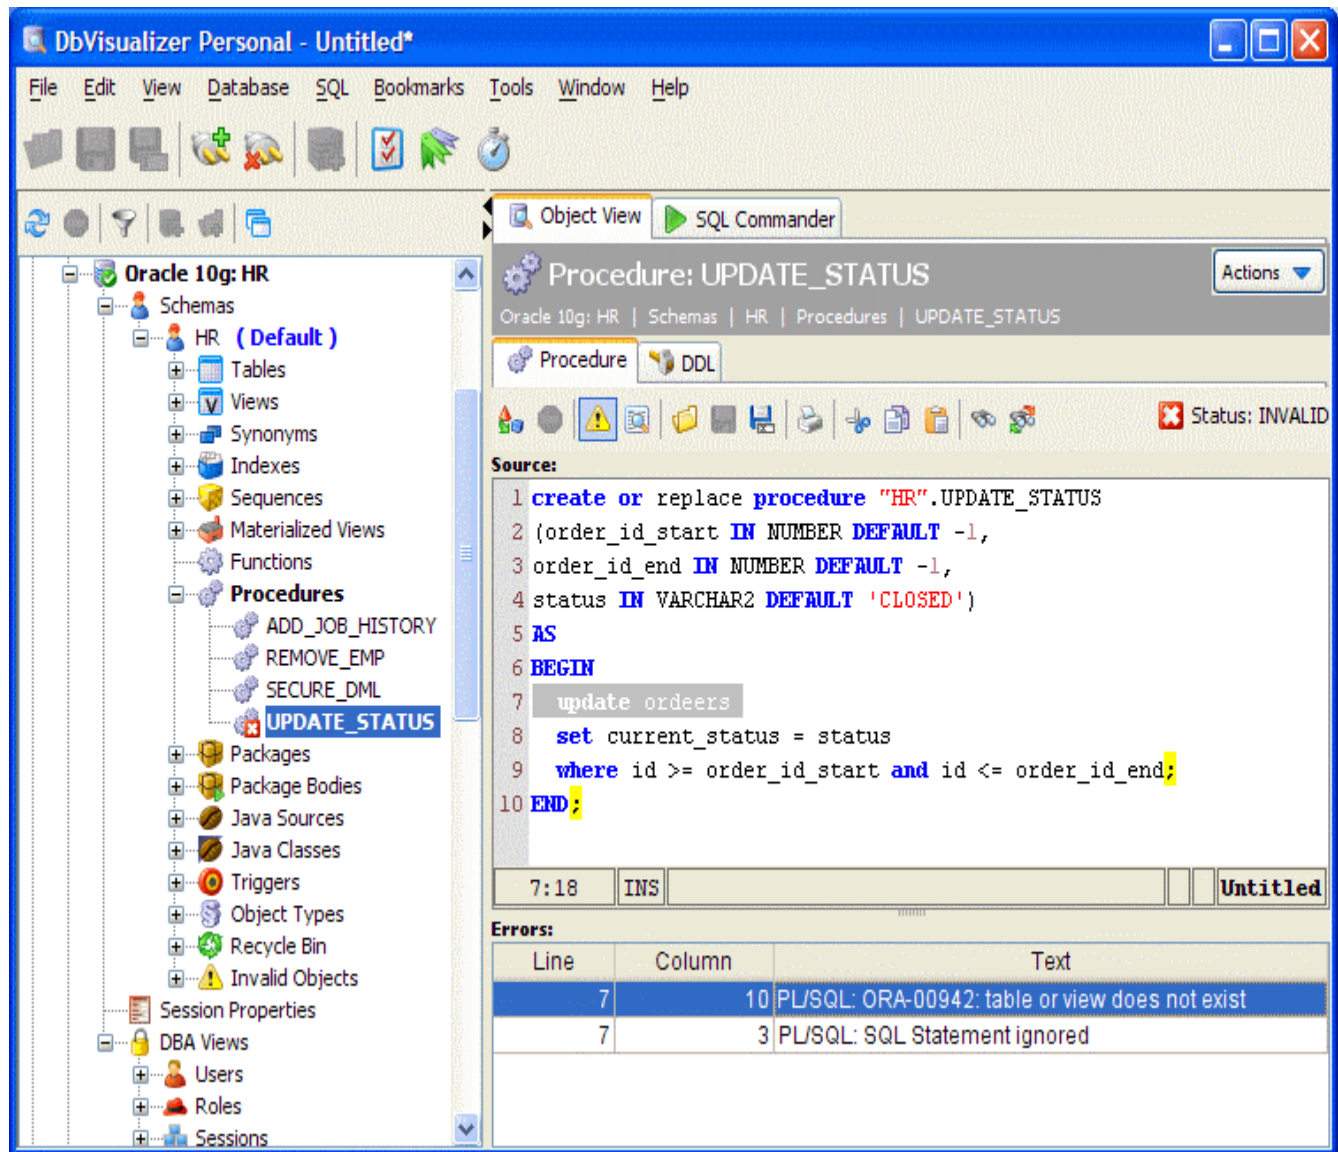


Figure: Compiling procedure with errors

If error(s) occur during compilation, the error list appears below the editor. It show the row number for the error in the source editor and an error message. When you click the error in the list, the corresponding row is highlighted in the editor.

In addition to the Status indicator in the editor, the object icon in the tree shows a little red cross for invalid procedures. The figure below shows the result of correcting the errors and recompiling the procedure:

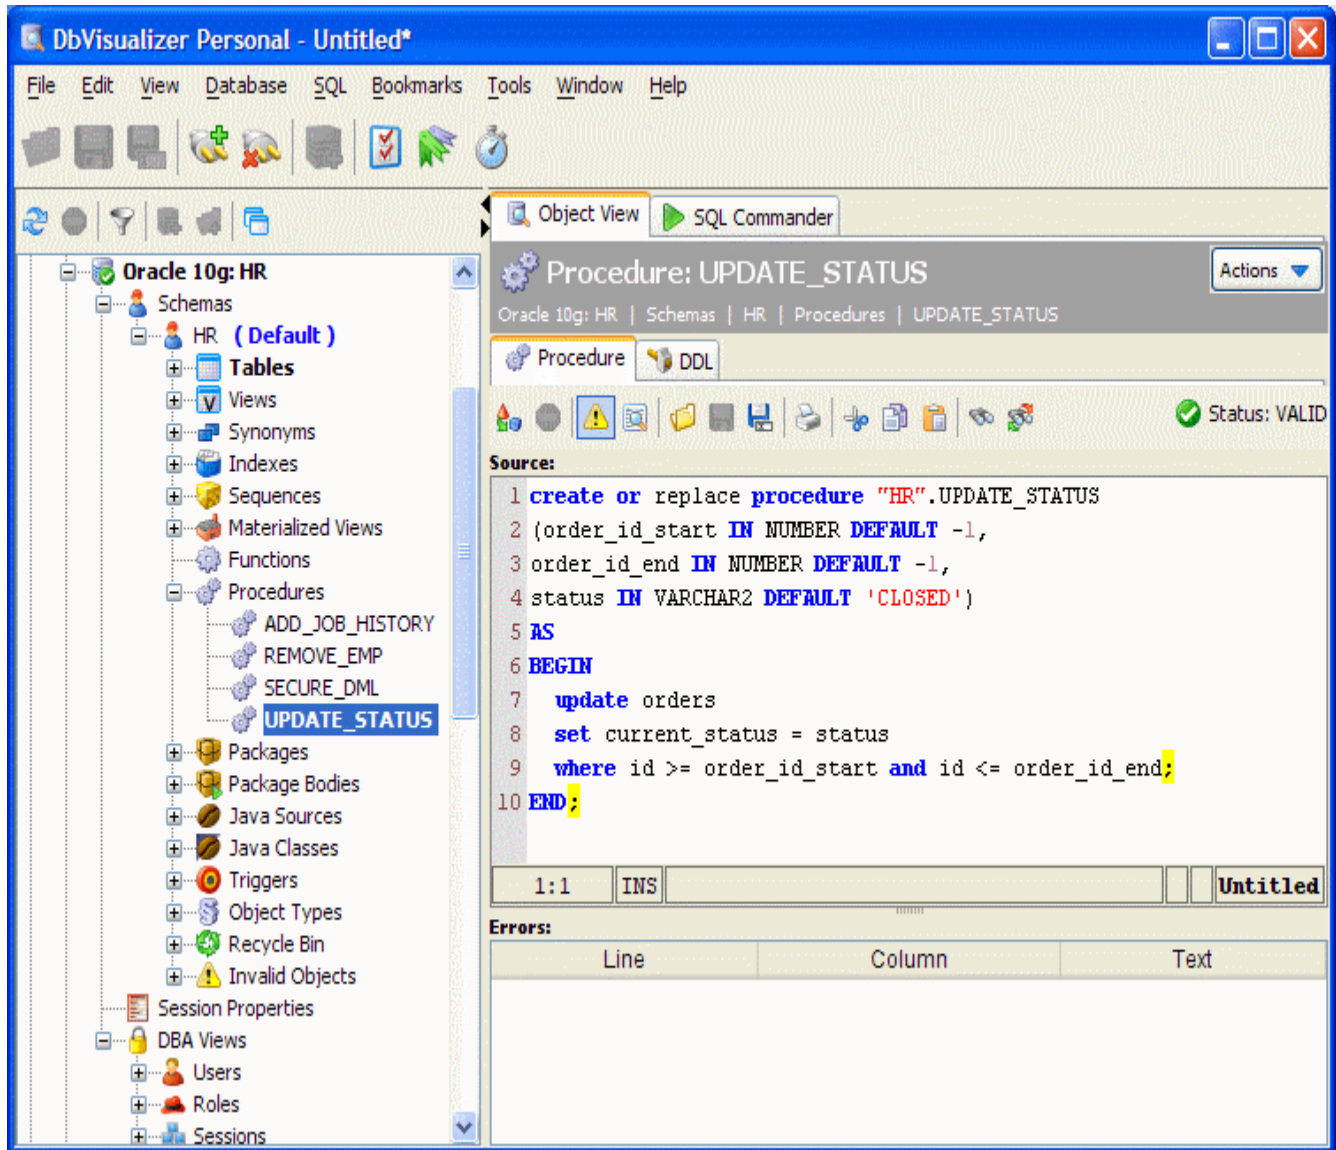


Figure: Compiling procedure with successful result

The status indicator now shows that the procedure is VALID.

Running in SQL Commander

You can now test the procedure in the SQL Commander as shown in the next screenshot.

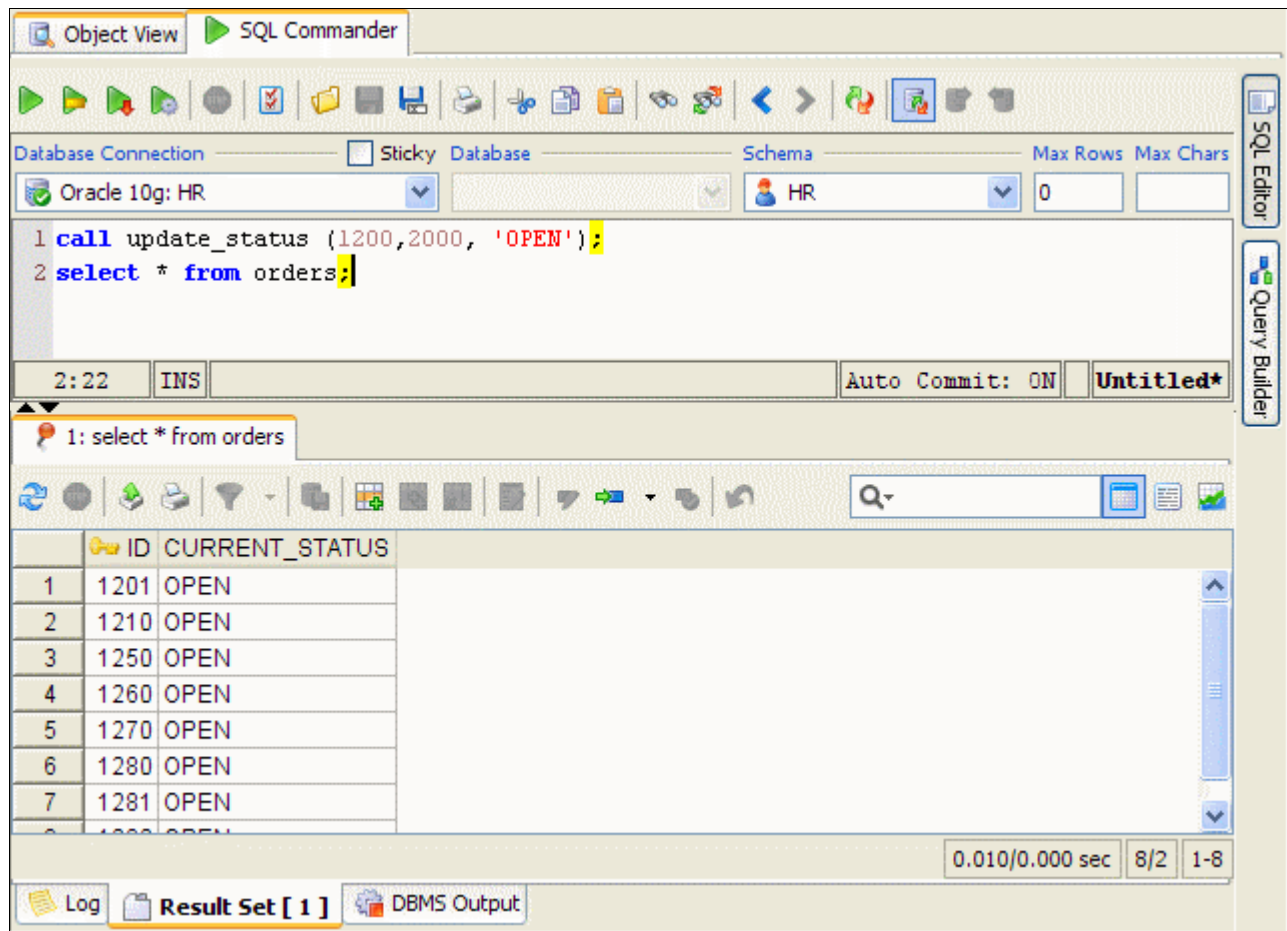


Figure: Running the procedure in SQL Commander

The figure shows the invocation of the update_status procedure with parameter values meaning that all IDs in the ORDERS table between 1200 and 2000 should be set to "OPEN". The second statement selects from the updated table.

SQL Bookmarks

Introduction

The purpose with the bookmark management is to offer a way to save SQL statements between invocations of DbVisualizer and make it easy to execute them. Another important requirement is to organize SQL statements in folders for structural and grouping purposes. The core of the bookmark management is the Bookmark Editor. It is here the bookmarks are organized.

The bookmark editor depends heavily on the SQL Commander, since, when requesting to execute an SQL Bookmark, the Bookmark Editor passes the SQL statement along with the connection data to the SQL Commander. You then use the SQL commander to edit and test the SQL until it is complete.

What's a bookmark in DbVisualizer?

An SQL Bookmark is generally an SQL statement that is saved between invocations of DbVisualizer. In addition, it also keeps related information needed to execute the SQL and present the result accordingly once it is requested.

- SQL statement
- Bookmark name
- Database Connection
- Catalog (aka Database)
- Schema
- Chart settings (optional)

There are different types of bookmarks. DbVisualizer automatically creates bookmarks in the following function areas:

- Each SQL that is executed in the SQL Commander is saved as an SQL bookmark in the History folder
- Each monitored SQL statement in the Monitor feature is an SQL bookmark. DbVisualizer creates SQL Bookmarks in the New folder when you use the operations in the Data tab to create row count related monitors

You can also create SQL Bookmarks manually in the Bookmark Editor, as well as from SQL in the SQL Commander with the operations available in the Bookmarks menu. The following sections describes this in more detail.



The Bookmarks Main Menu

The bookmarks main menu in the DbVisualizer window contains the following choices:



Figure: The Bookmark main menu

All menu choices except the Bookmark Editor choice are disabled if you are not in the SQL Commander tab.

Menu Choice	Description
Bookmark Editor	Requests to display the Bookmark Editor.
Add Bookmark to Folder	<p>This choice has a sub menu in which all folders are displayed. This list displays the paths for all folders (i.e the folder hierarchy from the root). The root folders are Personal, New or History (read more about these in the sections below). Once a folder has been selected the following dialog is displayed. Here you can change the default name and add an optional note.</p> 
Replace Bookmark	<p>This option is used to replace the chosen SQL bookmark with the SQL and connection data that is in the current SQL Commander editor. The Replace Bookmark sub menu consists of the root folders and, as the last menu entry, the name of the last SQL bookmark that was passed from the Bookmarks Editor, if any. If you want to replace the data for that SQL Bookmark just select its name in the menu.</p> 
Get Bookmark	Get Bookmark shows the same menu hierarchy as Replace Bookmark, but works in the reverse direction; it fetches the chosen SQL Bookmark and inserts it into the current SQL Commander editor.
Execute Bookmark	Same as Get Bookmark but this one also executes the SQL statement(s)

Bookmark Editor

The Bookmark Editor is at the core of the bookmark management, used for organization of SQL Bookmarks in folders and to do various adjustments.

Bookmark list

The editor presents the bookmarks in a tree list with the same structure as the tree that appears in the Bookmarks main menu options. The tree has three root folders that cannot be changed, moved or removed. There is basically no difference between these root folders except that they are used in different contexts in DbVisualizer.

- Personal
This root folder is supposed to hold the structure of favorite SQL bookmarks. By putting SQL bookmarks in folders, you get a better organization and overview of your bookmarks. All nodes in this root folder are manually maintained.
- New
When you create Row Count Monitors in the Database Objects->Data tab, DbVisualizer adds these monitors (as SQL Bookmarks) in the New root folder.

- History
All SQL statements or scripts that you execute in the SQL Commander are automatically added in the History root folder. The latest executed statement appear first in the list.

(The number after the root folder names indicates the number of SQL bookmarks the root folder contains).

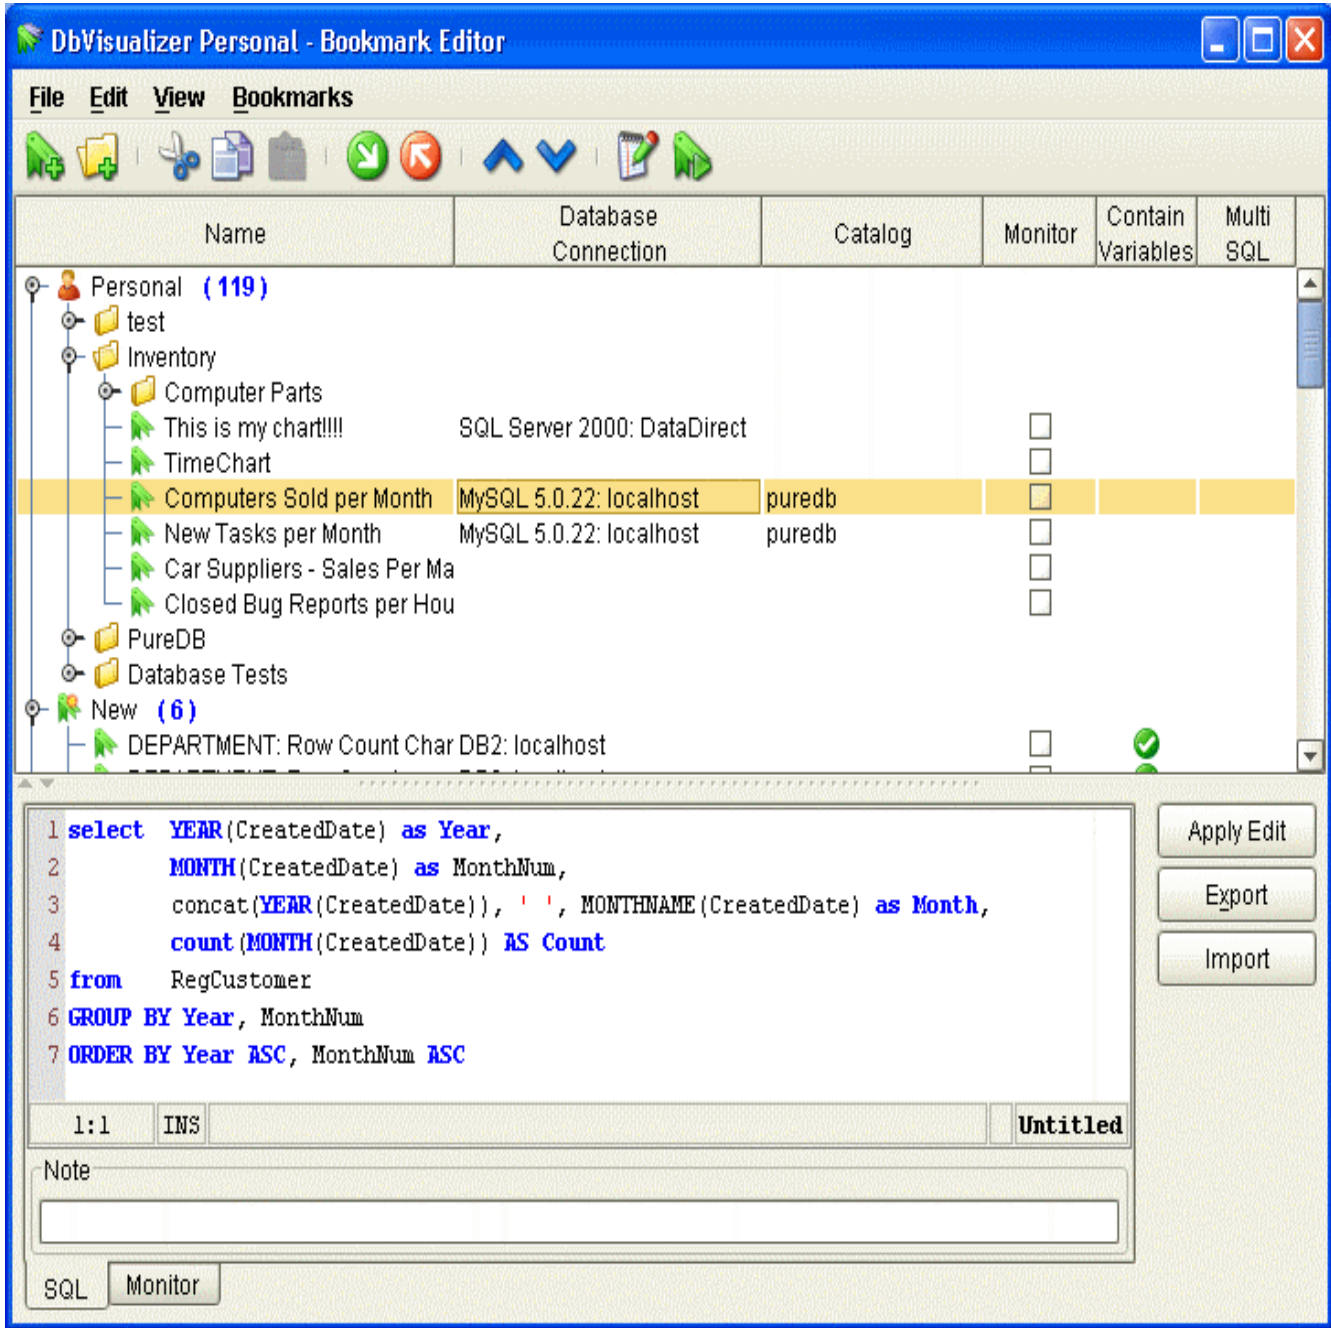


Figure: The Bookmark Editor

You cannot create folders or SQL bookmarks in the New or History root folders. The way to work with these folders is to copy the SQL Bookmarks you want from them into the appropriate location in the Personal root folder.

For each SQL Bookmark, the tree holds the following information:

Column in list	Description
Name	The name of the node (folder or SQL bookmark). Modify the name by selecting the column and click once to get into editor mode. The Edit->Change Name menu choice can be used for the same purpose. When an SQL Bookmark is created by DbVisualizer, the default name is the first 40 characters of the SQL statement.
Database Connection	Double-clicking on the Database Connection column displays a list of all defined database connections. The list indicates whether a connection is established or closed. You can choose a connection in the list if you want to use another database connection for an SQL Bookmark.
Catalog	This column lists the catalog (aka database) that was selected when the bookmark was created. You can change the catalog by double-clicking on the name and pick another from a list of accessible catalogs. Note: The list of catalogs is empty when the Database Connection is not established.
Schema	The schema that was selected when the bookmark was created. You can change it in the same way as you change the catalog.
Monitor	Check this box to enable the SQL Bookmark to be used as a monitor and thereby appear in the Monitor main tab. SQL statements that returns results are the most obvious candidates for being monitored.
Contain Variables	This column is read only and indicates whether the SQL statement includes any DbVisualizer variables. (I.e <code>\$\$variable name\$\$</code>)
Multi SQL	This column is also read only and indicates whether the SQL statement is composed of several SQL statements (aka script). This is determined by looking for statement delimiters in the SQL.

New and History root folders

The number of SQL bookmarks that may be added by DbVisualizer to the New and History root folders are specified in the Tool Properties->Bookmarks category. Bookmarks in these folders can be removed one by one or each folder can be cleared using the File->Clear all New entries or File->Clear all History entries.

SQL Editor

Monitor information

The monitor sub tab controls the total number of rows that will be kept in the result grid until rows are automatically removed. This feature is specific to the monitor feature. Please see [Charts and Monitors](#) for more information.

The Note field

You can use the note field to write a short description of the SQL Bookmark. This note appears as a tooltip in the Bookmarks main menu.

Executing an SQL bookmark or folder of SQL bookmarks

The SQL editor in the bookmark editor can be used to modify the SQL but it is not the place to execute SQL statements. Instead, you use pass the bookmark to the SQL Commander execute it there. Select the Edit->Copy to SQL Commander or Edit->Execute Bookmark menu operation to copy the selected SQL bookmark into the SQL Commander. The SQL Commander is then used to execute and edit the SQL. Once you are satisfied with it, you can select the last entry in the Bookmarks menu save the edits:



Figure: The Bookmark->Replace sub menu

If the last entry displays "No current bookmark," it indicates that the currently edited SQL was not passed from the Bookmark Editor. You can use the other menu choices to locate the bookmark to replace or save it as a new bookmark.

The Copy to SQL Commander and Execute Bookmark operations also operate on a folder. In this case, a script is created from all direct child SQL Bookmarks that are located in that folder. Each SQL statement is delimited by the delimiter as specified in Tool Properties.

Name	Database Connection	Catalog	Monitor	Contain Variables	Multi SQL
Personal (44)					
CustDB					
Inventory					
Add users					
Drop table	DB2: db2jcc		<input type="checkbox"/>		
Create table	DB2: db2jcc		<input type="checkbox"/>		
Insert master	DB2: db2jcc		<input type="checkbox"/>		
Insert master	DB2: db2jcc		<input type="checkbox"/>		
Insert master	DB2: db2jcc		<input type="checkbox"/>		
Computer perif.					
Organisation					
Consistency Check					
Database tests					
New (4)					
History (92)					

Figure: Selecting a folder for execution

As an example, choosing the Copy to SQL Commander for a folder of bookmarks may result in the following statements being added in the SQL Commander:




Figure: The SQL Commander editor

Note that the Database Connection and Catalog lists are empty. You need to select these from the lists when a script of SQL bookmarks is passed from the Bookmark Editor.

Tool Properties

Customizing DbVisualizer

DbVisualizer is highly customizable. You can control formatting, layout and the way DbVisualizer interacts with databases. The default settings are good enough for normal use, but sometimes it is necessary to modify these properties. This chapter guides you through all the properties.

The Tool Properties window divides properties into two groups:

- **General Settings**
These settings controls DbVisualizer in general, such as fonts, colors, data formats, etc.
- **Database Settings**
These settings are per supported database type and defines properties that are used in database specific operations. When you set a database property in Tool Properties, it applies to all database connections defined for that database type. To set a property for one specific connection, use the Connection Properties, available in the Object Details area when you select a connection.

The user preferences (XML) file

All properties are saved in an XML file. The exact location of this file is platform dependent. The location on your system is listed in the first, General category, in the Tool Properties window. The XML file contains, in addition to all properties, also the information about drivers, database connections, bookmarks, etc. We recommend that you do not edit this file manually; even though it is quite easy to do so, even a simple typo of an element name may cause problems. It's safer to edit all properties from the DbVisualizer GUI.

DbVisualizer automatically creates a backup copy of the XML file when the application is started. The location of this file is the same as for the standard XML file, but a .bak suffix is appended to the filename. The standard XML file might get broken for various reasons. If you see a warning message that the XML file can not be read when you launch of DbVisualizer, simply copy the backup file to the standard location and restart the application. If you move the XML file from its standard location, or if you remove it, DbVisualizer will automatically create a new one.

Tip: the -up command line argument is used to identify the file name (and path) to an alternate XML file.

General Settings

The General settings tab collects all categories that are used to control the general aspects of DbVisualizer.

Use the buttons at the bottom of the window when you have made some changes: Click Ok to save the changes and close the window, the Apply button to save the changes but keep the window open, and the Cancel button to revert all changes. To reset the properties to the factory defaults, use the Defaults button.

Changes are tracked on a per category basis. If you have made changes and click on another category, you are asked whether the changes should be applied or not. When you click Defaults (for both the General and the Database properties), you can reset either all properties or just the properties for currently selected category.

This is a screenshot of the General category tree.

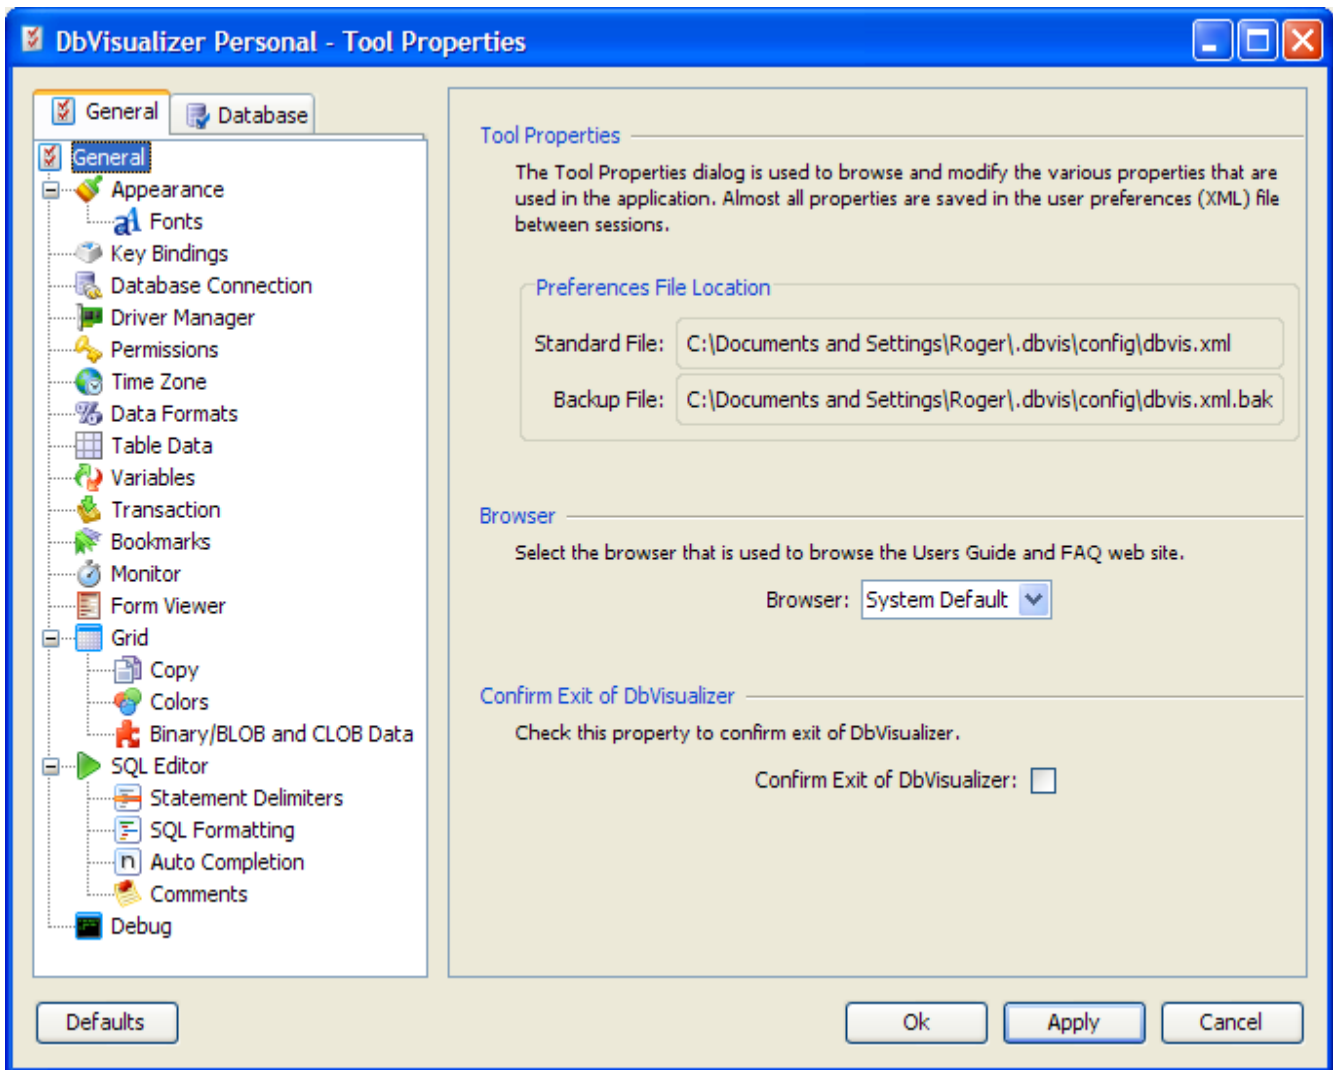


Figure: The Tool Properties window showing the tree with General categories

Appearance

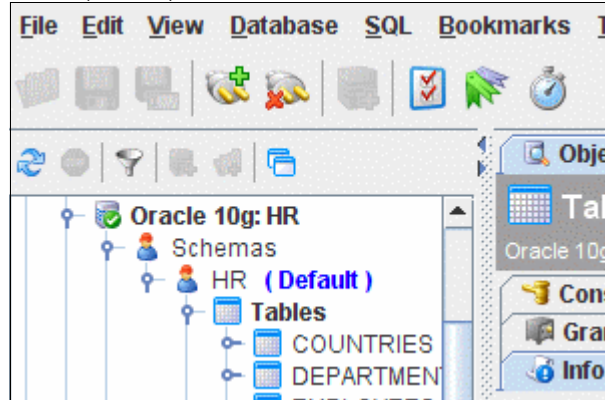
Property	Description
----------	-------------

Controls which look and feel to use.

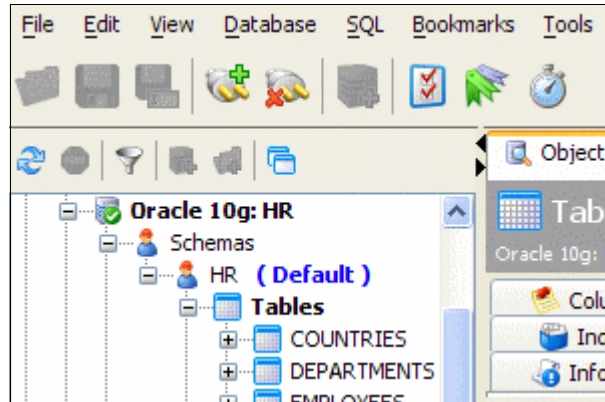
Note 1: You must restart DbVisualizer after you have selected a new look and feel.

Note 2: Some look and feels are platform specific and do not appear on all OS'es

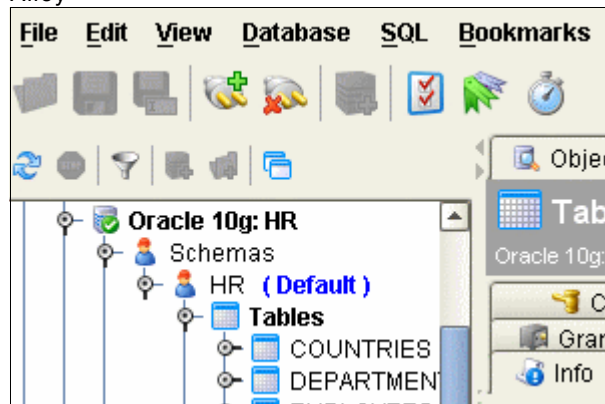
Metal (Ocean)




Windows



Alloy



GTK+



Look and Feel

Icon Sizes	The Menus, Main Tool Bars, Sub Tool Bars settings are used to control the size of the icons.
Show Tab Icons	Specifies whether an icon will appear in the header of all object view tabs.

Fonts

Individual fonts can be defined for SQL Editors, Grids and Text output data. The Application Font settings is used to control the font for all other components in the user interface, such as labels. Increasing the application font size is useful at demos or presentations. Anti-Aliased Fonts is supported by some look and feels and when enabled it gives a much smoother appearance of text in the application. Anti-Aliased font is not supported by the SQL editor.

Key Bindings

You can define key bindings for almost all operations and editor commands in DbVisualizer. Key bindings are grouped in Key Maps. DbVisualizer includes a set of predefined key maps targeted for the supported operating systems. These key maps cannot be deleted or modified. To customize key bindings, copy an existing key map and make your changes.

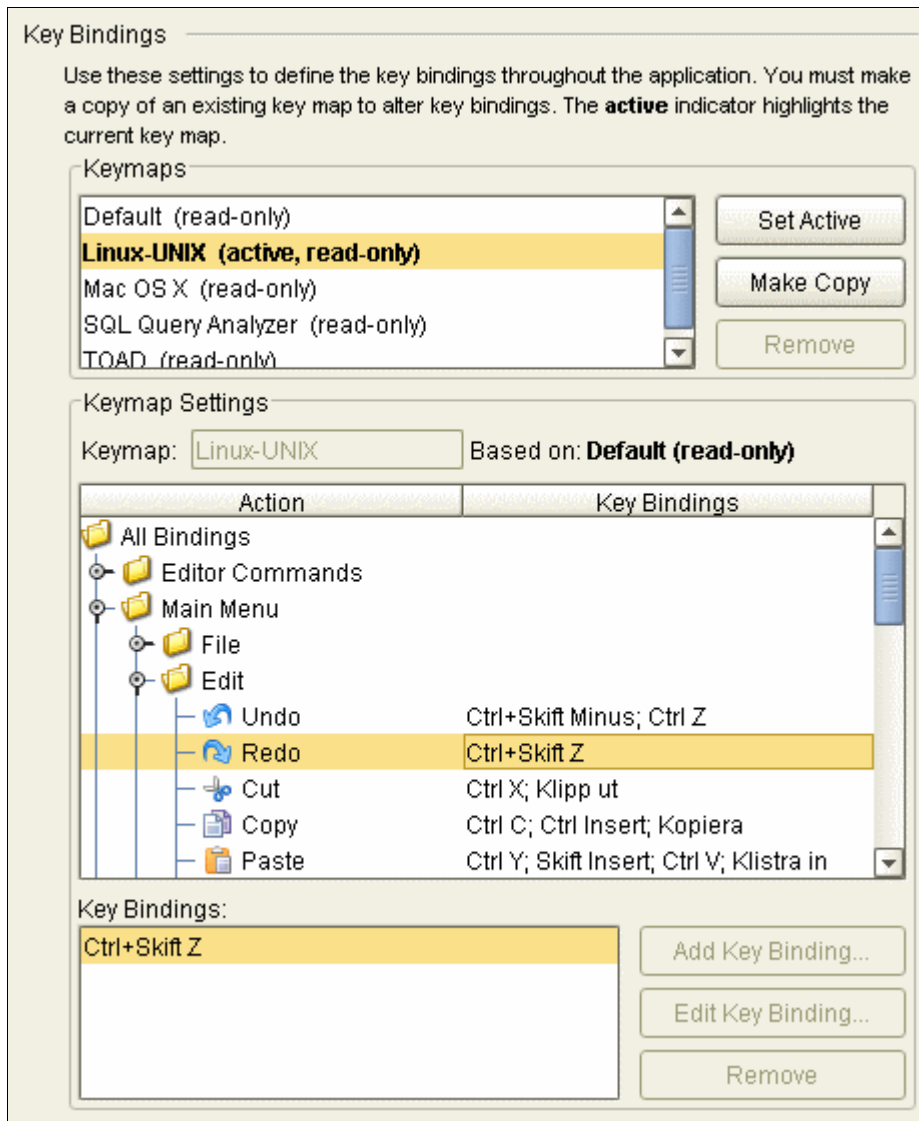


Figure: The key binding editor

All user defined key maps are stored in your `$HOME/.dbvis/config/keymaps` directory. A key map file contain only the differences between the copied key map and the current.

To create a new key map, select the map you want to copy and click the Make Copy button. Set a name on the new key map and activate it with the Set Active button. The newly created key map now has the exact same key bindings as the parent key map.

Note: Key maps must be uniquely named.

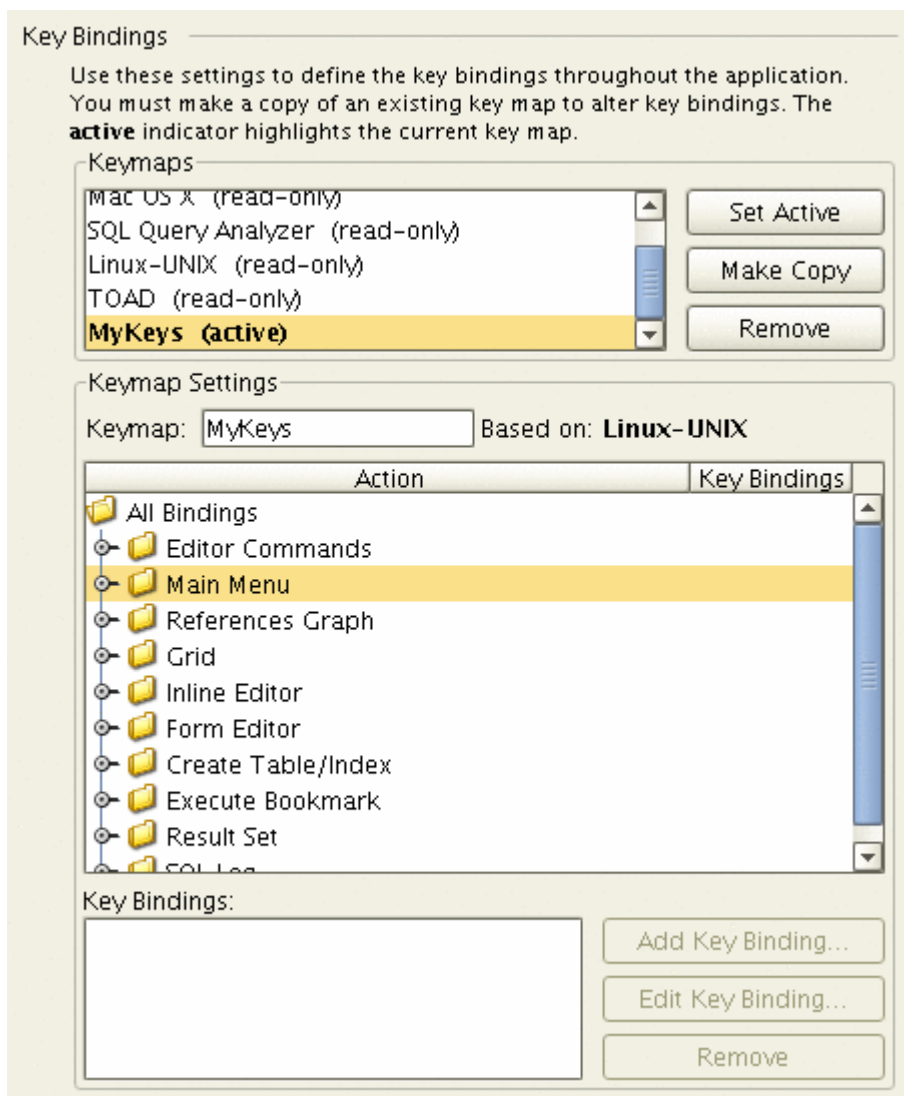


Figure: User defined key map

The action list is organized in folders. The Editor Commands folder lists all actions available in the SQL Commander editor and their current key bindings. The Main Menu folder contain sub folders, each representing a main window menu. The other folders group feature specific actions, such as actions to control the references graph, form editor, etc.

To modify the key bindings for an action, select the action from the action list. The current key bindings are listed in the Key Bindings list.

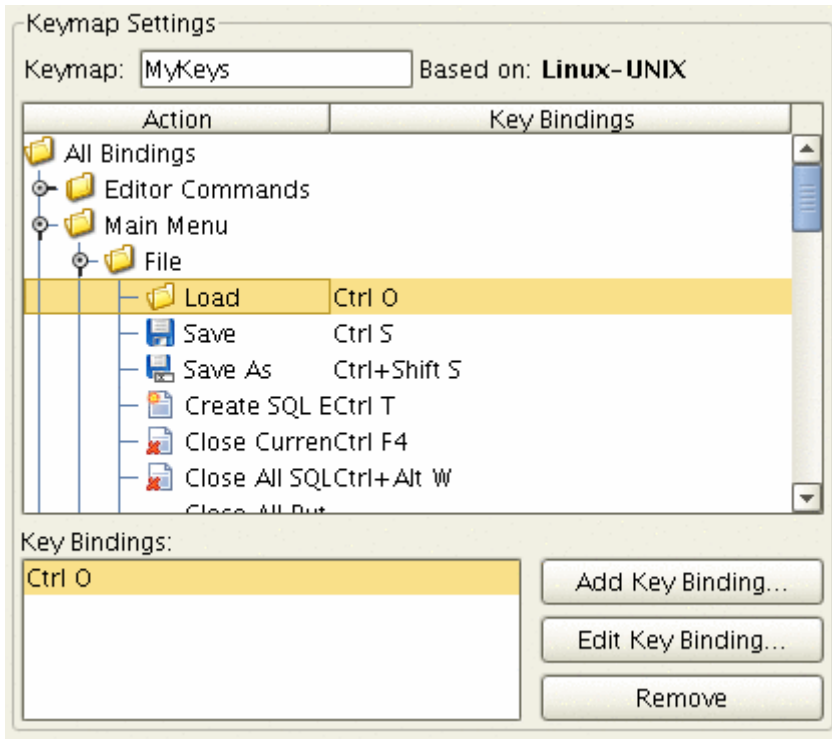


Figure: User defined key map

To add an additional key binding, press Add Key Binding or press Edit Key Binding to edit the selection.

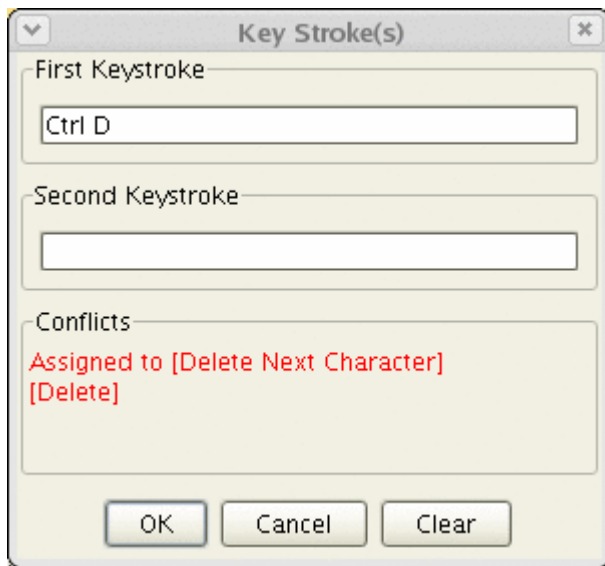


Figure: Key stroke dialog

The key stroke dialog controls whether a key binding is already assigned somewhere else. If there is a conflict with another binding, the Conflicts are shows the names of the actions that are conflicting. The modifier keys Shift, Alt, Ctrl and Command can be used to form the final key binding.

Note: Menu items and tooltips shows the first defined key binding in the list.

Database Connection

Property	Description
Ask When Creating Database Connection	If enabled, you will be asked if you want to use the Database Connection Wizard to create new connections.
Run "Connect All" at Startup	If enabled, the Connect All operation is automatically run when you launch DbVisualizer, connecting all Database Connections marked as being included in the Connect All operation (see the Database properties further down for more on this).
Confirm "Disconnect All"	If enabled, a dialog to be displayed before disconnecting all current database connections when using the Disconnect All operation.

Driver Manager

The Driver Manager searches specified folders for JDBC drivers and helps you make them available for use by DbVisualizer, see the [Load JDBC Driver and Get Connected](#) section for details. In the Driver Manager properties category, you can specify if you want the Driver Manager to run automatically at start-up, when new files are discovered in the specified driver folders, or when driver related errors are encountered. You can also specify the folders to search and files to exclude, if any.

Permissions

The Permission functionality is a security mechanism, where you can specify that certain database operations must be confirmed. You configure permissions per connection mode (Development, Test and Production) for feature areas described in the following sections.

Note: The permission feature is part of DbVisualizer and does not replace the authorization system in the actual database.

SQL Commander Permissions

For the SQL Commander, you can pick the permission type from a drop-down list for each SQL command:

- Allow
This permission type means that you can run the SQL statement without any confirmation
- Deny
This permission type means that the SQL statement is not executed at all.
- Ask
This permission type means that when executing an SQL statement, or a script of statements, the SQL Commander asks you whether the actual SQL command(s) should be executed.

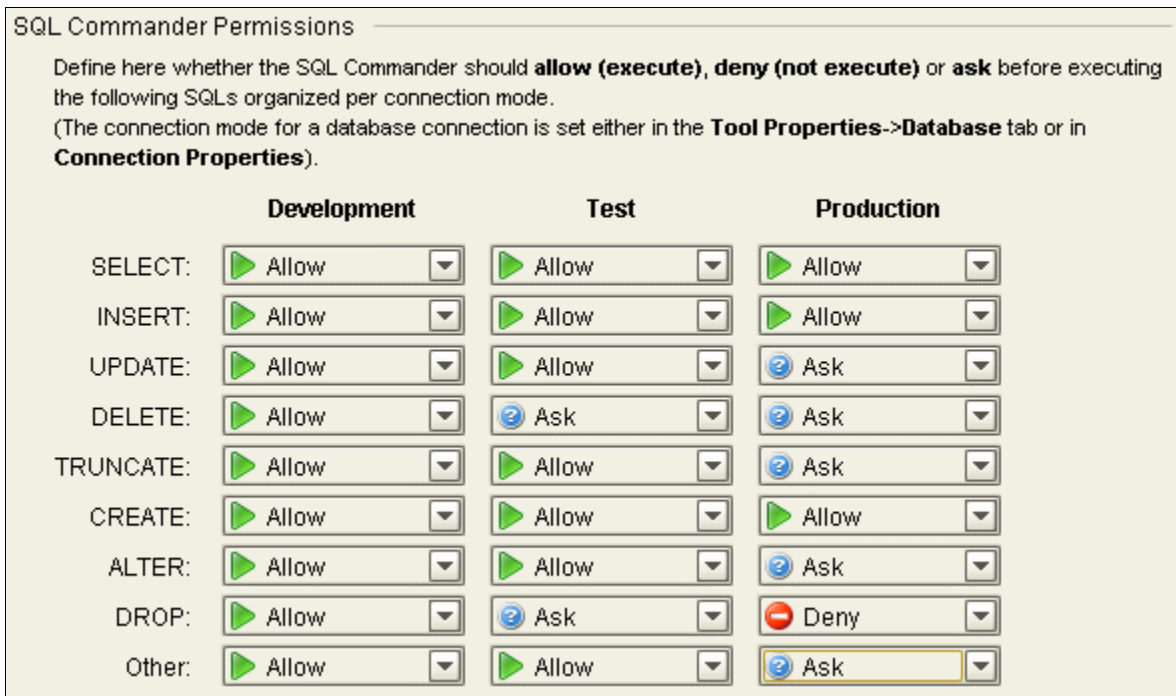


Figure: SQL Commander Permissions

Inline Editor Permissions

The permission types for the inline editor are:

- Confirm
A confirmation window is displayed, and you can accept the operation or cancel it
- No Confirm
The SQL operation is performed without any confirmation

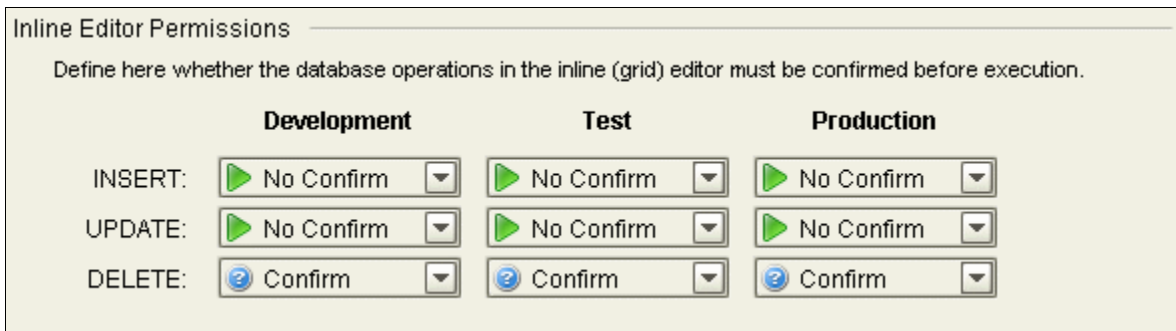


Figure: Inline Editor Permissions

Time Zone

In the Time Zone properties category, you can change the time zone for the DbVisualizer process, and thereby how date and time data is interpreted. DbVisualizer uses the OS time zone by default, which is usually what you want to use.

Changing the time zone is only of interest if you work with a database running with a different time zone than the time zone set on the client where you run DbVisualizer. One example is when working with a database that uses the UTC/GMT time zone to normalize all date/time data.

Data Formats

Property	Description
Date Format	Specifies the date format to use throughout the application (i.e., in grids, forms and during editing). More information below.
Time Format	Specifies the time format to use throughout the application (i.e., in grids, forms and during editing). More information below.
Timestamp Format	Specifies the timestamp format to use throughout the application (i.e., in grids, forms and during editing). More information below.
Numbers Format	Specifies how numbers will be formatted.
Decimal Number Format	Specifies how decimal numbers will be formatted.
Null String	Specifies the string representation of the null value. This string is the readable form of null and appears in grids, forms, exports and during editing.

Date, Time and Timestamp formats

The lists for date, time and timestamp format contain collections of standard formats. If these formats are not suitable, you can enter your own format in the appropriate field. The tokens used to define the format are listed in the right-click menu when the field has focus.

G - Era Designator
y - Year
M - Month in year
w - Week in year
W - Week in month
D - Day in year
d - Day in month
F - Day of week in month
E - Day in week
H - AM/PM marker
H - Hour in day (0-23)
k - Hour in day (1-24)
K - Hour in AM/PM (0-11)
h - Hour in AM/PM (1-12)
m - Minute in hour
s - Second in minute
S - Millisecond
z - Time zone
Z - Time zone

Figure: The date and time right click menu

The complete documentation for these tokens is available at the following web page: [SimpleDateFormat](#).

Number formats

The lists for number and decimal number contain collections of standard formats. If these formats are not suitable, you can enter your own format in the appropriate field. The tokens used to define the format are listed in the right-click menu when the field has focus, and complete documentation for these tokens is available at the following web page: [DecimalFormat](#).

Table Data

Property	Description
Show Table Row Count	Specifies if the number of rows in a table will be displayed in the header of the table in the Database Objects->Data tab. Enabling this property will cause an extra round trip to the database (i.e., a minor performance penalty)
Highlight Primary Key Columns	Specifies if Primary Key columns will be indicated in the Database Objects->Data tab, Variable Substitution dialog, SQL Commander Result grids and in the References Graph.
Include Variables in SQL	Specifies if the right-click menu operations in the Data tab will create SQL statements that include DbVisualizer variables or if the generated statements are plain SQL. Letting DbVisualizer generate statements with variables results in the variable substitution dialog being displayed when these statements are executed in the SQL Commander.
Max Rows at First Display	Set the number of rows that will be fetched for a table in the Data tab when a table is first displayed.

Variables

Variables can be used in the SQL executed in the SQL Commander and in Connection details. Before executing an SQL statement or connecting a database connection, a dialog is displayed, asking for replacement values.

These settings define a character sequence that identifies a variable and another sequence that delimits different parts of a variable.

Property	Description
Variable Identifier	The identifier for a variable. A variable starts and ends with this identifier. Default is "\$\$".
Variable Delimiter	The delimiter used to identify the parts of a variable. Default is " ".

Transaction

Property	Description
Pending Transactions at Disconnect	Specifies what DbVisualizer does on exit from the application, when the auto commit setting is disabled.

Bookmarks

Property	Description
Number of Bookmarks Limit	Specifies the number of SQL bookmarks that the New and History bookmark objects may keep before the lists are truncated.

Monitor

Property	Description
Start Monitors Automatically	Check to enable start of monitors automatically when database connections are established.

Form Viewer

Property	Description
Right Aligned Numbers	If enabled, numbers are displayed as right-aligned in the Form Editor/Viewer.

Grid

Property	Description
Auto Resize Column Widths	If Auto Resize is enabled, DbVisualizer automatically sizes each grid column based on the widest cell value. If Consider Column Header is also enabled, the header widths are also considered when calculating the column widths.
Show Grid Row Header	If enabled, a row header is shown also for read-only result set grids, such as monitoring result set grids.
Max Column Separator Width	This setting is used only when Auto Resize Column Widths is enabled and specifies a maximum visual column width for grids.
Image Thumbnail Size	Specifies the size (in pixels) for the widest side of a thumbnail shown for image data in the Data Form Viewer.
Meaning of setting Max Chars	<p>The Max Chars property in the Database Objects Data tab and in the SQL Commander is used to control the max number of characters for text values. If the number of characters for a text column is more than this setting, the column is colored in a light red color and the value is truncated as specified by this property:</p> <ul style="list-style-type: none">• Truncate Values Truncate the original value to be less than the setting of Max Chars. Note: this affects any subsequent edits and SQL operations that use the value since it's truncated. This setting is only useful to save memory when viewing very large text columns.• Truncate Values Visually Truncate the visible value only and leave the original value intact. This is the preferred setting since it will not harm the original value. The disadvantage is that more memory is needed when dealing with large text columns.

Copy

The Copy category groups properties that control the result of using Copy Selection and Copy Selection (With Column Header) via the grid right-click menu, the corresponding key bindings, and drag and drop.

Property	Description
Column Delimiter	Specifies the delimiter between columns in a multi column copy
End of Line Delimiter	Specifies the new line control characters for multi row copy requests

Colors

The Colors category is used to define alternative background, foreground and grid colors for grid components.

Binary/BLOB and CLOB Data

Property	Description
BLOB	Specifies how BLOB and binary data values are represented in grids. Setting this property to By Value results in performance penalties and the memory consumption increases dramatically.
CLOB	Specifies how BLOB and binary data values are represented in grids. Setting this property to By Value results in performance penalties and the memory consumption increases dramatically.

SQL Editor

The editor category controls various settings specific for the SQL Commander editor.

Property	Description
Tabs	Specifies settings for the tab keyboard key: Tab Size (the number of characters a tab character corresponds to), Whitespace(s) per Tab (by how many characters to indent when the tab key is pressed), and Expand Tab to Whitespace (if enabled, always insert space characters when the tab key is pressed). If Expand Tab to Whitespace is disabled, a tab character is inserted when the tab key has been clicked as many times as it takes to indent to the value specified by Tab Size, i.e., if Whitespace(s) per Tab is set to 4 and Tab Size is set to 8, clicking the tab key twice results in a tab character.
Recent Files Limit	Specifies the max number of files listed in the File->Load Recent sub menu.
Confirm Close of Unsaved Editors	If enabled, DbVisualizer asks you whether to save the text in an SQL editor with modified content (any editor; not only editors loaded from file) when you close the editor.
Set "Sticky" for SQL Editor(s)	If enabled, the Sticky flag is automatically set for all new SQL Editors, which means that the database connection details only can be changed manually.

Statement Delimiters

Statement delimiters define how a script should be divided into specific SQL statements in the pre-processing phase.

Property	Description
SQL Statement Delimiter 1	Defines the character(s) used to delimit one SQL statement from another in a SQL script
SQL Statement Delimiter 2	Defines the additional character(s) used to delimit one SQL statement from another in a SQL script. If there is no need for more than one SQL statement delimiter, set this one to the same as delimiter 1.
Allow "go" as Delimiter	Specifies whether go as the first word on a single line should be interpreted as a statement delimiter.
Begin Identifier	Defines the character(s) that identifies the start of an anonymous SQL block
End Identifier	Defines the character(s) that identifies the end of an anonymous SQL block

SQL Formatting

The SQL formatting category groups properties to control the SQL formatting feature in the SQL Commander. To see the effect of each property, modify it, press Apply and format the SQL in the SQL Commander.

Auto Completion

This category is used to define the visual appearance of the auto completion popup in SQL Editors.

Property	Description
Sort Tables List	Enable this to always present tables sorted in the auto completion popup
Sort Columns List	Enable this to always present column names sorted in the auto completion popup
Display Automatically	Enable this and the auto completion popup is automatically displayed whenever possible
Instant Substitution	Enable this and the auto completion feature substitutes directly if there is only one matching entry
Display Delay	Specifies the time in milliseconds until the auto completion popup is displayed automatically

Comments

Property	Description
Single Line Identifier 1	Specifies the character(s) that identifies the beginning of a one line comment
Single Line Identifier 2	Specifies the additional character(s) that identifies the beginning of a one line comment
Block Comment Begin Identifier	Specifies the character(s) that identifies the start of a multi line comment block
End	Specifies the character(s) that identifies the end of a multi line comment block

Debug

The Debug category holds properties that control the amount of output that is produced when setting various debug modes. Normally only error messages are displayed in the default debug destination, which is the Tools-

>Debug Window. The support team often refer to the debug properties when we want more information in a problem situation.

Property	Description
Debug Output Destination	Specifies the destination for all debug messages. It is not advisable to set this to Off, since this means that also error messages will be ignored. Standard Out is only useful if the debug mode of the DbVisualizer launcher is enabled.
Debug DbVisualizer	Defines the amount of logging that will be produced, where the Debug setting results in the most amount of output and Error in the least amount. Setting Detail Level to Full produces the most detail, but it also consumes more resources.
Debug JDBC Drivers	This property enables any debug output produced by a JDBC driver to be captured. The amount of output depends on the actual drivers.

Read more about using the debug output in the [Problem Resolution](#) section.

Database Settings

Database settings extends the General settings with properties that may have different values per supported database type. You specify the database type for a connection by choosing the appropriate type from the Database Type list in the Connection tab. If there is no matching entry, use the Generic database type.

The database type specific properties in the Tool Properties apply to all connections of the specific database type. You can also override these properties in the Connection Properties tab for a specific connection, in case you need to use different values for connections of the same database type.

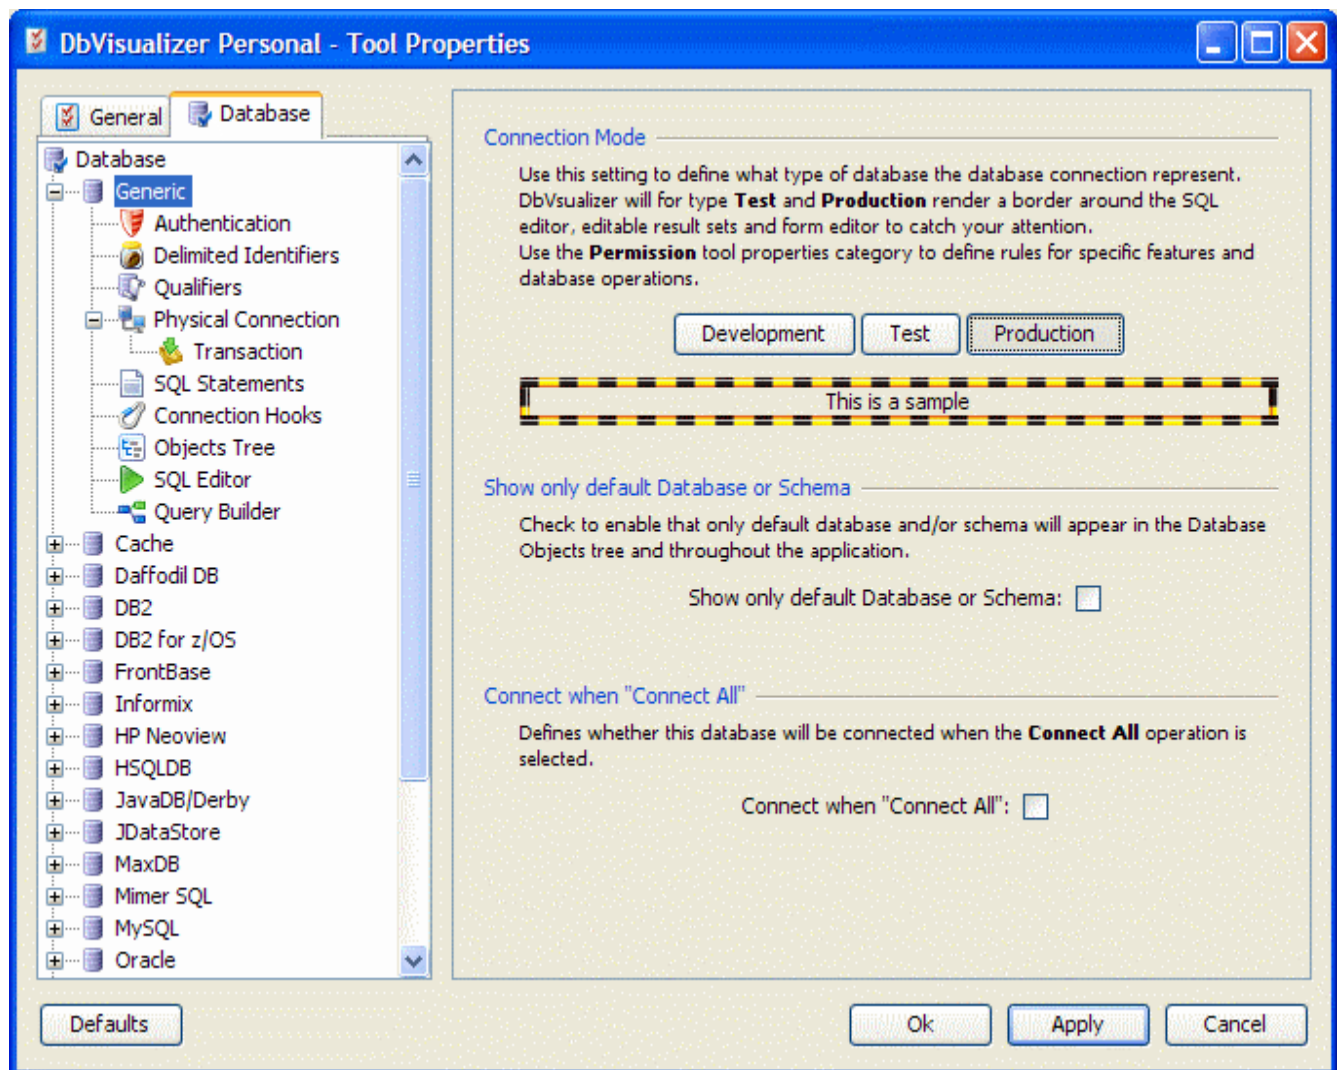


Figure: The Tool Properties window showing the tree with Database categories

The following properties are displayed when selecting a database type in the tree.

Property	Description
Connection Mode	Specifies the connection mode for the database connection: Development, Test or Production. Permissions are based on connection mode. For the Test and Production modes, DbVisualizer displays a border around areas where database content can be edited, to bring

	your attention to the fact that you are connected to a database where others may be affected by your changes.
Show only default Database or Schema	Enable this if you only want the default database or schema listed in the database objects tree.
Connect when "Connect All"	The Connect All feature allows you to connect to multiple database connections with a single click. Enable this property to include database connections of this type when using the Connect All feature.

Authentication

Property	Description
Save Password	If enabled, DbVisualizer saves the password for the database connection between invocations. (The password is saved encrypted)
Clear Password at Disconnect	If enabled, the password is cleared at disconnect
Require Userid	If enabled, you are asked to enter a userid whenever the database connection is established
Require Password	If enabled, you are asked to enter a password whenever the database connection is established

Delimited Identifiers

Delimited identifiers are identifiers which do not need to follow the rules of regular database object identifiers. Usually, delimited identifiers are used when you need to use SQL reserved words, spaces and mixed case sequences in an identifier.

Property	Description
Begin Identifier	Defines the start character for a delimited identifier. Normally, this is a double quote (")
End Identifier	Defines the end character for a delimited identifier. Normally, this is a double quote (")
Scripting	Enable this to use delimited identifiers in the Scripting features
Auto Completion/Query Builder	Enable this to use delimited identifiers in the auto completion and query builder features

Qualifiers

These properties control whether table and column names should be qualified when DbVisualizer generates SQL statement.

Property	Description
Qualify with Schema/Database: Scripting	Enable this to qualify object names with the schema/database in the Scripting features
Qualify with Schema/Database: Auto Completion/Query Builder	Enable this to qualify object names with the schema/database in the auto completion and query builder features.
Qualify Columns: Auto Completion/Query	Enable this to qualify column names with the table name in the

Builder

auto completion and query builder features.
Note: When you specify a table name alias, it is always used as a column name qualifier, regardless of this property setting.

Physical Connection

The Physical Connection category controls whether DbVisualizer should use only one physical connection with the database server or if physical connections will be acquired when needed. The Use Single Shared Physical Database Connection is disabled by default. If enabled then briefly it means that whenever establishing a connection DbVisualizer will assign one physical database connection for the objects tree and one per every SQL editor in the SQL Commander. The physical connection for a SQL editor is not acquired directly when the editor is created but rather when doing the first execute in it.

If enabling Use Single Shared Physical Database Connection then only one physical connection will be used for that database. DbVisualizer will then share the physical connection among all features communicating with the database. If using a single physical connection and auto commit is off then a confirmation dialog may appear when launching features that require transaction control and if there are uncommitted changes in the database.

Transaction

Property	Description
Auto Commit	Defines if each executed SQL statement will be auto committed or not. This setting applies for all SQL statements that are executed in the SQL Commander.
Ask when Auto Commit is Off	If auto commit is off then this setting when enabled will show a confirmation dialog if there are uncommitted changes produced by the last execution in the SQL Commander.
Transaction Isolation	Attempts to change the transaction isolation level for all database connections. Note: If this property is changed during a transaction, the result is JDBC driver specific.
Commit Batch Size	Specifies after how many rows DbVisualizer commits the transaction when saving a batch of changes in the table data editor.

SQL Statements

This category controls the SQL templates that DbVisualizer uses internally throughout the application. Each SQL template is composed of the standard SQL and variables. Variables are identified with \$\$...\$\$\$. DbVisualizer relies on a number of predefined variables, listed in the SQL Templates area right-click menu:

catalog
catalogseparator
schema
schemaseparator
table
table-name
where-columns
quoted-where-columns
columns
values
column-values
create-columns
index-type
index
unique
index-columns
create-primary-key
DbVis-Date
DbVis-Time

Figure: All predefined variables

A specific predefined variable can be used in one or more of the SQL templates. Using a variable in a SQL statement that is not valid will result in the variable appearing as-is when the statement is executed.

There is normally no reason to modify the SQL templates, nor the variable identifier or delimiter settings. There might however be circumstances when edits are needed:

- To put quotes or brackets around table names
- To change the variable identifier or variable delimiter since the default settings may interfere with object names in the database
- To modify the appearance of the where clause or the list of columns

Property	Name	Description
SQL Templates	SELECT ALL	Command used when selecting all rows for a table
	SELECT ALL WHERE	Command used when selecting some rows for a table
	SELECT COUNT	Command used to get the number of rows in a table
	INSERT INTO	Command used to insert a new row into a table
	UPDATE WHERE	Command used to update an existing row in a table
	DELETE WHERE	Command used to delete a specific row in a table
	DROP TABLE	Command used to drop a specific table
	CREATE TABLE	Command used to create a new table with an optional primary key
	CREATE INDEX	Command used to create an index for a specific table
	Monitor Row Count	Command used to get the number of rows in a table and the current time stamp

Monitor Row Count Change	Command used to get the row count difference in a table compared to the previous execution. The calculated row count and the current time stamp is returned
--------------------------	---

Connection Hooks

Connection hooks defines optional SQL commands that are sent to the database at connect and just before disconnect. They are typically used to initialize the database session with custom settings and to clean up various resources at disconnect.

Property	Description
Run SQL at Connect	Defines the SQL to be executed just after the connection has been established
Run SQL at Disconnect	Defines the SQL to be executed just before the connection will be disconnected

Objects Tree

Property	Description
Custom Object Tree Labels	Here you can define custom tree labels for the data nodes in the database objects tree. The Object Type must match the corresponding type in the actual database profile, see more below.

The label for a *data node* (e.g., a table or view node, as opposed to a node that just groups nodes, such as the Tables node) is typically the name of the database object the node represents, e.g., the table or view name. In some cases, you may want to extend the label to include other information, such as the name of the schema that the object belongs to. To do this, you can use a custom tree label, defined in the Objects Tree properties category.

You need two pieces of information to define a custom label: the Object Type name for the data node, and the names of the variables that hold the information you want to use in the label. You find this information in the `<ObjectsTreeDef>` element in the database profile XML file (described in detail in the [Plug-in Framework](#) section) for the database type you want to modify.

Using the database profile for the JavaDB/Derby database type as an example, a stripped down version of the `<ObjectsTreeDef>` element looks like this:

```
<ObjectsTreeDef id="derby">
  <GroupNode type="Schemas" label="Schemas">
    <DataNode type="Schema" label="${derby.getSchemas.Schema}">
      <SetVar name="schema" value="${derby.getSchemas.Schema}"/>
      <SetVar name="schemaId" value="${derby.getSchemas.Schema Id}"/>
      [...]
    <GroupNode type="Tables" label="Tables">
      <DataNode type="Table" label="${derby.getTables.Table Name}" isLeaf="true">
        <SetVar name="objectname" value="${derby.getTables.Table Name}"/>
        <SetVar name="rowcount" value="true"/>
        <SetVar name="acceptInQB" value="true"/>
        [...]
      </DataNode>
    </GroupNode>
  </DataNode>
  [...]
</ObjectsTreeDef>
```

In this example, there is one `<DataNode>` element with a `type` attribute set to `Schema`, with a nested `<DataNode>` element with a `type` attribute set to `Table`. These two elements represent data nodes, for the schema and table node, respectively, and the `type` attribute value is the Object Type name you need to bind the custom label to an object type.

Each `<DataNode>` element also has a number of nested `<SetVar>` elements, declaring the variables you can use in the custom label value. All variables declared for the object type node and those declared for a parent `<DataNode>` element can be used in the label. So, if you want the label for table nodes in the tree to show both the schema name and the table name, you add a custom label declaration like this:

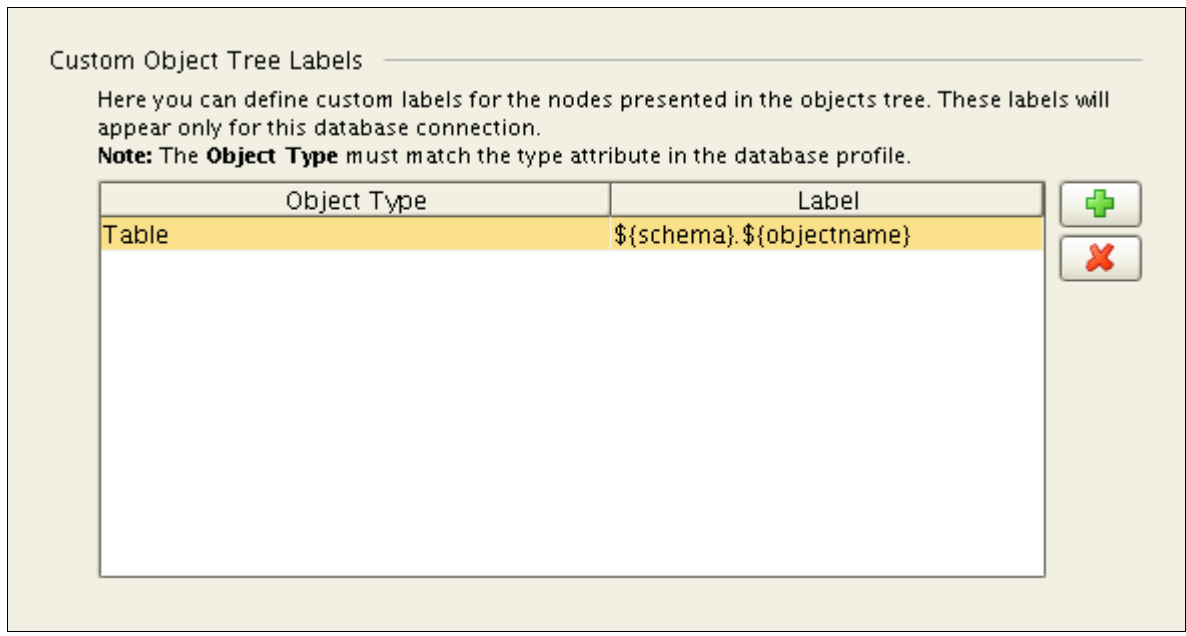


Figure: Custom label declaration

SQL Editor

Property	Description
Remove New Line Characters	Specifies whether any new line characters should be removed from any SQL statement executed in the SQL Commander, and in the implicit SQL execution functionality in DbVisualizer. Some drivers/databases such as DB2 require that no new line characters are part of any executed SQL.
Set Current Schema	<p>If enabled, changing the schema in the SQL Commander also changes the default schema for the database connection, so that unqualified table names in any SQL statement are associated with the selected schema. If this property is not enabled, changing the schema only affects the schemas used for auto-completion.</p> <p>Note: Only a few databases supports setting the default schema for an opened connection. This property is only shown for database types that support it.</p>

Query Builder

Property	Description
Query Builder Auto-Join Properties	With auto-join enabled, the Query Builder automatically joins tables as they are included in the query, based on the specified column matching rule: FK/PK declarations or columns with matching names in different tables.
Generate JOIN clauses in Query Builder	<p>Specifies whether the Query Builder generates joins as JOIN clauses or WHERE conditions.</p> <p>JOIN clause:</p> <pre>SELECT * FROM HR.EMPLOYEES emp INNER JOIN HR.DEPARTMENTS dept ON (emp.DEPARTMENT_ID = dept.DEPARTMENT_ID)</pre> <p>WHERE condition:</p> <pre>SELECT * FROM HR.EMPLOYEES emp, HR.DEPARTMENTS dept WHERE (emp.DEPARTMENT_ID = dept.DEPARTMENT_ID)</pre>

Database Specific settings

DbVisualizer provides more support for some databases than for others, and so requires extended configuration capabilities for these databases.

Data Types (Oracle)

With Oracle, the DATE data type should sometimes be handled as TIMESTAMP. Enable Handle DATE as TIMESTAMP and DbVisualizer automatically treat DATE as TIMESTAMP.

Data Types (DB2 and JavaDB/Derby)

DB2 and JavaDB/Derby supports a data type named CHAR FOR BIT DATA. If you want to see values of this type as text, enable this property.

Explain Plan (Oracle, SQL Server and DB2)

The explain plan feature supported for Oracle, SQL Server and DB2 can be configured to highlight certain threshold levels.

Property	Description
Color Critical Nodes	If enabled, critical nodes in the explain plan are highlighted.
Critical Threshold	Specifies the threshold for when a node should be handled as critical
Warning Threshold	Specifies the threshold for when a node should be handled as a warning

Explain Plan (Oracle)

The explain plan feature for Oracle can be configured to define the management of the underlying plan table in which the explain plan result is stored.

Explain Plan (DB2)

The explain plan feature for DB2 can be configured to define the management of the underlying plan tables in which the explain plan result is stored.

System Tables (Oracle)

Select here whether the database profile for Oracle should retrieve database information from the DBA or ALL system tables.

Note: If choosing DBA, make sure the appropriate privileges are granted for the user you are connecting as.

Export, Import and Print

Introduction


The Export feature writes data that has been fetched and presented in DbVisualizer to a file. The Export Wizard dialog looks different depending on whether a [grid](#), [graph](#) or [chart](#) data is being exported. The following sections describe the options available for each of these cases. There are major differences between DbVisualizer Free and DbVisualizer Personal when exporting grid data. This document explains the complete functionality in the Personal edition, some of which is not available in DbVisualizer Free.

The Import feature reads data stored in CSV (Character Separated Values) format from files.

The Printing feature prints grid and graph data to a printer or a file.

Exporting very large result sets using the standard export feature may fail due to running out of memory, since all data must first be presented in DbVisualizer. The [@export](#) client side command in the SQL Commander solves this problem, since it exports the data on the fly while it is fetched from the database.

Export Grid data

The Export wizard is launched using the Export button in the grid toolbar () or from the grid's right-click menu. If you want to export just some of the grid rows and columns instead of all data in the grid, select the data to export and launch the wizard with the Export Selection right-click menu choice.

Settings page

The first wizard page is the Settings page, containing general properties for how the exported data should be formatted. All settings in the settings page can be saved to a file for later use in the export wizard or in the SQL Commander when exporting result sets using the [@export editor command](#).

The screenshot shows the 'Export Grid' settings dialog box in DbVisualizer Personal. The window title is 'DbVisualizer Personal - Export Grid'. The dialog is titled 'Export Grid' and 'Settings'. It is divided into several sections:

- Output Format:** Radio buttons for CSV (selected), HTML, SQL, XML, and XLS. An 'Encoding' dropdown menu is set to 'Big5'.
- Data Format:** Fields for Date, Time, Timestamp, Number, Decimal Number, Null Value, and Binary/BLOB. Each has a dropdown menu and an example value.
 - Date: Unformatted (Example: 2007-07-02)
 - Time: Unformatted (Example: 11:05:09)
 - Timestamp: Unformatted (Example: 2007-07-02 11:05:09.624)
 - Number: Unformatted (Example: 9126183)
 - Decimal Number: Unformatted (Example: 9126183.531815)
 - Null Value: (null)
 - Binary/BLOB: Size
- Quote Text Data:** Radio buttons for None (selected), Single, Double, and ANSI. Example: O'Learys -> O'Learys.
- Options:** Checkboxes and dropdowns for:
 - Column Delimiter: TAB
 - Row Delimiter: UNIX/Linux/Mac OS X - LF
 - Include Column Names: checked
 - Row Comment Identifier: empty text box
 - Remove Newline Characters: unchecked
 - Include Original SQL: unchecked

At the bottom, there are buttons for 'Settings' (with a dropdown arrow), '< Back', 'Next >', and 'Cancel'.

Figure: The grid export wizard

Read the sections below for detailed information on each field and what settings that can be made.

Output Format

Grid data can be exported in the following formats.

Format	Description
CSV	<p>The CSV format (Character Separated Values) is used to export the grid of data to a file in which each column is separated with a character or several. It is even possible to specify the row delimiter (aka newline sequence of characters).</p> <pre>5,Hepp,59248 15,Hopp,41993 16,Hupp,44115</pre> <p>The above example use a "," as the column delimiter and a "\n" sequence as the row delimiter (invisible above).</p>
HTML	<p>The data is exported in HTML format using the <TABLE> and associated tags.</p>
SQL	<p>The SQL format simply creates an SQL INSERT statement for each row in the grid. It also uses the column names from the grid to define the column list in the SQL statement.</p> <pre>insert into table1 (Column1, Column2, Column3) values (5, 'Hepp', 59248); insert into table1 (Column1, Column2, Column3) values (15, 'Hopp', 41993); insert into table1 (Column1, Column2, Column3) values (16, 'Hupp', 44115);</pre>
XML	<p>The XML format is handy when importing or using the exported data in an XML enabled application. The default structure of the XML format is:</p> <pre><ROWSET> <ROW> <Column1>5</Column1> <Column2>Hepp</Column2> <Column3>59248</Column3> </ROW> <ROW> <Column1>15</Column1> <Column2>Hopp</Column2> <Column3>41993</Column3> </ROW> <ROW> <Column1>15</Column1> <Column2>Hupp</Column2> <Column3>44115</Column3> </ROW> </ROWSET></pre> <p>Alternatively, you can choose between the commonly used XmlDataSet and FlatXmlDataSet formats.</p>

Encoding

The encoding choice specifies which character encoding to use for the data. It is also used to set the encoding in the HTML and XML headers. The default choice is based on your systems default encoding.

Data Format

The data format settings defines how the data for each of the data types will be formatted.


Quote Text Data

Defines if text data should appear between quotes or not. Selecting the ANSI choice will automatically prefix any single quotes with an additional one.

Options

The options section is used to define settings that are specific for the selected output format.

CSV




The screenshot shows a dialog box titled "Options" with a light beige background. It contains the following settings:

- Column Delimiter: A dropdown menu with a semicolon (;) selected.
- Row Delimiter: A dropdown menu with "UNIX/Linux/Mac OS X - LF" selected.
- Include Column Names: A checked checkbox.
- Row Comment Identifier: An empty text input field.
- Remove Newline Characters: An unchecked checkbox.
- Include Original SQL: An unchecked checkbox.

Figure: CSV specific export options

HTML



The screenshot shows a dialog box titled "Options" with a light beige background. It contains the following settings:

- Title: A text input field containing "DbVisualizer export output".
- Description: A large empty text area.
- Include Original SQL: An unchecked checkbox.

Figure: HTML specific export options

SQL




The screenshot shows a dialog box titled "Options" with a light beige background. It contains the following fields and controls:

- Table Name:** A text input field containing "HR.COUNTRIES".
- Statement Separator:** A dropdown menu with a downward arrow, currently showing a semicolon (;).
- Row Comment Identifier:** A text input field containing "--".
- Include Create DDL:** A checkbox that is currently unchecked.
- Include Original SQL:** A checkbox that is currently unchecked.

Figure: SQL specific export options

XML



The screenshot shows a dialog box titled "Options" with a light beige background. It contains the following fields and controls:

- XML Style:** Three radio buttons: "DbVisualizer" (selected), "XmlDataSet", and "FlatXmlDataSet".
- Description:** A large empty text area.
- Include Original SQL:** A checkbox that is currently unchecked.

Figure: XML specific export options

Settings

Clicking the Settings button reveals a menu with options for saving and loading settings to and from a file.

- **Use Default Settings**
Use this choice to initialize the settings with default values. Some of the settings will be fetched from the general tool properties dialog.
- **Load**
Use this choice to open the file choose dialog, in which you can select a settings file
- **Save As**
Use this choice to save the settings to a file
- **Copy Settings to Clipboard**
Use this choice to copy all settings to the system clipboard. These can then be pasted into the SQL Commander to define the settings for [@export editor commands](#).

Data page

Clicking the Next button in the wizards moves you to the Data page. Use the columns list to control which columns to export and how to format the data for each columns. The list is exactly the same as the column headers in the original grid, i.e., if a column was manually removed from the grid before launching the Export wizard, then it will not appear in this list.

The screenshot shows the 'Export Grid' wizard's 'Data' page. The 'Columns' section contains a table with the following data:

Export	Name	Type	Is Text	Value
<input checked="" type="checkbox"/>	EMPLOYEE_ID	BigDecimal	<input type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	FIRST_NAME	String	<input checked="" type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	LAST_NAME	String	<input checked="" type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	EMAIL	String	<input checked="" type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	PHONE_NUMBER	String	<input checked="" type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	HIRE_DATE	Timestamp	<input checked="" type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	JOB_ID	String	<input checked="" type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	SALARY	BigDecimal	<input type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	COMMISSION_PCT	BigDecimal	<input type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	MANAGER_ID	BigDecimal	<input type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	DEPARTMENT_ID	BigDecimal	<input type="checkbox"/>	\$\$value\$\$

The 'Table Rows' section shows:

Total Number of Rows in Grid:

Number of Rows to Export:

Figure: The grid export wizard

The Table Rows fields show you how many rows are available and let you specify the number of rows to export. This setting along with the Add Row button is especially useful when you use the test data generation feature described in the next section.

Here follows information about the columns in the list.

Field	Description
Export	Defines whether the column will be exported or not. Uncheck it to ignore the column in the exported output.
Name	The name of the column. This is only used if exporting in HTML, XML or SQL format. Column headers are optional in the CSV output format.

Type	The internal DbVisualizer type for the column. This type is used to determine if the column is a text column (i.e., if the data should be enclosed by quotes or not).
Text	Specifies if the column is considered to be a text column (this is determined based on the type) and so if the value should be enclosed in quotes.
Value	The default "\$\$value\$\$" variable is simply be substituted with the column value in the exported output. You can enter additional static text in the value field. This is also the place where any test data generators are defined.

Generate Test Data

The test data generator is useful when you need to add random column data to the exported output.

The Value field specifies the data to be in the exported output. By default, it contains the \$\$value\$\$ variable, which is simply replaced with the real column value during the export process. You can also add static values before and after the \$\$value\$\$ field, to be exported as entered.

Alternatively, you can use test data generator variables in the Value field. The choices are available in the right-click menu when you edit the Value field.

The screenshot shows the 'Export Grid' dialog box with a table of columns. The 'NEWCOLUMN' row is selected, and a right-click context menu is open over it. The menu contains the following options:

- Generate random number
- Generate random string of random size
- Generate random value from a list of values
- Generate sequential number

Below the table, the 'Table Rows' section shows 'Total Number of Rows in Grid: 107' and 'Number of Rows to Export: 107'. At the bottom, there are buttons for 'Settings', '< Back', 'Next >', and 'Cancel'.

Figure: Right-click menu with the test data generator functions

Function Name	Function Call	Example
Generate random number	\$\$var randomnumber(1, 2147483647)\$\$	Generates a random number between 1 and 2147483647
Generate random string of random size	\$\$var randomtext(1, 10)\$\$	Generates random text with a length between 1 an 10 characters
Generate random value from a list of values	\$\$var randomenum(v1, v2, v3, v4, v5)\$\$	Picks one of the listed values in random order
Generate sequential number	\$\$var number(1, 2147483647, 1)\$\$	Generates a sequential number starting from 1. The generator re-starts at 1 when 2147483647 is reached. The number is increased with 1 every time a new value is generated.

Test data generator example

Here is an example of how to use the test data generators to try out planned changes to the data. Consider this initial data:


	EMPNO 	ENAME	JOB	MGR	HIREDATE	SAL	DEPTNO
1	7369	SMITH	CLERK	7902	2005-01-24 12:11:08.0	800	20
2	7499	ALLEN	SALESMAN	7698	2005-01-24 12:11:08.0	1600	30
3	7521	WARD	SALESMAN	7698	2005-01-24 12:11:08.0	1250	30
4	7566	JONES	MANAGER	7839	2005-01-24 12:11:08.0	2975	20
5	7654	MARTIN	SALESMAN	7698	2005-01-24 12:11:08.0	1250	30
6	7698	BLAKE	MANAGER	7839	2005-01-24 12:11:08.0	2850	30
7	7782	CLARK	MANAGER	7839	2005-01-24 12:11:08.0	2450	10
8	7788	SCOTT	ANALYST	7566	2005-01-24 12:11:08.0	3000	20
9	7839	KING	PRESIDENT	(null)	2005-01-24 12:11:08.0	5000	10
10	7844	TURNER	SALESMAN	7698	2005-01-24 12:11:08.0	1500	30
11	7876	ADAMS	CLERK	7788	2005-01-24 12:11:08.0	1100	20
12	7900	JAMES	CLERK	7698	2005-01-24 12:11:08.0	950	30
13	7902	FORD	ANALYST	7566	2005-01-24 12:11:08.0	3000	20
14	7934	MILLER	CLERK	7782	2005-01-24 12:11:08.0	1300	10
15	7939	MILLER	CLERK	7782	2005-01-24 12:11:08.0	1300	10

Figure: Sample of grid data

The Data page that produced the grid in the previous figure looks like this:

Columns

Export	Name	Type	Is Text	Value
<input checked="" type="checkbox"/>	EMPNO	BigDecimal	<input type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	ENAME	String	<input checked="" type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	JOB	String	<input checked="" type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	MGR	BigDecimal	<input type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	HIREDATE	Timestamp	<input checked="" type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	SAL	BigDecimal	<input type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	COMM	BigDecimal	<input type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	DEPTNO	BigDecimal	<input type="checkbox"/>	\$\$value\$\$

Buttons: ↑, ↓, Add Row

Figure: The export window

After the changes, the JOB column should not appear in the output and the new JOB_FUNCTION should contain abbreviated job function codes. To test this, we simply uncheck the Export field for JOB entry and set the Value for the JOB_FUNCTION to use the Generate random value from a list of values function.

Columns

Export	Name	Type	Is Text	Value
<input checked="" type="checkbox"/>	EMPNO	BigDecimal	<input type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	ENAME	String	<input checked="" type="checkbox"/>	\$\$value\$\$
<input type="checkbox"/>	JOB	String	<input checked="" type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	JOB_FUNCTION	String	<input type="checkbox"/>	\$\$var randomenum(eng, adm, fin)\$\$
<input checked="" type="checkbox"/>	MGR	BigDecimal	<input type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	HIREDATE	Timestamp	<input checked="" type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	SAL	BigDecimal	<input type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	COMM	BigDecimal	<input type="checkbox"/>	\$\$value\$\$
<input checked="" type="checkbox"/>	DEPTNO	BigDecimal	<input type="checkbox"/>	\$\$value\$\$

Buttons: ↑, ↓, Add Row

Figure: Customized columns list with a generator function

Previewing the data (or exporting it) in CSV format results in this:

```
EMPNO, ENAME, JOB_FUNCTION, MGR, HIREDATE, SAL, COMM, DEPTNO
7369, SMITH, adm, 7902, 2005-01-24 12:11:08, 800, (null), 20
7499, ALLEN, adm, 7698, 2005-01-24 12:11:08, 1600, 300, 30
7521, WARD, eng, 7698, 2005-01-24 12:11:08, 1250, 500, 30
7566, JONES, adm, 7839, 2005-01-24 12:11:08, 2975, (null), 20
7654, MARTIN, eng, 7698, 2005-01-24 12:11:08, 1250, 1400, 30
7698, BLAKE, eng, 7839, 2005-01-24 12:11:08, 2850, (null), 30
7782, CLARK, eng, 7839, 2005-01-24 12:11:08, 2450, (null), 10
7788, SCOTT, eng, 7566, 2005-01-24 12:11:08, 3000, (null), 20
7839, KING, eng, (null), 2005-01-24 12:11:08, 5000, (null), 10
7844, TURNER, eng, 7698, 2005-01-24 12:11:08, 1500, 0, 30
```

```

7876, ADAMS, fin, 7788, 2005-01-24 12:11:08, 1100, (null), 20
7900, JAMES, eng, 7698, 2005-01-24 12:11:08, 950, (null), 30
7902, FORD, eng, 7566, 2005-01-24 12:11:08, 3000, (null), 20
7934, MILLER, fin, 7782, 2005-01-24 12:11:08, 1300, (null), 10
7939, MILLER, fin, 7782, 2005-01-24 12:11:08, 1300, (null), 10
...

```

Preview

The third wizard page is the Preview page, showing the first 100 rows of the data as it will appear when it is finally exported. This is useful to verify the data before performing the export process. If the previewed data is not what you expected, just use the back button to modify the settings.

Output Destination

The final wizard page is the Output Destination page. The destination field specifies the target destination for the exported data.

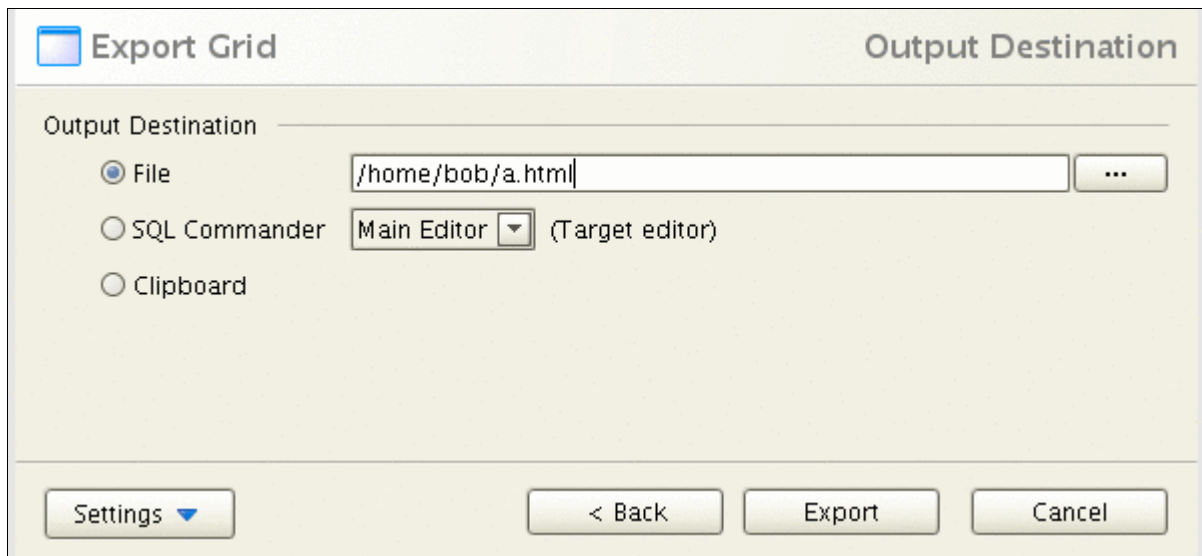


Figure: The output destination and final page for grid export

Destination	Description
File	This option outputs the data to a named file.
SQL Commander	This destination will transfers the export data to the SQL Commander editor. It is primarily useful when exporting the SQL output format.
Clipboard	Exporting to the (system) clipboard is convenient if you want to use the exported data in another application without the extra step of exporting to file first. CSV formatted data can even be pasted into a spreadsheet application such as Excel or StarOffice, and the cells in the grid will appear as cells in the spreadsheet. Read more about the CSV format in the Format section.

Export Text data

The wizard for exporting result sets in Text format is very simple, as it is only possible to specify where the exported output should go.

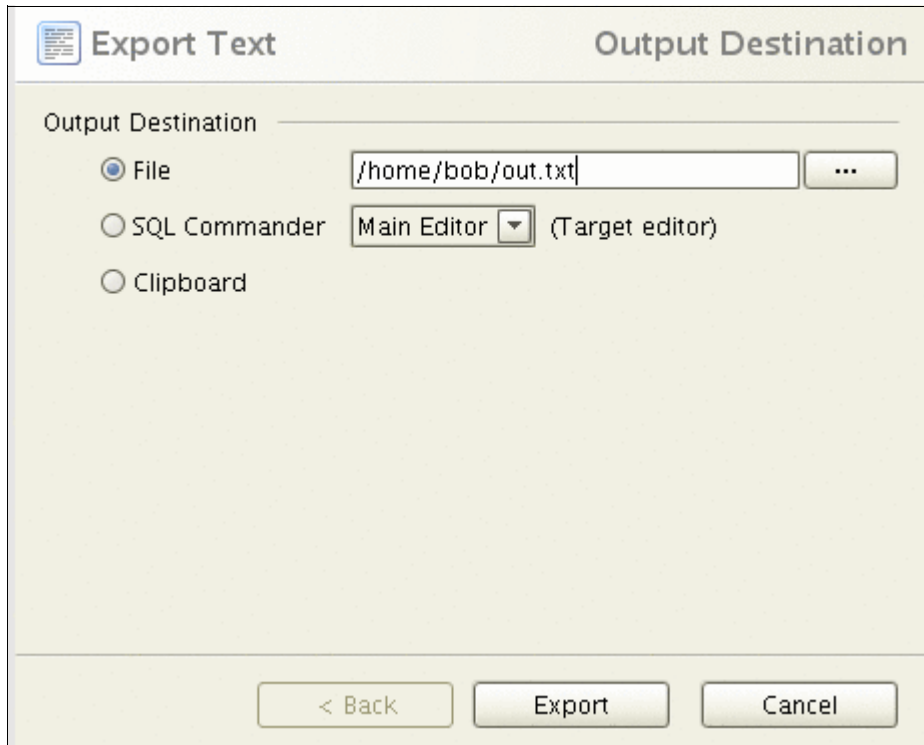


Figure: Export window for text format result sets

Export Graph data

Exporting references graphs exports the graph with the same zoom level as it appears on the screen. The Export wizard pages when exporting a graph looks like this:

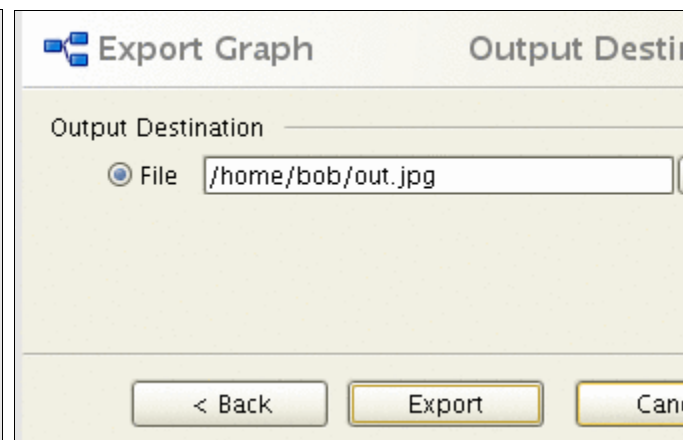
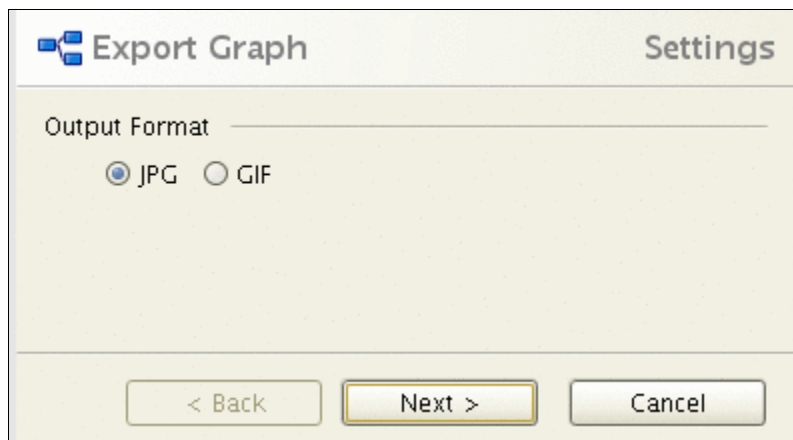


Figure: Export window for graphs

The graph can be exported to a File in the JPEG or GIF formats.

Export Chart data

The options when exporting charts are the same as for graphs, but also let you set the size and orientation to use for the chart in the file.

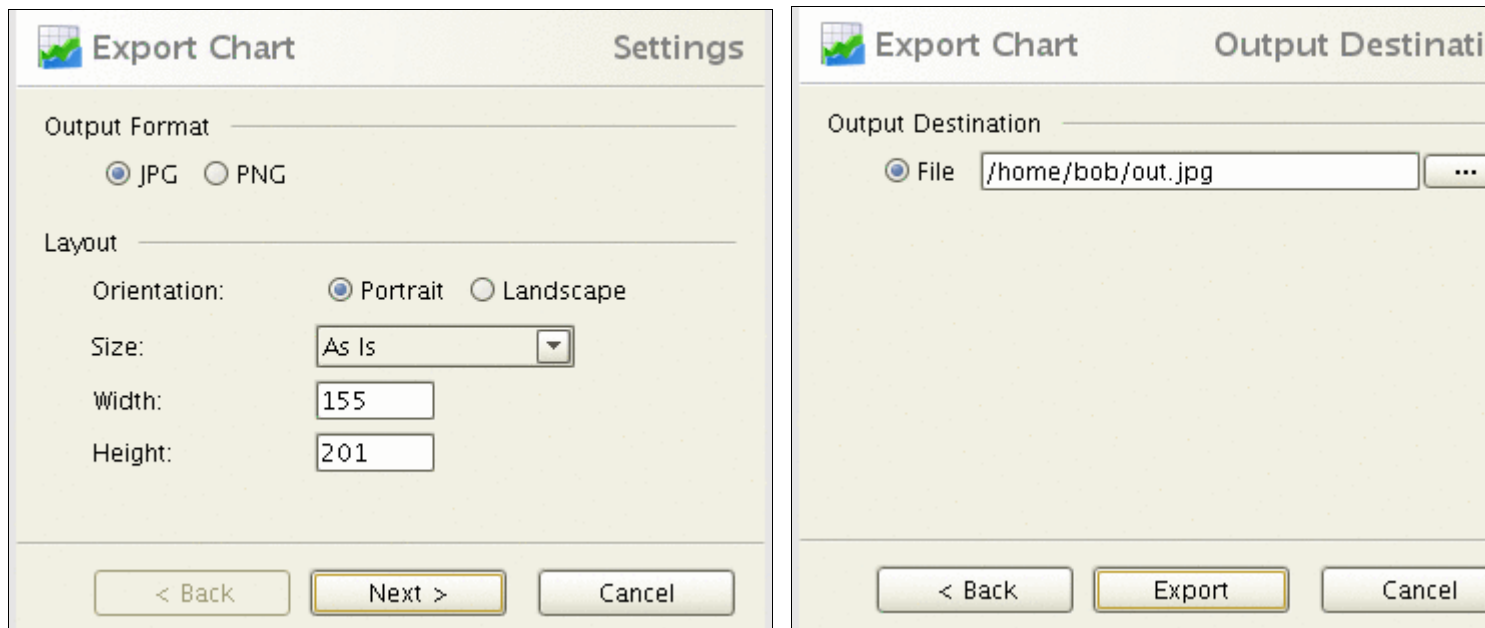


Figure: Export window for charts

A chart can be exported to a File in the JPEG and PNG formats. The optional Layout settings are used to control the size of the image. The default width and height are the same as the size of the chart as it appear on the screen. The Size list when clicked shows a list of well known paper formats. The Width and Height are changed to match the selected size. Whether setting the width and height manually or selecting a predefined size, the exported image is scaled accordingly.

Import Table Data

The Import Table Data feature is used to import files containing data organized as rows of columns with separator characters between them. The destination for the imported data can be a database table or a grid in DbVisualizer. The grid option is convenient for smaller files, as the features available for a DbVisualizer grid can then be used to do various things with the data. An example is that a CSV file can easily be converted into an XML file or a HTML document by using the data import feature to grid and then use the Export Wizard in the grid to output the grid data in the desired format.

The Import Wizard is launched via the right-click menu for table objects or via the Actions menu.

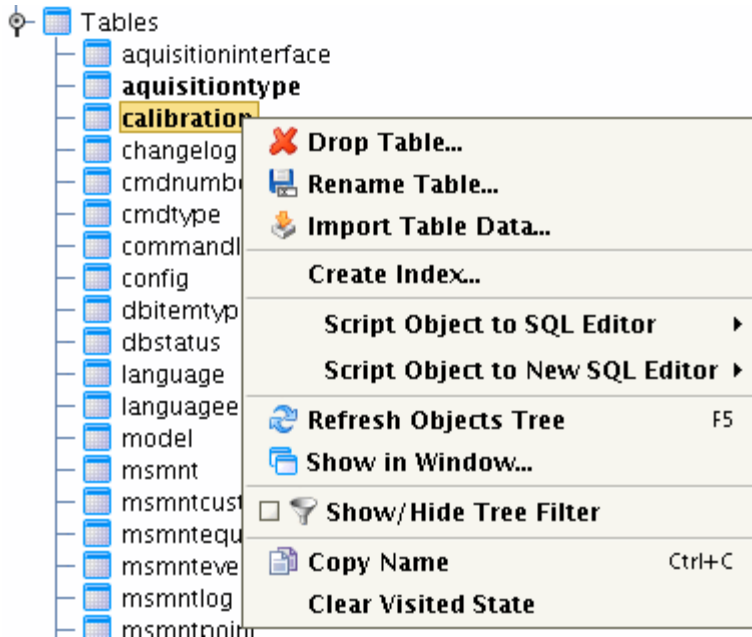



Figure: Import Table Data action in the right-click menu for table objects

Note 1: The first row in the source file can be used to name the columns.

Note 2: The Import Wizard can not be used to import binary data.

Source File

In the first wizard page, select the source file to import and then click the Next button.



Import CSV File Source File

Enter location and encoding of the file that should be imported.

Input File

File: ...

Encoding: ▼

Figure: The Source File import wizard page

Settings

In the Settings page, you specify how the data in the file is organized. The Data section at the bottom of the page shows a preview of the parsed data in the Grid tab and the original source file in the File tab. If a row in the Grid tab is red, it indicates that the row will be ignored during the import process. This happens if setting any of the Options settings results in rows not being qualified.

In the Delimiters section, define the character that separates the columns in the file. If you enable Auto Detect, DbVisualizer tries the following characters:

- comma ","
- tab "TAB"
- semicolon ";"
- percent "%"

Use the Options section to further define how the data should be read.

Import CSV File Settings

Specify options how the columns in the file should be identified. Use the Data Grid to verify that the file is properly read.

Delimiters

Column Delimiter: Auto Detect String

Options

Header in First Row:

Skip Empty Row(s):

Skip First Row(s):

Skip Rows Starting With:

Text Quoted Between:

Data

Grid File

EMPLOYEE_ID	FIRST_NAME	LAST_NAME	EMAIL	PHONE_NUMBER	HIRE_DATE
100	Steven	King	SKING	515.123.4567	1987-06-17 00:00:00
101	Neena	Kochhar	NKOCHHAR	515.123.4568	1989-09-21 00:00:00
102	Ilex	De Haan	LDEHAAN	515.123.4569	2005-06-29 23:10:53
103	Alexander	Hunold	AHUNOLD	590.423.4567	1990-01-03 00:00:00
104	Bruce	Ernst	BERNST	590.423.4568	1991-05-21 00:00:00
105	David	Austin	DAUSTIN	590.423.4569	1997-06-25 00:00:00

Preview Rows: Column Widths:

Figure: The Settings wizard page

The following shows the preview grid with some rows in red. The reason is that the Skip First Row(s) and Skip Rows Starting With are set, i.e., the first two rows and the rows starting with 103 will not be imported.

Import CSV File Settings

Specify options how the columns in the file should be identified. Use the Data Grid to verify that the file is properly read.

Delimiters

Column Delimiter: Auto Detect String

Options

Header in First Row:

Skip Empty Row(s):

Skip First Row(s):

Skip Rows Starting With:

Text Quoted Between:

Data

EMPLOYEE_ID	FIRST_NAME	LAST_NAME	EMAIL	PHONE_NUMBER	HIRE_DATE
100	Steven	King	SKING	515.123.4567	1987-06-17 00:00:00
101	Neena	Kochhar	NKOCHHAR	515.123.4568	1989-09-21 00:00:00
102	Ilex	De Haan	LDEHAAN	515.123.4569	2005-06-29 23:10:53
103	Alexander	Hunold	AHUNOLD	590.423.4567	1990-01-03 00:00:00
104	Bruce	Ernst	BERNST	590.423.4568	1991-05-21 00:00:00
105	David	Austin	DAUSTIN	590.423.4569	1997-06-25 00:00:00

Preview Rows: Column Widths:

Figure: The Settings wizard page

Data Formats

The Data Formats page is used to define formats for some data types. The first row in the preview grid contains a data type drop-down boxes. Just select the appropriate type for each column.

Import CSV File Data Formats

Define formats for the data types that appear in the file. Then use the drop down lists in the Data Grid to map data types with the columns in the file.

Data Formats

Date: Example: 2005-06-29

Time: Example: 23:47:02

Timestamp: Example: 2005-06-29 23:47:02

Thousand Separator:

Decimal Separator:

Null Values: Example: (null), NULL, nada

Boolean True: Example: true, 1, yes

Boolean False: Example: false, 0, no

Data

Grid

EMPLOYEE_ID	FIRST_NAME	LAST_NAME	EMAIL	PHONE
Number	String	String	String	String
101	Neena	Kochhar	NKOCHHAR	515.123
102	Ilex	De Haan	LDEHAAN	515.123
103	Alexander	Hunold	AHUNOLD	590.423
104	Bruce	Ernst	BERNST	590.423

Preview Rows: Column Widths:

Figure: The Settings wizard page

The following is displayed when selecting the drop-down box in the preview grid.

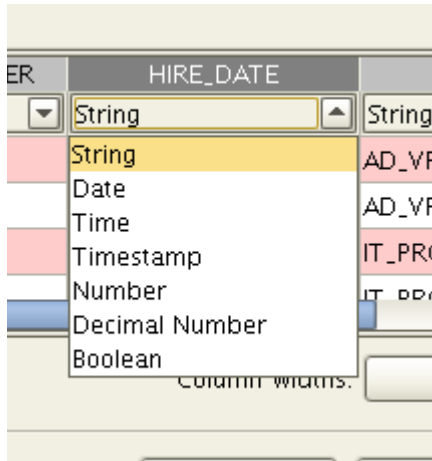



Figure: The data type drop down

Import Destination

The Import Destination page provides two options: Grid and Database Table. The Grid choice is used to import the data into a grid that will be presented in its own window in DbVisualizer.

When the Database Table choice is selected, the page shows information about the table in which the data will be imported. The Map Table Columns with File Columns grid shows the columns in the selected database table and the columns in the source file. DbVisualizer automatically assigns the columns in the source file with the first columns in the target table. You can then manually reassign them, using the drop-down lists in the File Column Name field. Choose the empty choice in the columns drop-down to ignore the column during import.

 Import CSV File
Import Destination

Select the destination the file should be imported to. Use **Database Table** to import the data into a table. Choosing **Grid** will import the file to a grid displayed in a separate window. You can then use the right click menu in the window to perform standard operations on the data. (The Grid choice is recommend for smaller i.e 50,000 rows).

Import Into

Grid
 Database Table

Database Table

Database Connection: Oracle 10g: system

Database:

Schema: HR

Table: EMPLOYEES

Map Table Columns with File Columns


Key	Table Column Name	Table Data Type	File Column Name
	EMPLOYEE_ID	NUMBER	EMPLOYEE_ID
	FIRST_NAME	VARCHAR2	FIRST_NAME
	LAST_NAME	VARCHAR2	LAST_NAME
	EMAIL	VARCHAR2	EMAIL
	PHONE_NUMBER	VARCHAR2	PHONE_NUMBER
	HIRE_DATE	DATE	
	JOB_ID	VARCHAR2	EMPLOYEE_ID
	SALARY	NUMBER	FIRST_NAME
	COMMISSION_PCT	NUMBER	LAST_NAME
	MANAGER_ID	NUMBER	EMAIL
	DEPARTMENT_ID	NUMBER	PHONE_NUMBER
			HIRE_DATE
			JOB_ID
			SALARY
			COMMISSION_PCT

Figure: The data type drop down

Import process

The last wizard page is used to start and monitor the import process. Here you can select whether all rows in the source file should be imported or only a portion. Errors that occur during the import process are presented in the log.

Import CSV File Import

This is the last page. Specify options for the import process and then press **Import** to start.

Options

Import all Rows

Import Only row(s)

Keep Window after Import

Stop on Error

Import Log

```
00:06:27 - Import started. Wait...
00:06:27 - [ Row: 23, File lines: 23/24 ] ORA-00001: unique constraint (HR.EMP_EMAIL_UK)
violated
00:06:28 - [ Row: 25, File lines: 25/26 ] ORA-00001: unique constraint (HR.EMP_EMAIL_UK)
violated
00:06:28 - [ Row: 70, File lines: 70/71 ] ORA-00001: unique constraint (HR.EMP_EMAIL_UK)
violated
00:06:28 - [ Row: 76, File lines: 76/77 ] ORA-00001: unique constraint (HR.EMP_EMAIL_UK)
violated
00:06:28 - [ Row: 81, File lines: 81/82 ] ORA-00001: unique constraint (HR.EMP_EMAIL_UK)
violated
00:06:29 - [ Row: 90, File lines: 90/91 ] ORA-00001: unique constraint (HR.EMP_EMP_ID_PK)
violated
00:06:29 - [ Row: 103, File lines: 103/104 ] ORA-00001: unique constraint
(HR.EMP_EMP_ID_PK) violated
00:06:29 - [ Row: 104, File lines: 104/105 ] ORA-00001: unique constraint
(HR.EMP_EMP_ID_PK) violated
00:06:29 - [ Row: 105, File lines: 105/106 ] ORA-00001: unique constraint
(HR.EMP_EMP_ID_PK) violated
00:06:29 - [ Row: 106, File lines: 106/107 ] ORA-00001: unique constraint
(HR.EMP_EMP_ID_PK) violated
00:06:29 - [ Row: 107, File lines: 107/108 ] ORA-00001: unique constraint
(HR.EMP_EMP_ID_PK) violated
00:06:29 - [ Row: 108, File lines: 108/109 ] ORA-00001: unique constraint
(HR.EMP_EMP_ID_PK) violated
00:06:29 - Import finished. 93 row(s) inserted and 12 failed.
```

< Back Import Cancel

Figure: The import process page

Print

The printing support in DbVisualizer supports printing of grids, graphs, charts and plain text, such as the content of an SQL Editor. The print dialog looks somewhat different depending on what is printed. In all cases, you launch the print dialog by clicking on the Print button in the toolbar for the object you want to print, or by choosing Print from the right-click menu. The right-click menu also contains a Print Preview choice, if you want to see what the printout will look like before you actually print.

Printer Setup

If you want to set the page orientation (e.g., portrait or landscape) and paper size, you must launch the Printer Setup dialog, using the File->Printer Setup main menu option, before you print. Printing varies widely between platforms, so even though the Print dialog (as opposed to the Printer Setup dialog) on some platforms also lets you choose a page orientation and other options, they may be ignored if specified in that dialog. The only supported way to specify the page orientation and other options is via the Printer Setup dialog.

Grid, Chart and Plain Text

For a grid, chart and plain text, DbVisualizer launches the platform's native Print dialog, so it looks different on different platforms. The two options available on all platforms are a choice of printer and the page range. On some platforms, the dialog may offer additional options, but they may be ignored by DbVisualizer. Use the Printer Setup dialog to set other options than which printer to use and the page range, as described above.

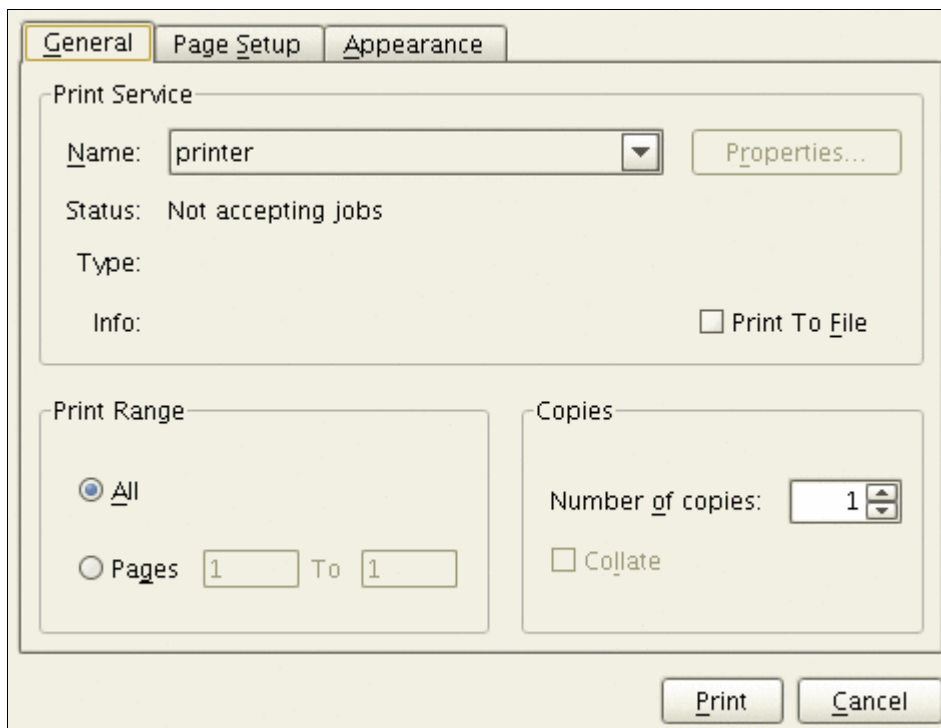


Figure: Standard print dialog

The figure above shows how the Print dialog looks on the Linux platform.


When you print a grid in DbVisualizer, the grid is printed as it is shown on the screen, i.e., with the table headers, sort and primary key indicator, etc. It is printed as a screenshot that may span several pages, depending on the number of rows and columns that are printed. For a grid, the right-click menu contains a Print Selection choice that you can use if you just want to print selected rows and columns.

An alternative to printing a grid as a screenshot is to export the grid to HTML and then use a web browser to print it.

Printing a chart scales the chart to the size of the paper. Plain text is printed as-is and may span multiple pages, both in height and width.

Graph

Printing a graph adds a custom dialog before the native Print dialog is displayed.



The dialog box contains the following elements:

- Poster Rows: 1
- Poster Columns: 1
- Add Poster Coords:
- Clip Area: Graph
- Buttons: Ok, Reset, Cancel

Figure: Print options when printing graphs

You can specify the number of rows (pages) and columns (pages) that the complete image will be split into. You can also select whether the view as it appears on the screen or the complete graph should be printed. When you click Ok, the native Print dialog is displayed, where you can select the printer.

Print Preview

Use the File->Print Preview feature to preview what the printout will look like before you actually print it.

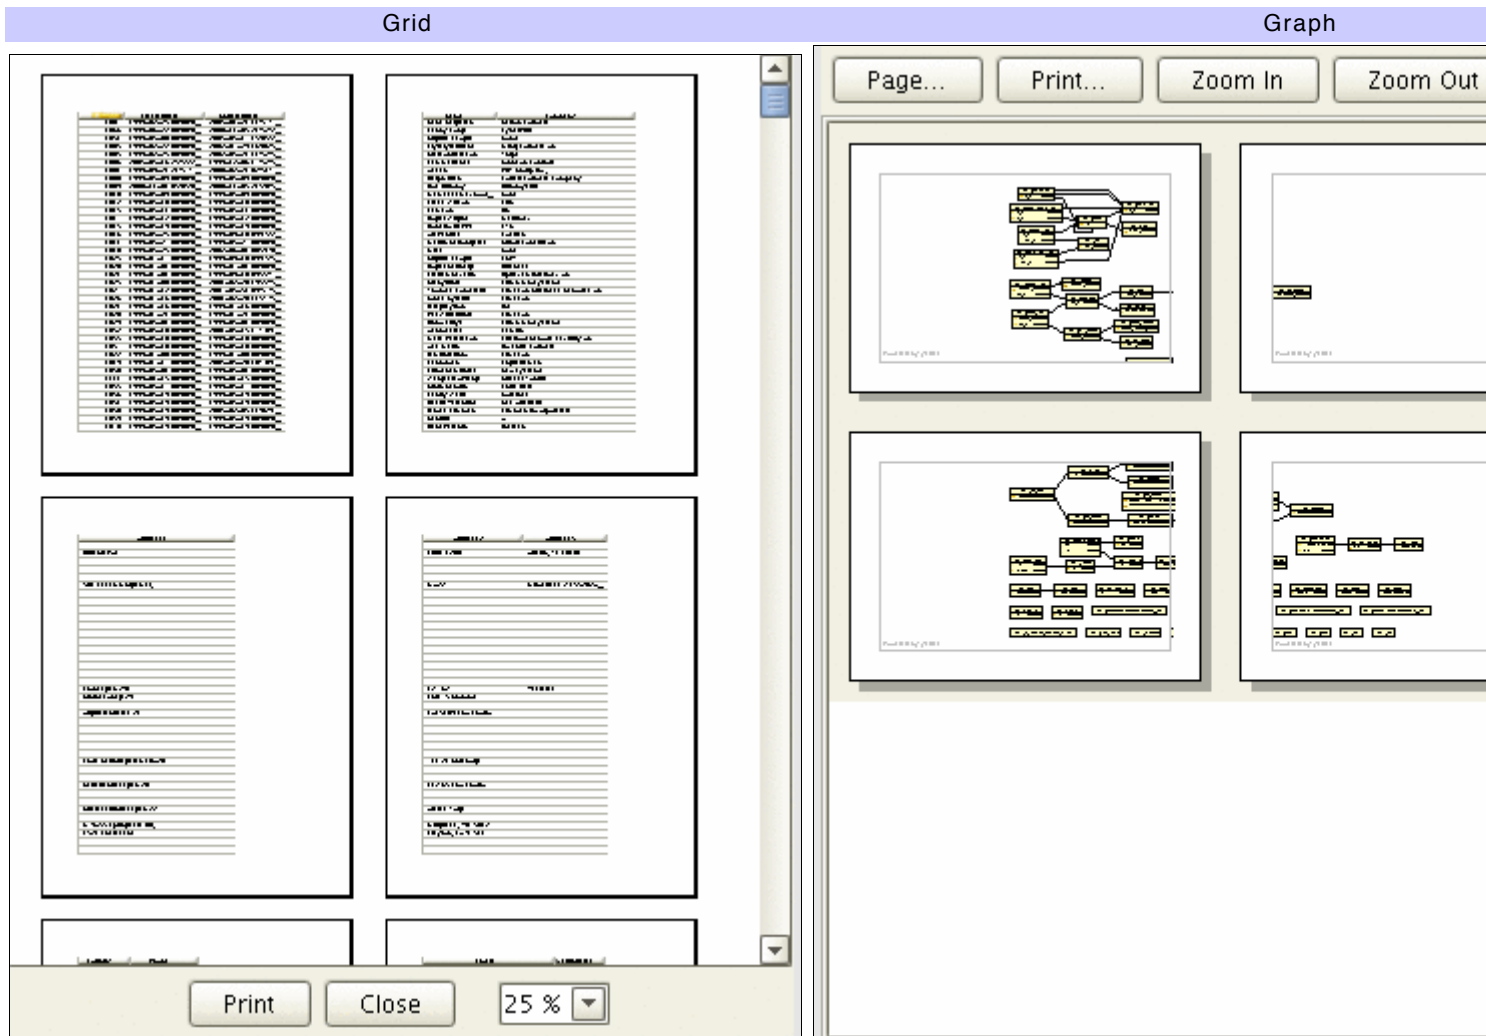


Figure: Grid and graph print previews

Plug-in Framework

Introduction

Note: The plug-in framework is supported only by the DbVisualizer Personal edition.

This document explains the database profile framework which is the base for how DbVisualizer presents information in the Database Objects tree and in the Object View. In addition, it is also used to define object actions, such as drop, rename, compile, create, comment, alter, etc.

What features in DbVisualizer relies on the database profile?

One of the most important and central features in DbVisualizer is the database objects tree, used to navigate databases, and the object view, showing details about specific objects. The general problem exploring any database is that they are all different with respect to the information describing what's in the database (also called system tables or database meta data). This basically means that it's rather complex to implement a multi-database support product, such as DbVisualizer, since each database must be handled specifically. All databases also support different object types, apart from the most common ones, such as table, view, index, etc.

The database profile framework is used to simplify the process of defining what information DbVisualizer will display and operate on for a specific database. Technically, a database profile is an XML document with all of the logic, structure and actions easily mapped to the visual components in DbVisualizer. Another great benefit of separating the database specific logic from the implementation of DbVisualizer is that anyone with some degree of domain knowledge can create a database profile. All that is needed is a text editor (preferably with XML support) and some ideas of what should be the final result.

A great source for inspiration (except for this document) is all the existing database profiles that comes with DbVisualizer. All database profiles are (and must be) stored in the `DBVIS-HOME/resources/profiles` directory.

The following figure illustrates which features in DbVisualizer are controlled by the database profile.

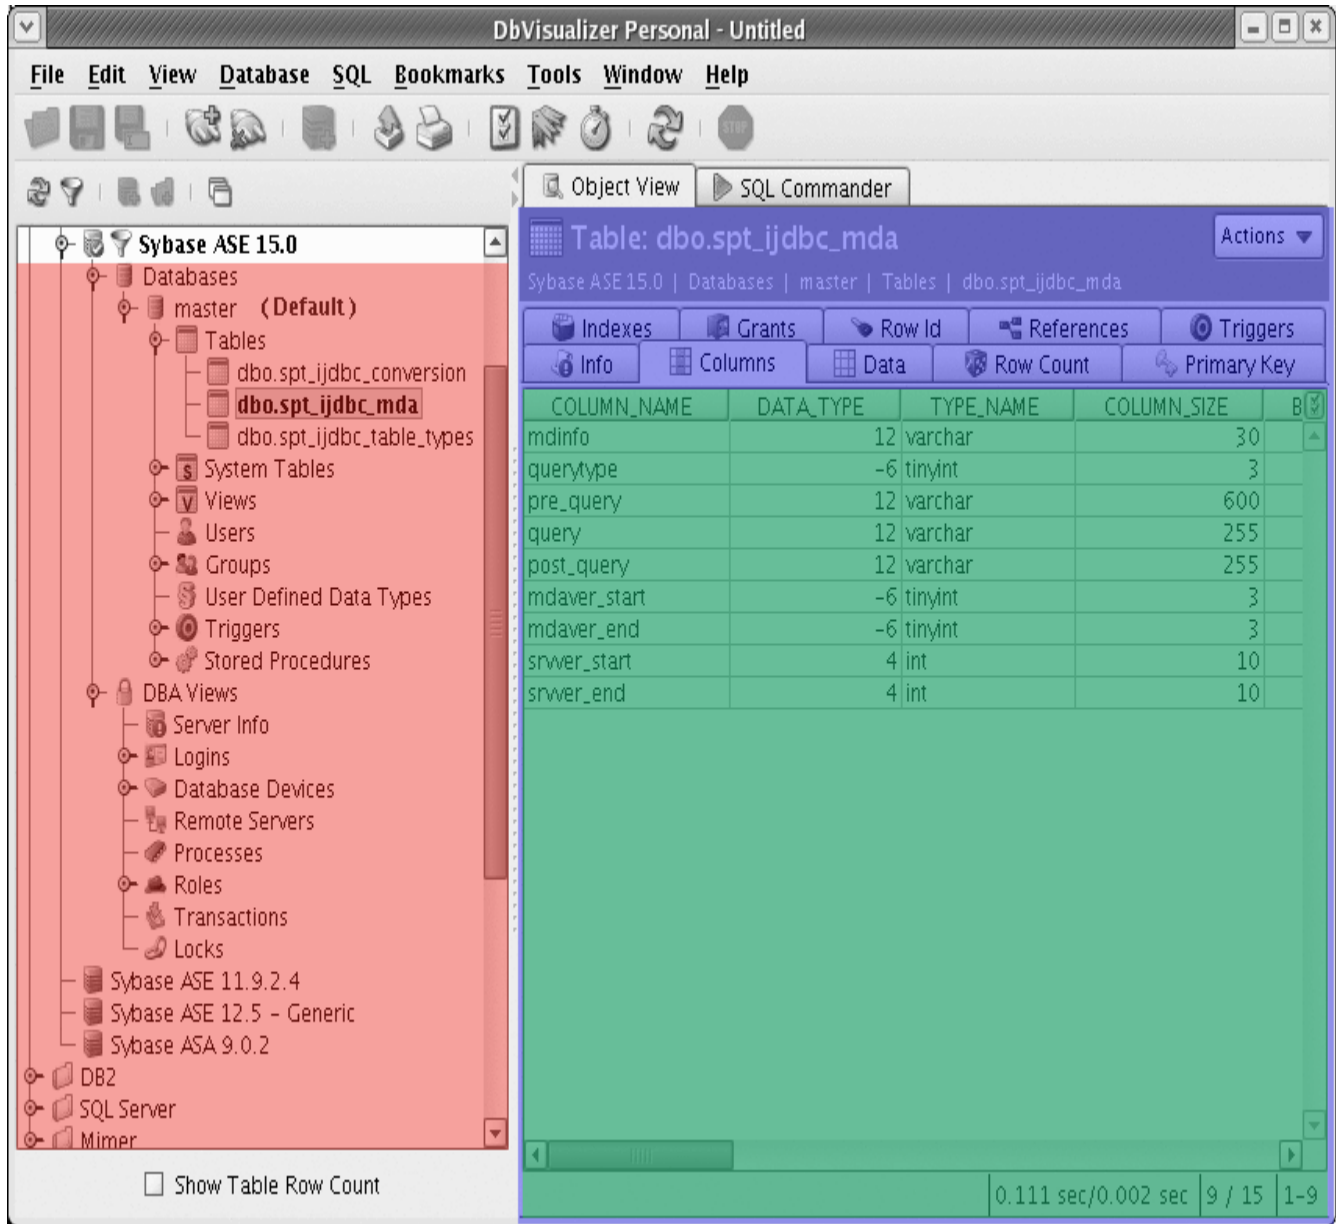


Figure: What the database profile controls in DbVisualizer

The red box at the left shows the database objects tree. This tree is used to navigate the objects in the database. Selecting an object in the tree shows the object view (blue box) for the selected object type. An object view may have several data views (green), showing object information. DbVisualizer shows these as labeled tabs. The green box in the screenshot shows the content of the data view labeled Columns. The type of viewer that is presenting the data in the screenshot is the grid viewer. Read more about all data viewers in the [Viewers](#) section.

Common to both the database objects tree and the object view are the SQL commands that are used to fetch the information from the database. The associated SQL is executed by DbVisualizer whenever a node in the tree is expanded (to expose any child objects) or when a node is selected (to fill the object data views).

Right-clicking the mouse on an object in the tree or clicking the Actions button in the object view shows a menu with all valid actions for the selected object. These are also defined per database profile and object type. Read more about the capabilities of actions in the [Definition of user actions](#) section.

How does DbVisualizer know what database profile to use?

DbVisualizer automatically load the appropriate database profile (XML file) based on the following:

1. The Database Type for the database connection is matched with the information in the DBVIS-HOME/resources/database-mappings.xml file to find out if there is a database profile available. If it finds one, it is used.
2. If there is no matching profile, the generic profile is used. This is very basic profile and shows only rudimentary information about the objects in the database. This is also the profile used in the DbVisualizer Free edition for all databases.

A specific database profile can be selected manually for a database connection. This is done in the database connection properties. Manually choosing a profile requires that the profile supports the actual database. If it doesn't, various errors will be reported once the database objects tree is explored. (Whenever the profile is changed, you must reconnect the database connection).

The name of the loaded profile is listed in the Connection tab status bar when the connection has been established. You can click the profile link to display the Database Profile list.



Figure: The status bar in the Connection tab when connected

XML structure

The mapping from the visual components in the user interface described earlier and the element definitions in the XML file is, briefly, as follows:

- The database objects tree (green box) is described by the ObjectsTreeDef root element. (The Database Connections node is mandatory and its appearance cannot be controlled by the profile).
- The object views (green and blue boxes) are described by the ObjectsViewDef root element.
- The commands used to execute the SQL to get the information for ObjectsTreeDef, ObjectsViewDef and optionally ObjectsActionDef definitions are defined by the Commands root element.
- All Actions for an object are defined in the ObjectsActionDef root element. (Actions are optional).

The XML for a database profile is quite simple, but there are a few things that need to be highlighted. All database connections loads a database profile from an XML file. If there is no matching database profile, the generic profile is used. This profile uses the standard JDBC metadata calls in order to obtain information about the structure and objects in the database. The generic profile is not one XML file, as the database specific profiles are, but instead four files:

- generic-commands.xml
- generic-actions.xml
- generic-tree.xml
- generic-view.xml

All these files are referred in the generic.xml file via include statements, i.e., each of the above files are included in the generic.xml file when it is loaded. The reason for this file organization is that the four files above can also be included and extended in a specialized profile. See later for more information.

The XML structure used to represent the database profile is as follows (click on the link to read more about each specific section):

- [Commands](#)
Defines the SQLs for the ObjectsTreeDef, ObjectsViewDef and optionally ObjectsActionDef.
- [ObjectsActionDef](#) (optional)
Defines actions for object types.
- [ObjectsTreeDef](#)
Defines the structure and what objects should be visible in the objects tree.
- [ObjectsViewDef](#)
Defines the object views for a specific object type.

XML skeleton

The following is a minimal database profile XML file, showing its structure.

```
<?xml version="1.0" encoding="UTF-8" ?>
  <!DOCTYPE DatabaseProfile SYSTEM "dbvis-defs.dtd" [
    <!ENTITY generic-commands SYSTEM "generic-commands.xml">
    <!ENTITY generic-view SYSTEM "generic-view.xml">
  ]>

  <DatabaseProfile desc="Profile for Sybase ASE"
    version="$Revision: 1.23 $"
    date="$Date: 2007/06/09 21:37:11 $"
    minver="5.0">

    <!-- ===== -->
    <!-- Definition of the commands -->
    <!-- ===== -->

    <Commands>
      &generic-commands;
      ...
    </Commands>

    <!-- ===== -->
    <!-- Definition of the object actions that are used by the tree -->
    <!-- ===== -->

    <ObjectsActionDef>
      ...
    </ObjectsActionDef>

    <!-- ===== -->
    <!-- Definition of the database objects tree structure -->
    <!-- ===== -->

    <ObjectsTreeDef id="sybase-ase">
      ...
    </ObjectsTreeDef>

    <!-- ===== -->
    <!-- Definition of the database objects views -->
    <!-- ===== -->

    <!-- Include the generic-view -->
    &generic-view;

    <ObjectsViewDef id="sybase-ase" extends="generic">
      ...
    </ObjectsViewDef>

  </DatabaseProfile>
```

Note: The name of the XML file (sybase-ase) and the values for the id attribute for the ObjectsTreeDef and ObjectsViewDef elements must be the same.

The first rows in the XML defines external dependencies and their URIs. The DOCTYPE identifier defines the DTD that is used to validate the XML. The ENTITY identifiers lists URIs for external references. In this case they identify the generic-commands.xml and generic-view.xml files. They can then be referenced in the XML as &generic-commands; and &generic-view;, which simply means that the related XML files are included in the final document when the profile is loaded.

The root of the database profile is the DatabaseProfile element. Continue to the next sections for information about the elements forming the database profile.

Tip: If you are using an XML editor to edit the profile it is very convenient to load the DTD in the editor, as you will then get color and error highlighting.

<DatabaseProfile>

The DatabaseProfile is the root element in the XML file. It is required and have the following attributes.

```
<DatabaseProfile desc="Profile for Sybase ASE"
  version="$Revision: 1.23 $"
  date="$Date: 2007/06/09 21:37:11 $"
  minver="5.0">
```

...

```
</DatabaseProfile>
```

The attributes specified for the DatabaseProfile element appear in the Database Profile list when selecting the connection properties for a database connection:

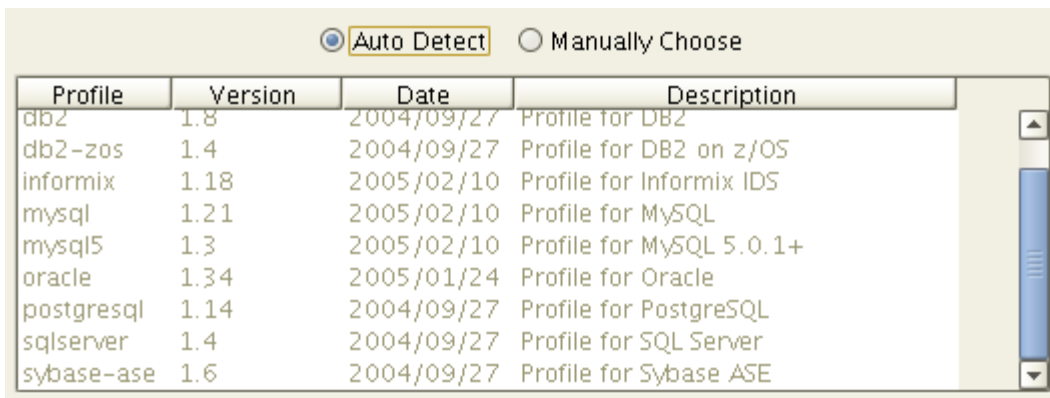


Figure: The list of available database profiles

<Commands> - The SQLs used to interact with the database

This element contains all Command elements with SQL sub element. A Command element is identified by a unique id attribute, which is then referred in ObjectsTreeDef, ObjectsViewDef and (optionally) ObjectsActionDef definitions.

```
<Commands>
```

```

&generic-commands;

<Command>
  ...
</Command>

</Commands>

```

The first statement in the <Commands> element is:

```
&generic-commands;
```

This means that the generic-commands entity defined at the top of the XML file is included in the XML i.e., all its definitions are accessible from the ObjectsTreeDef, ObjectsViewDef and ObjectsActionDef. If you don't plan to use any of the generic command, simply ignore this include statement.

<Command>

The Command element specifies the SQL associated with the command. In most cases, the SQL should return a result set with 0 or several rows. (The exception is actions which not necessarily need to return a result set, e.g., a "drop" action). The following command queries for login information in Sybase ASE.

```

<Command id="sybase-ase.getLogins">
  <SQL>
    <![CDATA[
select name, suid, dbname, fullname, language, totcpu,
totio, pwdate from master.dbo.syslogins
]]>
  </SQL>
</Command>

```

The id for this command is sybase-ase.getLogins. The reason for prefixing the id with the name of the profile is for maintainability. Since the generic-commands.xml file is included in most profiles, it is good to use unique prefixes for all commands so that they do not conflict with the commands in the generic-commands.xml file.

Result set

The result set for the previous query looks as follows:

name	suid	dbname	fullname	language	totcpu	totio	pwdate
sa	1	master	(null)	(null)	0	0	2005-02-24 23:59:14
probe	2	subsystemdb	(null)	(null)	0	0	2005-02-25 00:01:15

The way DbVisualizer handles the result set depends on whether the command is executed as a request in the database objects tree (ObjectsTreeDef) or in the object view (ObjectsViewDef). If executed in the database objects tree, each row in the result set will be represented by a new node in the tree. If executed in the object view, it is the viewer component that decides how the result will be presented. For more information on how a result set is used in the ObjectsTreeDef or ObjectsViewDef, read the specific sections.

Another important difference between the database objects tree and the object view is that the tree is a hierarchical structure of objects while the object view presents information about a specific object. An object that is inserted in the database objects tree is a 1..1 mapping to a row from the actual result set. The end user will see these objects (nodes) by some descriptive label, as defined in the ObjectsTreeDef. However, all data for the row from the original result set is stored with the object in the tree and may be used in the label, variables, conditions, etc. This is not the case in the ObjectsViewDef.

The following example put some light on this. Consider the previous result set and that it's used to create objects in the database objects tree. The end user will see the following in DbVisualizer. The visible name for each row is the name column in the result set.




Figure: Sample of the Logins node having two child nodes

Each of the sa and probe nodes have all their respective data from the result set associated with the nodes. The data is referenced as commandId.columnName, i.e., sybase-ase.getLogins.name, sybase-ase.getLogins.dbname, etc. All associated data for the sa node in the example is listed next:

```

sybase-ase.getLogins.name = sa
sybase-ase.getLogins.suid = 1
sybase-ase.getLogins.dbname = master
sybase-ase.getLogins.fullname = (null)
sybase-ase.getLogins.language = (null)
sybase-ase.getLogins.totcpu = 0
sybase-ase.getLogins.totio = 0
sybase-ase.getLogins.pwdate = 2005-02-24 23:59:14
  
```

The DataNode definition presenting sa and probe in the previous screenshot example use the associated data for the label as follows:

```
label="${sybase-ase.getLogins.name}"
```

<Input> - Setting command input

There are two types of Commands: with and without dynamic input. The difference is that dynamic input Commands accepts input data that is typically used to form the WHERE clause in SELECT SQLs. The previous example illustrates a static SQL (without dynamic data).

To allow for dynamic input, just add variables at the positions in the statement that should get dynamic values. The following is an extension of the previous example that allows for dynamic input.

```

<Command id="sybase-ase.getLogins">
  <SQL>
    <![CDATA[
select name, suid, dbname, fullname, language, totcpu,
totio, pwdate from master.dbo.syslogins
where name = '${name}' and suid = '${suid}'
    ]]>
  </SQL>
</Command>
  
```

The example above adds two input variables: \${name} and \${suid}. Values for these variables should then be supplied wherever the command is referred for execution via the Input element.

The following is an example from the ObjectsTreeDef and its use of the sybase-ase.getLogins command:

```

<GroupNode type="Logins" label="Logins">
  <DataNode type="Login" label="${sybase-ase.getLogins.Name} isLeaf="true">
    <SetVar name="objectname" value="${sybase-ase.getLogins.Name}">
    <Command idref="sybase-ase.getLogins">
      <Input name="name" value="sa">
      <Input name="suid" value="${sybase-ase.getProcesses.suid}">
    
```

```

    </Command>
  </DataNode>
</GroupNode>

```

(Note that the Command element refers the command via the idref attribute which will be matched with the corresponding id for the Command).

There is no magic with this definition, since the `$(name)` variable in the final SQL will be replaced with string "sa".

The value for the `$(suid)` definition will in this case get the value of the `sybase-ase.getProcesses.suid` when the SQL is executed. So where is this variable defined? As explained in the [Result Set](#) section, all the data for a row in the result set is associated with the objects in the database objects tree. In addition, it is possible to use all the data kept by the current object and all its parent objects (as presented in the objects tree) in the input to commands. So to evaluate the `$(sybase-ase.getProcesses.suid)` variable, DbVisualizer first looks for the variable in the current object. If it doesn't exist, it continues to look through the parent objects until it reaches the root, which is the Connections object in the objects tree. If the variable is not found, it will be set to the string representation for null, which is (null) by default. Whenever a matching variable is found, DbVisualizer uses its value and stops searching.

<Output> - Redefine command output

As mentioned earlier, a specific column value in a result set row is referenced by the name of the column prefixed by the command id. Sometimes this is not desirable and the Output definition can be used to change this behavior. The following identifies a column in the result set by its index number, starting from 1, and then force its name to be set to the value of the id attribute.

```

<Output>
  <Column id="sybase-ase.getLogins.Name" index="1">
  <Column id="sybase-ase.getLogins.suid" index="2">
</Output>

```

The Output element can also be used to alter the structure of columns in the result set by adding, renaming or removing columns.

```

<Output>
  <Column modelaction="add" index="THIS_IS_A_NEW_COLUMN" value="Rattle and Hum">
  <Column modelaction="rename" index="ADDR" name="ADDRESS">
  <Column modelaction="rename" index="2" name="PHONE">
  <Column modelaction="drop" index="MOBILE_PHONE">
  <Column modelaction="drop" index="4">
</Output>

```

(The rename and drop actions accepts either the name of the column or index number starting from the left at index 1).

- The add operation adds a new column to all rows. The value attribute accepts variables using the `$(...)` syntax.
- The rename operation simply renames a column.
- The drop operation drops the specified column.

The rename operation is primarily used when building a custom command that is supposed to be used by a viewer that requires predefined input by specific column names. Read more in the [ObjectsViewDef](#) section.

<ObjectsTreeDef> - Definition of the Database Objects Tree

The ObjectsTreeDef element section controls how the database objects tree should be presented and which

commands should be executed to form its content (nodes). The mapping between the graphical representation in DbVisualizer and its ObjectsTreeDef XML is as straight forward as it can be:

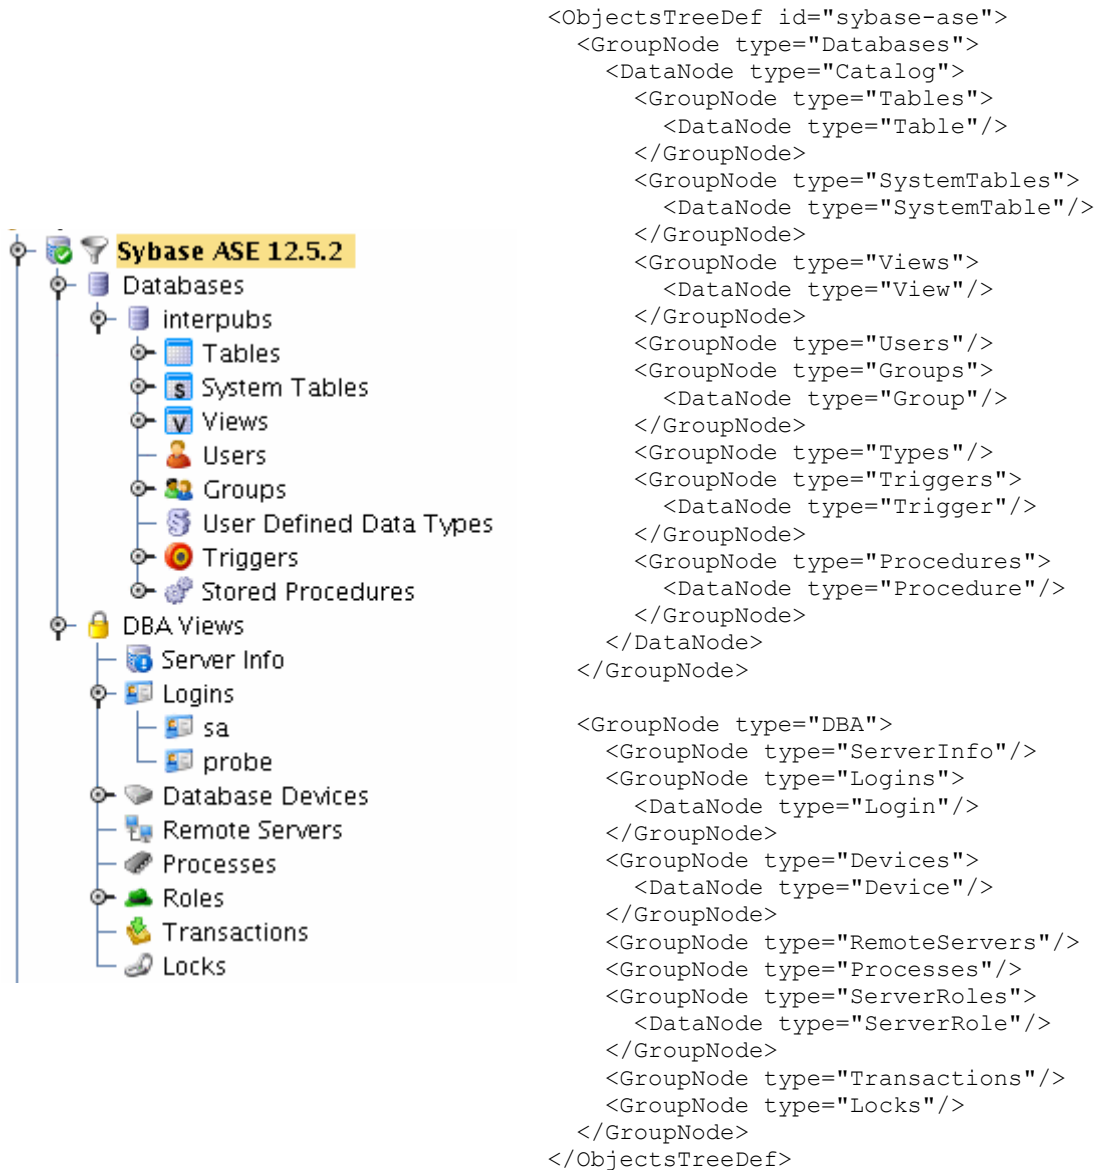


Figure: The visual database objects tree and its XML definition

The screenshot shows all nodes representing the GroupNode definitions in the ObjectsTreeDef. One exception is the Logins object, which has been expanded (sa and probe child objects) to illustrate what DataNode objects look like. The ObjectsTreeDef in the example has been simplified to show only the type attribute. (The label of the nodes as they appear in the visual tree is not listed in the ObjectsTreeDef example). The type attribute is primarily used internally in the profile as an identifier between the ObjectsTreeDef and the ObjectsViewDef. The type is also visible in the DbVisualizer GUI, in the tooltip for a tree node and in the object view header. The type is also used to identify the icon used to represent the object type.

There are no limitation on the number of levels in the ObjectsTreeDef. A good rule of thumb is, however, to keep it simple, clean and intuitive.

The DataNode definitions are the most important objects in the ObjectTreeDef. They also define which object tree filters are available for each object type, if overlaid icons should appear (and the criteria), etc. Read the next sections for details.

<GroupNode> - Static objects used for grouping

The GroupNode element represents a static object in the tree. These don't have any associated SQL and appear only once where they are defined. A GroupNode is primarily used for structural and grouping purposes. The GroupNode element has the following attributes.

```
<GroupNode type="SystemTables" label="System Tables" isLeaf="false">
  ...
</GroupNode>
```

The isLeaf attribute is optional and controls whether the GroupNode may have any child objects or not. It can always be set to true, but the effect in the visual database objects tree is then that the expand icon to the left of the group node icon will always be displayed, even if it can never have any child objects. The default setting for isLeaf is false.

Note: If isLeaf is set to false and there are child Group and/or Data -nodes, these will not appear. The result may cause some frustration during the design...

<DataNode> - Dynamic objects created via SQL

The DataNode element feeds the tree with nodes produced by a Command. The example in the [Command](#) section querying for all logins in Sybase ASE look as follow in the ObjectsTreeDef:

```
<GroupNode type="Logins" label="Logins">
  <DataNode type="Login" label="{sybase-ase.getLogins.Name}" isLeaf="true">
    <Command idref="sybase-ase.getLogins"/>
  </DataNode>
</GroupNode>
```

First, there is a GroupNode element with the purpose to group all child objects in a Logins node. The DataNode has, in this example, the same attributes as the GroupNode, the type is however "Login" instead of "Logins" (as it is for the GroupNode). This difference is important when the user selects one of the objects, since the Object View shows the appropriate views based on the object type.

The DataNode definition can be seen as a template, as the associated command fetches rows of data from the database and DbVisualizer uses the DataNode definition to create one node per row in the result set.

The label attribute for the data node is somewhat different, as it introduces the use of a variable (or several). The real value for the label will, in this example, be the value in the Name column produced by the sybase-ase.getLogins command, as you can see in the Command definition (variable names are automatically prefixed with the command id).

The Command element uses the idref attribute to identify the command that should be executed. The command in this case and in the [Result set](#) section produces a result set with 2 rows and 8 columns. The result will be two nodes each, with the label of the Name column in the result set.




Figure: Sample of the Logins node having two child nodes

The label can be changed by setting it to any other valid variable or a combination of several variables. (It's even possible to specify static text in the label):

```
label="{sybase-ase.getLogins.Name} ({sybase-ase.getLogins.dbname}) "
```

The example above results in the following labels:

```
sa (system)
probe (subsystemdb)
```

The complete set of attributes for the DataNode element is:

<code>type="value"</code>	- The type of node (required)
<code>actiontype="value"</code>	- Object type used for object actions (optional)
<code>label="value"</code>	- The visual label (required)
<code>isLeaf="true/false"</code>	- Specifies if the node can have child objects (default true)
<code>sort="col1,col2"</code>	- A comma separated list of names/variables used for sorting
<code>drop-label-not-equal="value"</code>	- Do not add the node if the label is not equal to this value or variable
<code>warnstate="condition"</code>	- If condition is true, show an overlay icon for the node
<code>errorstate="condition"</code>	- If condition is true, show an overlay icon for the node
<code>stop-label-hot-equal="value"</code>	- The node will be a leaf if the label doesn't match this value or variable
<code>is-empty-output="continue/stop"</code>	- If result set is empty, use this to control whether child GroupNode/DataNodes should be added anyway or ignored

The Command definition in the example above is simple, since it doesn't use any variables in the SQL. Continue reading the next section for details about passing input data to commands.

<Command>

Commands are referenced in the DataNode definition by the `idref` attribute. Sometimes it is required that a specific DataNode must supply input to a command. This is done by adding Input elements as children to the Command.

```
<DataNode type="Login" label="${sybase-ase.getLogins.Name}" isLeaf="true">
  <Command idref="sybase-ase.getLogins">
    <Input name="name" value="sa">
    <Input name="suid" value="${sybase-ase.getProcesses.suid}">
  </Command>
</DataNode>
```

The value for a variable specified in an Input element is evaluated using the strategy outlined in the [Result set](#) section.

<Filter>

The Filter element is specific for Command elements that appear in the ObjectsTreeDef section. A filter defines which data for a DataNode that is allowed to use in filters. This filter functionality is commonly referred to as the Database Objects Tree Filtering in DbVisualizer. The filtering setup appears below the database objects tree, and the following example shows that filtering may be specified for these object types:

- Catalog
- Table
- System Table
- View

For each of the Filter definitions, one or several columns can be as part of the filtering criteria.

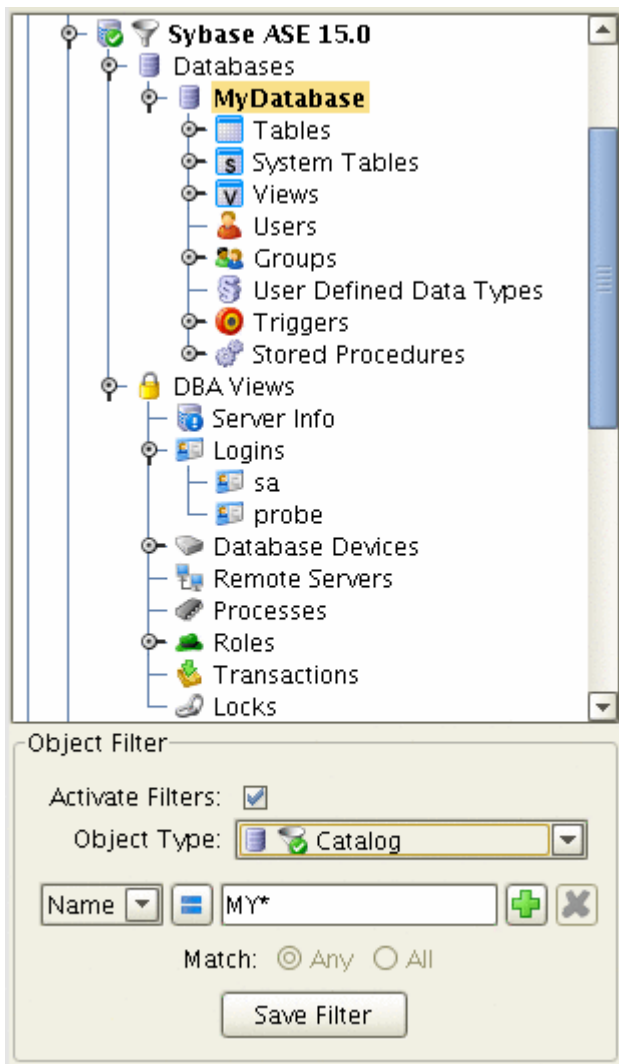


Figure: Screen shot showing the filter pane

```

<DataNode type="Views" label="{sybase-ase.getViews.Name}" isLeaf="true">
  <Command idref="sybase-ase.getViews">
    <Filter type="View" name="View Table">
      <Column index="TABLE_NAME" name="Name"/>
    </Filter>
  </Command>
</DataNode>

```

The previous filter definition specifies a filter for the View object type. The name specifies the name of the filter as it appears in the object type drop-down list. The nested Column element defines the index, which should be either a column name in the result set or an index number for the column. The name attribute specifies the name of the column as it appears in the filter pane.

Several Column elements may be specified for a Filter element.

<SetVar>

The SetVar element is needed in the ObjectsTreeDef for DataNode's. Some object types have special meaning in

DbVisualizer. Two examples are the Catalog and Schema object types. For DataNode objects, you must use SetVar elements to identify them, with name attributes set to "catalog" and "schema", respectively.

```
<DataNode type="Catalog" label="{getCatalogs.TABLE_CAT}" isLeaf="false">
  <SetVar name="catalog" value="{getCatalogs.TABLE_CAT}">
</DataNode>
```

All non Catalog or Schema DataNode's must use SetVar to set the "objectname" variable:

```
<DataNode type="Views" label="{sybase-ase.getViews.Name}" isLeaf="true">
  <SetVar name="objectname" value="{sybase-ase.getViews.Name}">
  <SetVar name="rowcount" value="true/false">
</DataNode>
```

The objectname variable is used to identify the object represented by the data node, so that it can be uniformly referenced in object views and object actions. Its value should be the identifier for the object as it is identified in the database, e.g., a table name or view name.

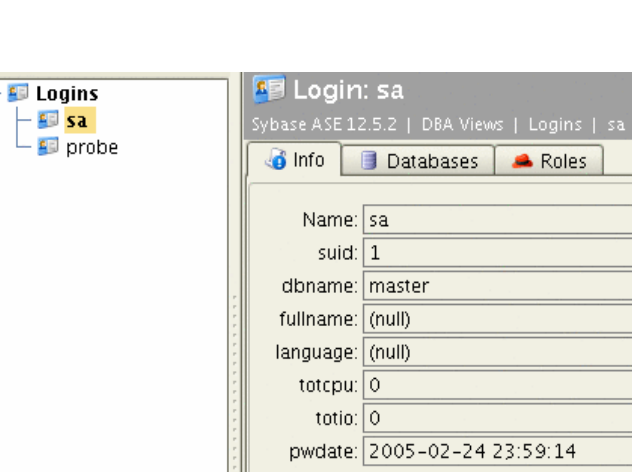
The rowcount variable is optional and controls whether the object supports getting row count information via the select count(*) SQL statement.

Another optional variable (not shown in the example above) is named acceptInQB. If set to true, nodes of this type can be used in the Query Builder. It should only be set to true for object types representing tabular data that can be queried with an SQL SELECT statement, such as tables, views, materialized views, etc.

<ObjectsViewDef> - Definition of the Object Views

The ObjectsViewDef element defines all views for the object types in the objects tree. These views are displayed in the Object View area for the selected object. Which views should appear when selecting a node in the tree is based on the object type for the tree node and the corresponding object view definition.

When an object is selected in the tree (sa in the screenshot below), its complete information is passed to the object view handler (right in the sample). This handler determines, based on the object type, which object view should be used to present the information. When the object view is found, all data views are created as tabs in the user interface. The selected object and its information is passed to each of the data views for processing and presentation. The following shows how the Object View look in DbVisualizer and its accompanying ObjectView definitions.



```
<ObjectView type="Logins">
  <DataView type="Logins" label="Logins"
    viewer="grid">
    <Command idref="sybase-ase.getLogins"/>
  </DataView>
</ObjectView>

<ObjectView type="Login">
  <DataView type="Info" label="Info"
    viewer="node-form"/>
  <DataView type="Databases" label="Databases"
    viewer="grid">
    <Command idref="sybase-ase.getLoginDatabases"/>
  </DataView>
  <DataView type="Roles" label="Roles"
    viewer="grid">
    <Command idref="sybase-ase.getLoginRoles"/>
  </DataView>
</ObjectView>
```

Figure: The visual database objects tree, object view and the XML definition

The screenshot shows both the Logins node and its child nodes, sa and probe. From the GroupNode and DataNode declaration examples in the previous sections, we know that these nodes are instances of the object types Logins (the Login node) and Login (the two sub nodes, sa and probe).

The ObjectView XML definitions shows the data views for these two types, Logins and Login. Clicking on the node labeled Logins in the tree will show the object view for the <ObjectView type="Logins"> definition while clicking on the node labeled sa or probe will show the object view for the <ObjectView type="Login"> .

The example shows sa being selected. Its DataView definitions are (by label):

- Info
- Databases
- Roles

These views are presented in DbVisualizer as tabs. The label of each tab is the label defined in the DataView and the icons are defined by the respective object type.

The ObjectsViewDef root element has the following attributes:

```
<!-- Include the generic-view -->
&generic-view;

<ObjectsViewDef id="sybase-ase" extends="generic" >
  ...
</ObjectsViewDef>
```

The first statement for the ObjectsViewDef elements is:

```
&generic-view;
```

This simply means that the generic-view entity defined at the top of the XML file is included in the XML, i.e., all its definitions are accessible as is. One example is the ObjectView definition in the generic-view.xml file for the Table object type. It contains a lot of DataView elements that identify all viewers for the Table. If you now want to use the generic Table DataView's but add a new Abbreviations data view, then simply extend the generic Table DataView. This is done by adding a extends="generic" attribute in the ObjectsViewDef element. By using the exact same object type in the extended ObjectView, you will then get this behavior. Read more about extending ObjectView's in the [Extending ObjectView](#) section.

<ObjectView>

The ObjectView element is associated with an object type and groups all DataView elements that appear when the object type is selected in the database objects tree. Here follows the ObjectView definition for the Login object type.

```
<ObjectView type="Login">
  ...
</ObjectView>
```

This element is simple as its only attribute is the type attribute. The type attribute value is used when a node is clicked in the database objects tree to map the object of the type clicked and its ObjectView.

<DataView>

The DataView element is as important as the DataNode is in the ObjectsTreeDef. It defines how the viewer should be labeled in DbVisualizer, which viewer (presentation form) it should use, commands and other things. The following is the DataView definitions for the Login object type. (The ObjectView element is part of the sample just for clarification).

```

<ObjectView type="Login">
  <DataView type="Info" label="Info" viewer="node-form"/>
  <DataView type="Databases" label="Databases" viewer="grid">
    <Command idref="sybase-ase.getLoginDatabases"/>
  </DataView>
  <DataView type="Roles" label="Roles" viewer="grid">
    <Command idref="sybase-ase.getLoginRoles"/>
  </DataView>
</ObjectView>

```

The elements are used to define how the object is presented in DbVisualizer, as described in the [introduction](#) of the ObjectsViewDef section. All three data view elements have a viewer attribute, which identifies how the data in the view should be presented, e.g., as a grid or a form. See the next section for a list of viewers.

Viewers

The viewer attribute for a DataView specifies how the data for the view should be presented. The following sections walk through the supported viewers.

The following sample illustrates the viewer attribute.

```

<ObjectView type="Login">
  <DataView type="Info" label="Info" viewer="node-form"/>
</ObjectView>

```

DataView definitions may be nested and the viewers are then presented with the nested DataView in the lower part of the screen.

grid

The grid viewer presents a result set in a grid, with standard grid features such as search, copy, fit, export, etc. The result set is presented exactly as it is produced by the Command and any optional Output processing.

Here is a sample of the XML for the grid viewer:

```

<DataView type="Columns" label="Columns" viewer="grid">
  <Command idref="oracle.getColumns">
    <Input name="owner" value="{schema}"/>
    <Input name="table" value="{objectname}"/>
  </Command>
</DataView>

```

And here is a screenshot of the Columns tab created from the previous definition.

Table: EMP					Actions
Oracle 10g: SCOTT Schemas SCOTT Tables EMP					
Row Id	References	Columns Comment	Constraints	Triggers	DDL
Info	Columns	Data	Row Count	Primary Key	Indexes Grants
OWNER	TABLE_NAME	COLUMN_NAME	DATA_TYPE	DATA_TYPE_MOD	
SCOTT	EMP	EMPNO	NUMBER	(null)	(nl)
SCOTT	EMP	ENAME	VARCHAR2	(null)	(nl)
SCOTT	EMP	JOB	VARCHAR2	(null)	(nl)
SCOTT	EMP	MGR	NUMBER	(null)	(nl)
SCOTT	EMP	HIREDATE	DATE	(null)	(nl)
SCOTT	EMP	SAL	NUMBER	(null)	(nl)
SCOTT	EMP	COMM	NUMBER	(null)	(nl)
SCOTT	EMP	DEPTNO	NUMBER	(null)	(nl)

0.060 sec/0.115 sec | 8 / 31 | 1-

Figure: The grid viewer

The nesting capability for grid viewers is really powerful, as it can be used to create a drill-down view of the data. Consider the scenario with a grid viewer showing all Trigger objects. Wouldn't it be nice to offer the user the capability to display the trigger source when selecting a row in the list? This is easily accomplished with the following:

```
<DataView type="Trigger" label="Triggers" viewer="grid">
  <Command idref="oracle.getTriggers">
    <Input name="owner" value="{schema}"/>
    <Input name="table" value="{objectname}"/>
  </Command>
  <DataView type="Source" label="Source" viewer="text">
    <Input name="dataColumn" value="text"/>
    <Input name="formatSQL" value="true"/>
    <Command idref="oracle.getTriggerSource">
      <Input name="owner" value="{OWNER}"/>
      <Input name="name" value="{TRIGGER_NAME}"/>
    </Command>
  </DataView>
  <DataView type="Info" label="Info" viewer="node-form"/>
</DataView>
```

- The first DataView definition defines the top grid viewer and the command to get the result set for it.
- The next DataView is the nested text viewer, specifying various input parameter for the viewer along with the command to get the source for the trigger. The difference here is that the input parameters for this command reference column names in the top grid. Since this viewer is nested, it will automatically be notified whenever an entry in the top grid is selected.
- The third nested DataView is presented as a tab next to the Source viewer, and presents additional

information about the selected trigger.

The following screenshot illustrates the above sample:

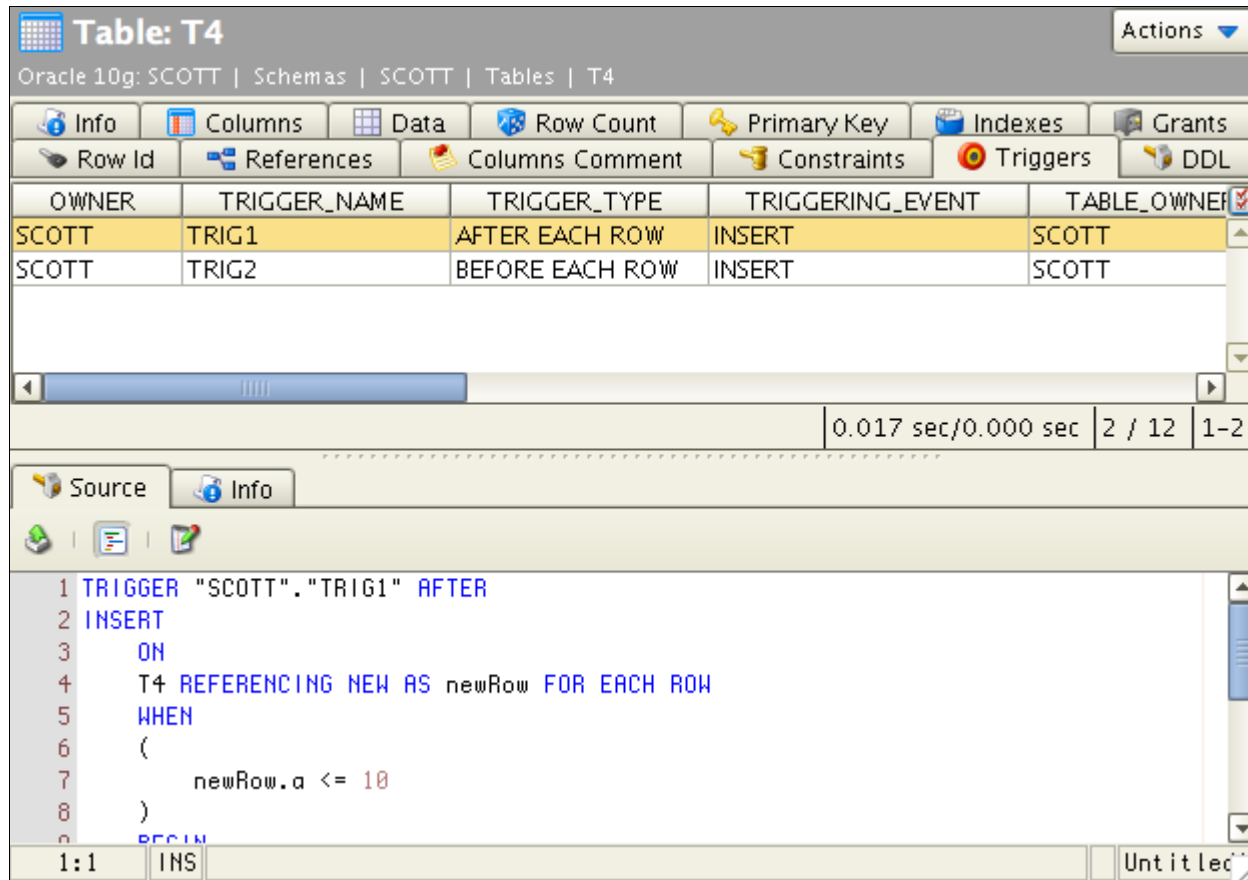


Figure: Example use of nested DataViews

Adding custom menu items in the grid

The `menutem` parameter specifies entries that should appear in the right-click menu in the grid. The value for the `menutem` is the label for the item, while the `child` input element specifies the SQL command that should be produced for all selected rows when the menu item is selected. The result of a custom menu item is that the grid viewer creates a statement that it copies to the SQL Commander; it will never execute the produced SQL in the scope of the viewer.

The following is an example with two menu items:

- Script: SELECT ALL
- Script: DROP TABLE

The variables in the SQL statement should identify column names in the result set. The user may select any columns in the visual grid and choose a custom menu item. It is only the actual rows that are picked from the selection as the columns are predefined by the `menutem` declaration.

The variables specified in these examples starts with `${schema=...}` and `${object=...}`. These defines that the first variable represents a schema variable while the second defines an object. This is needed for DbVisualizer to determine whether delimited identifiers should be used and if identifiers should be qualified, as defined in the connection properties for the database.

```

<Input name="menuItem" value="Script: SELECT ALL">
  <Input name="command" value="select * from ${schema=OWNER}${object=TABLE_NAME}"/>
</Input>

<Input name="menuItem" value="Script: DROP TABLE">
  <Input name="command" value="drop table ${schema=OWNER}${object=TABLE_NAME}"/>
</Input>

```

Here is a sample:

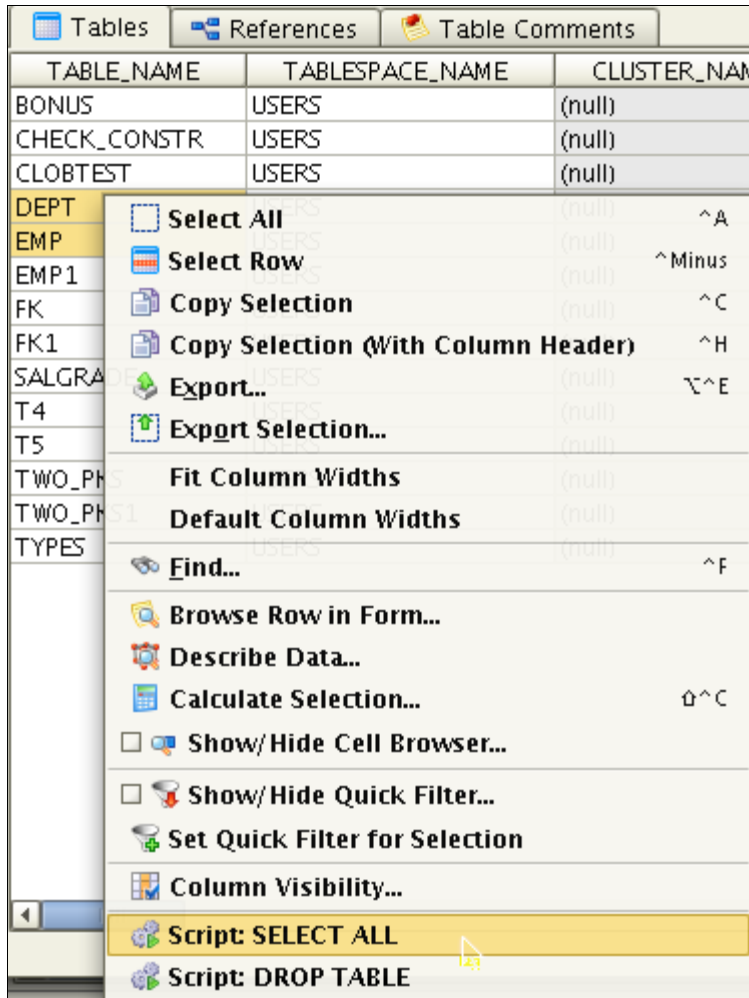


Figure: Custom menu items in grid viewer

Note: The result of selecting a menu item defined as a menuItem input parameter is that the specified command is copied to the current SQL editor.

Setting initial max column width

Some result sets may contain columns with very wide data. The following parameter sets an initial maximum column width for all columns in the grid.

```

<Input name="columnWidth" value="<pixels>"/>

```

text

The text viewer presents data from one column in a result set in a text browser (read only editor). This viewer is typically used to present large chunks of data, such as source code, SQL statements, etc. If the result set contains several rows, the text viewer reads the data in the column for each row and present the combined data.

Here is a sample of the XML for the text viewer:

```
<DataView type="Source" label="Source" viewer="text">
  <Input name="dataColumn" value="text"/>
  <Input name="formatSQL" value="true"/>
  <Command idref="oracle.getTriggerSource">
    <Input name="owner" value="{schema}"/>
    <Input name="name" value="{objectname}"/>
  </Command>
</DataView>
```

And here is a screenshot of the Source tab based on the previous definition.

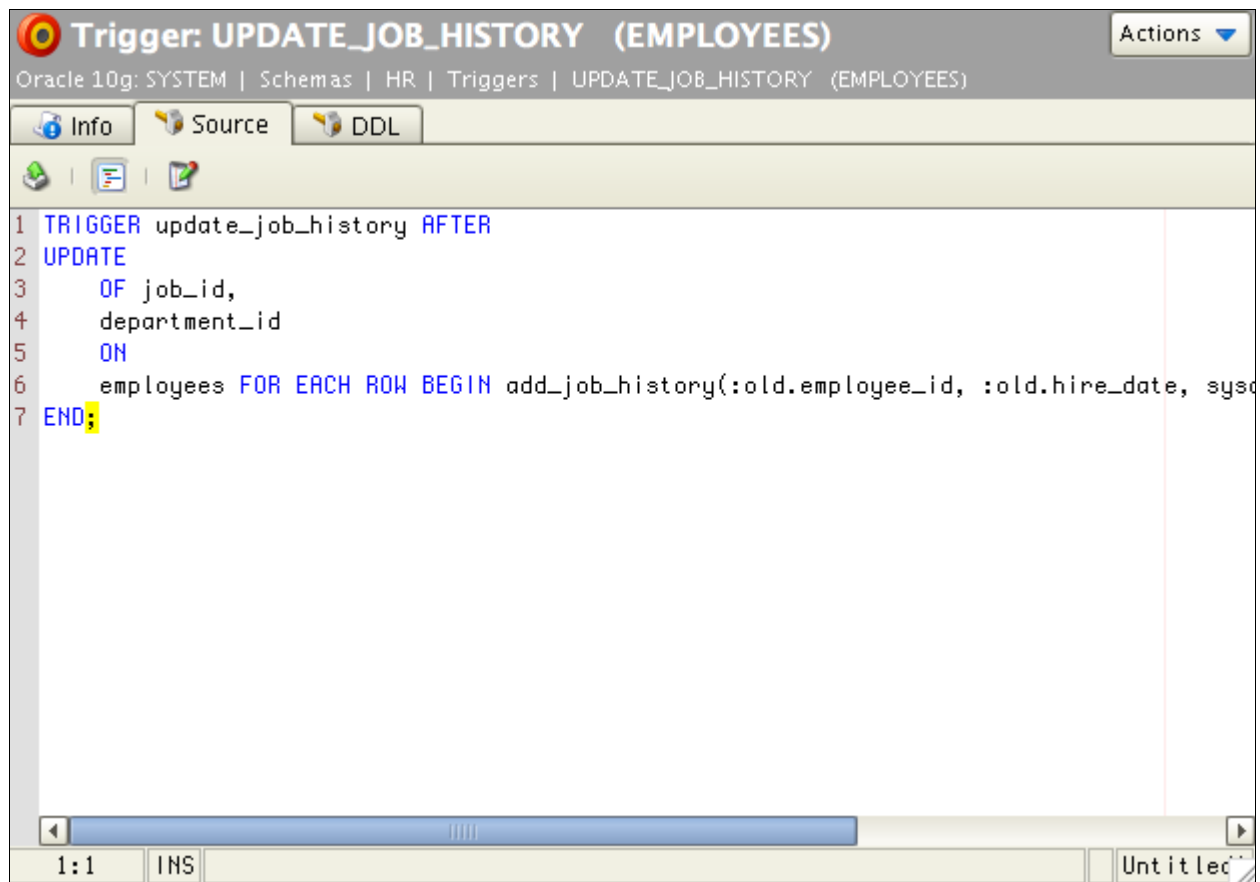


Figure: The text viewer

Specify what column to browse

By default, the text viewer uses the data in first column. This behavior can be controlled by using the dataColumn input parameter. Simply specify the name of the column in the result set or its index (starting at 1 from the left).

```
<Input name="dataColumn" value="<column-name>"/>
```

Enable SQL formatting of the data

The text viewer includes the SQL Formatting toolbar button, which when pressed formats the content in the viewer. The formatSQL input parameter is used to control whether formatting should be enabled by default. If formatSQL is not specified, no initial formatting is made.

```
<Input name="formatSQL" value="<true/false>"/>
```

form

The form viewer presents row(s) from a result set in a form. If several rows are in the result, they are presented in a list. Selecting one row from the list presents all columns and data for that row in a form.

Here is a sample of the XML for the form viewer:

```
<DataView type="Info" label="Info" viewer="form">  
  <Command idref="oracle.getTable">  
    <Input name="owner" value="{schema}"/>  
    <Input name="table" value="{objectname}"/>  
  </Command>  
</DataView>
```

And here is a screenshot of the Info tab based on the previous definition.

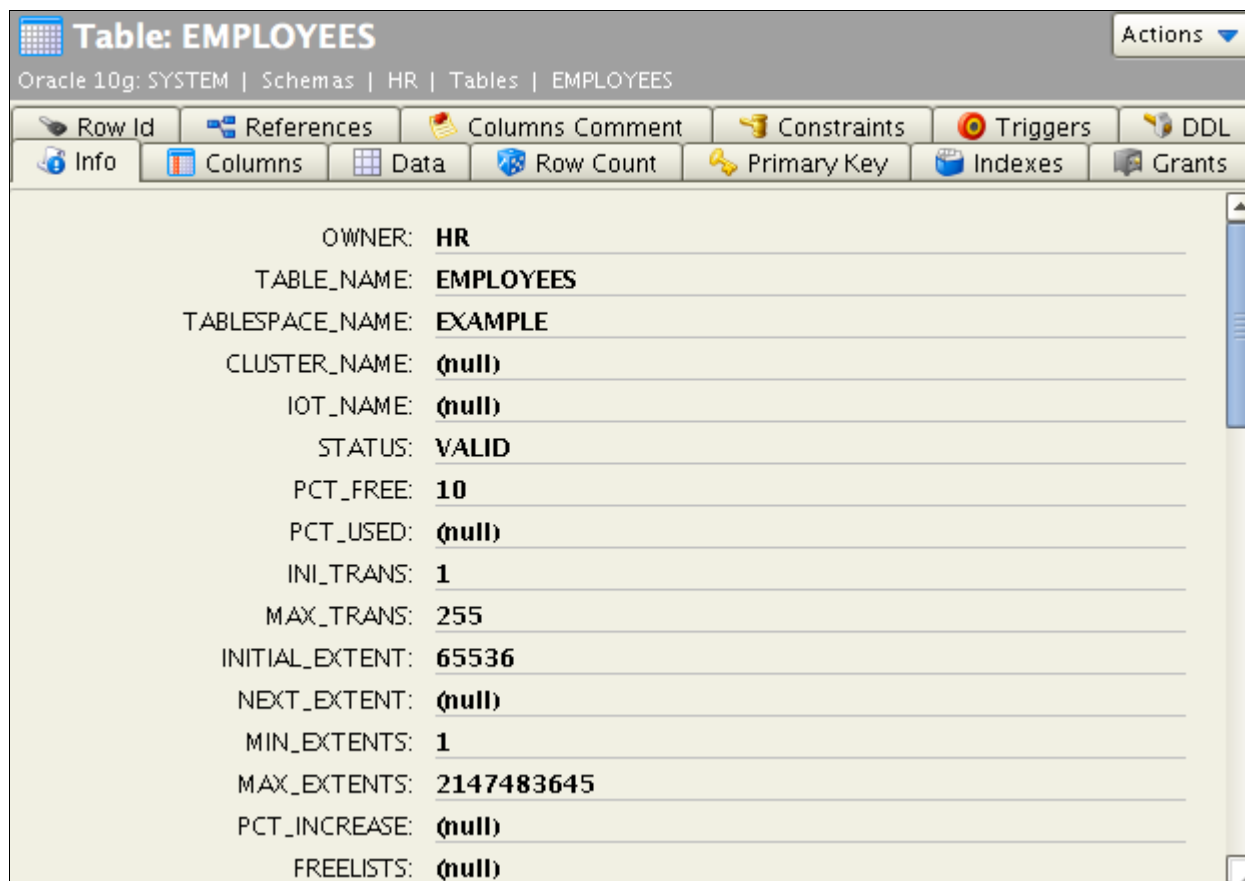


Figure: The form viewer

node-form

The node-form viewer presents all data associated with the selected object (variables).

Here is a sample of the XML for the node-form viewer:

```
<DataView type="Constraint" label="Constraint" viewer="node-form"/>
```

And here is a screenshot of the Constraint tab based on the previous definition.

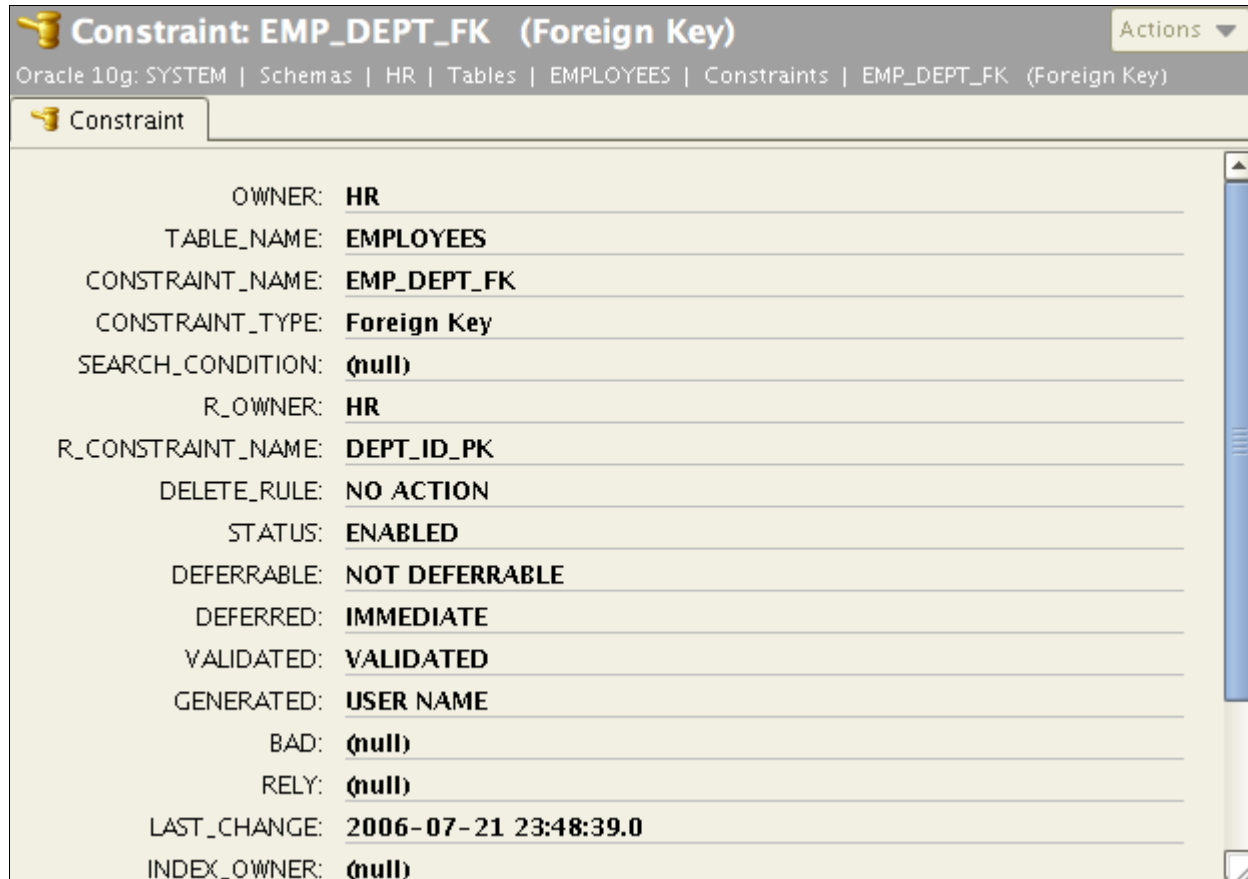


Figure: The node-form viewer

table-refs

The table-refs viewer shows the references graph for the current object (this must be an object supporting referential integrity constraints, such as a Table),

Here is a sample of the XML for the table-refs viewer:

```
<DataView type="References" label="References" viewer="table-refs"/>
```

And here is a screenshot of the References tab based on the previous definition.

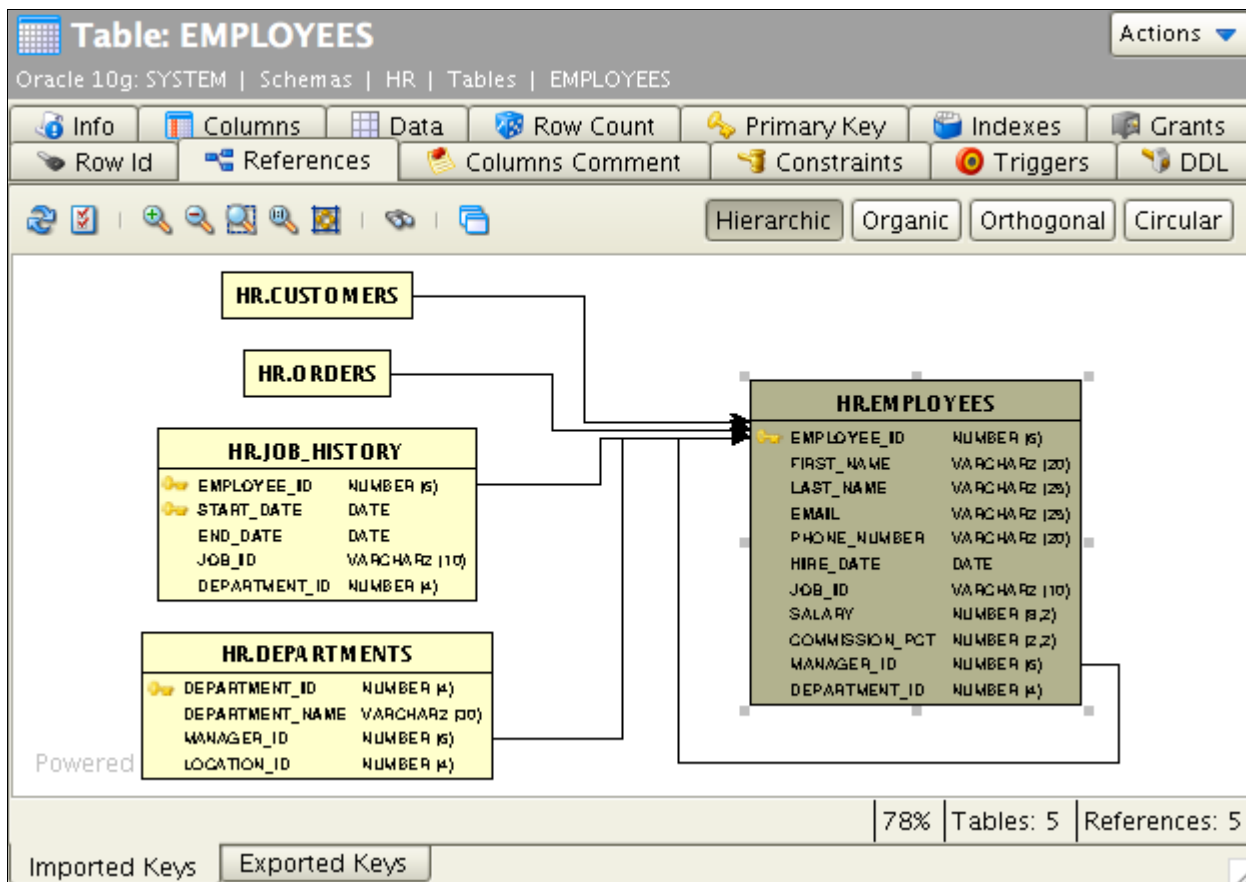


Figure: The table-refs viewer

tables-refs

The tables-refs viewer shows the references graph for several tables in the result set (the result set must contain objects supporting referential integrity constraints, such as a Table).

Here is a sample of the XML for the tables-refs viewer:

```
<DataView type="References" label="References" viewer="tables-refs">
  <Command idref="getTables">
    <Input name="catalog" value="{catalog}"/>
    <Input name="schema" value="{schema}"/>
    <Input name="table" value="{objectname}"/>
    <Input name="type" value="{tableType}"/>
  </Command>
</DataView>
```

And here is a screenshot of the References tab based on the previous definition.

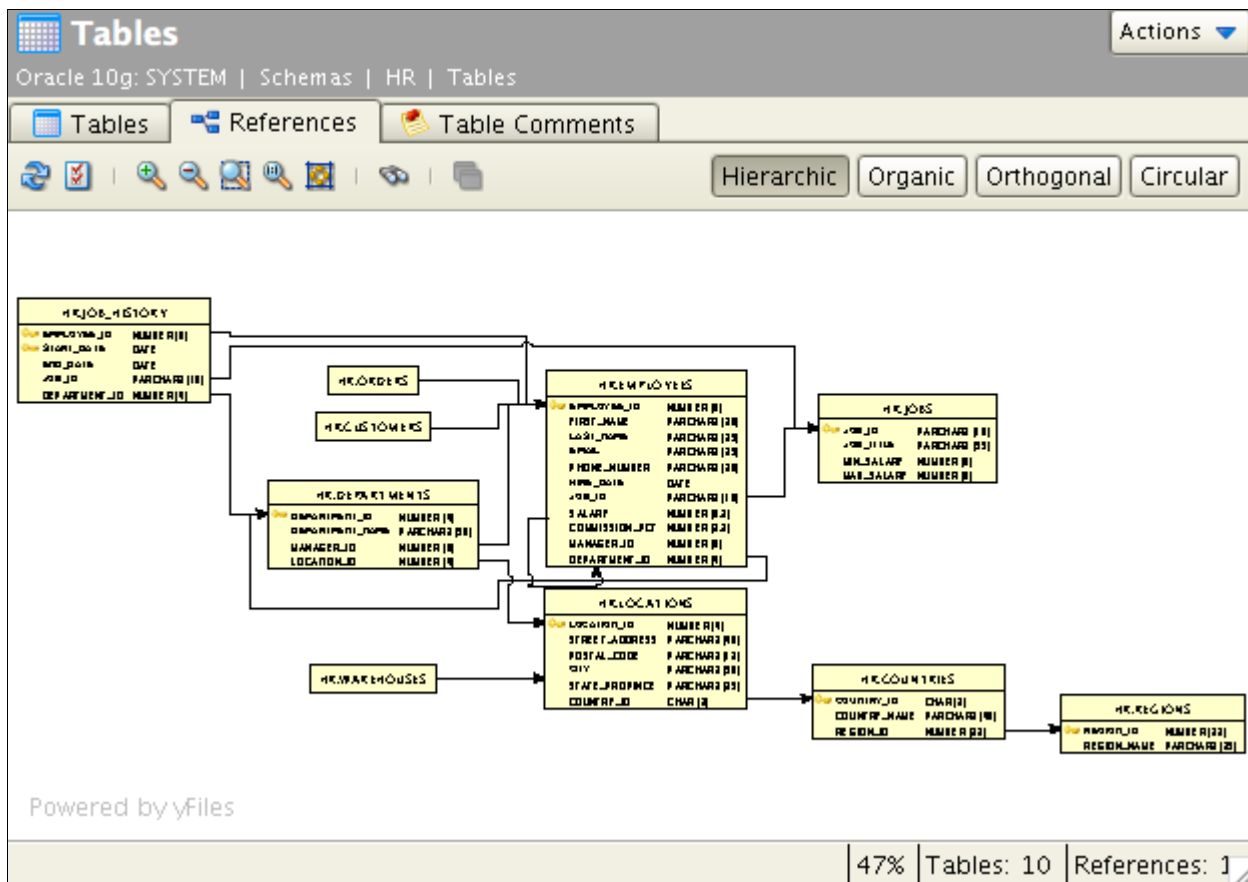


Figure: The tables-refs viewer

table-data

The table-data viewer shows the data for a table in a grid with editing features.

Note: information presented in the grid is obtained automatically by the viewer via a traditional `SELECT * FROM <schema>.table` statement, i.e., the object type having this viewer defined must be able to support getting a result set via this SQL statement.

Here is a sample of the XML for the table-data viewer:

```
<DataView type="Data" label="Data" viewer="table-data"/>
```

And here is a screenshot of the Data tab based on the previous definition.

Table: EMP								Actions
Oracle 10g: SCOTT Schemas SCOTT Tables EMP								
Row Id	References	Columns Comment	Constraints	Triggers	DDL			
Info	Columns	Data	Row Count	Primary Key	Indexes	Grants		
	EMPNO	ENAME	JOB	MGR	HIREDATE	SAL		
1	7369	SMITH	CLERK	7902	1980-12-17 00:00:00.0	800	(nul	
2	7499	ALLEN	SALESMAN	7698	1981-02-20 00:00:00.0	1600		
3	7521	WARD	SALESMAN	7698	1981-02-22 00:00:00.0	1250		
4	7566	JONES	MANAGER	7839	1981-04-02 00:00:00.0	2975		
5	7654	MARTIN	SALESMAN	7698	1981-09-28 00:00:00.0	1250		
6	7698	BLAKE	MANAGER	7839	1981-05-01 00:00:00.0	2850	(nul	
7	7782	CLARK	MANAGER	7839	1981-06-09 00:00:00.0	2450	(nul	
8	7788	SCOTT	ANALYST	7566	1987-04-19 00:00:00.0	3000	(nul	
9	7839	KING	PRESIDENT	(null)	1981-11-17 00:00:00.0	5000	(nul	
10	7844	TURNER	SALESMAN	7698	1981-09-08 00:00:00.0	1500		
11	7876	ADAMS	CLERK	7788	1987-05-23 00:00:00.0	1100	(nul	
12	7900	JAMES	CLERK	7698	1981-12-03 00:00:00.0	950	(nul	
13	7902	FORD	ANALYST	7566	1981-12-03 00:00:00.0	3000	(nul	
14	7934	MILLER	CLERK	7782	1982-01-23 00:00:00.0	1300	(nul	

Max Rows: 6000 Max Chars: -1 0.015 sec/0.006 sec 14 / 8 1-1

Figure: The table-data viewer

table-rowcount

The table-rowcount viewer shows the row count for a (table) object.

Note: The row count is obtained automatically by the viewer via a traditional `SELECT COUNT(*) FROM <schema>.table` statement, i.e., the object type having this viewer defined must be able to support getting a result set via this SQL statement.

Here is a sample of the XML for the table-rowcount viewer:

```
<DataView type="RowCount" label="Row Count" viewer="table-rowcount"/>
```

And here is a screenshot of the Row Count tab based on the previous definition.

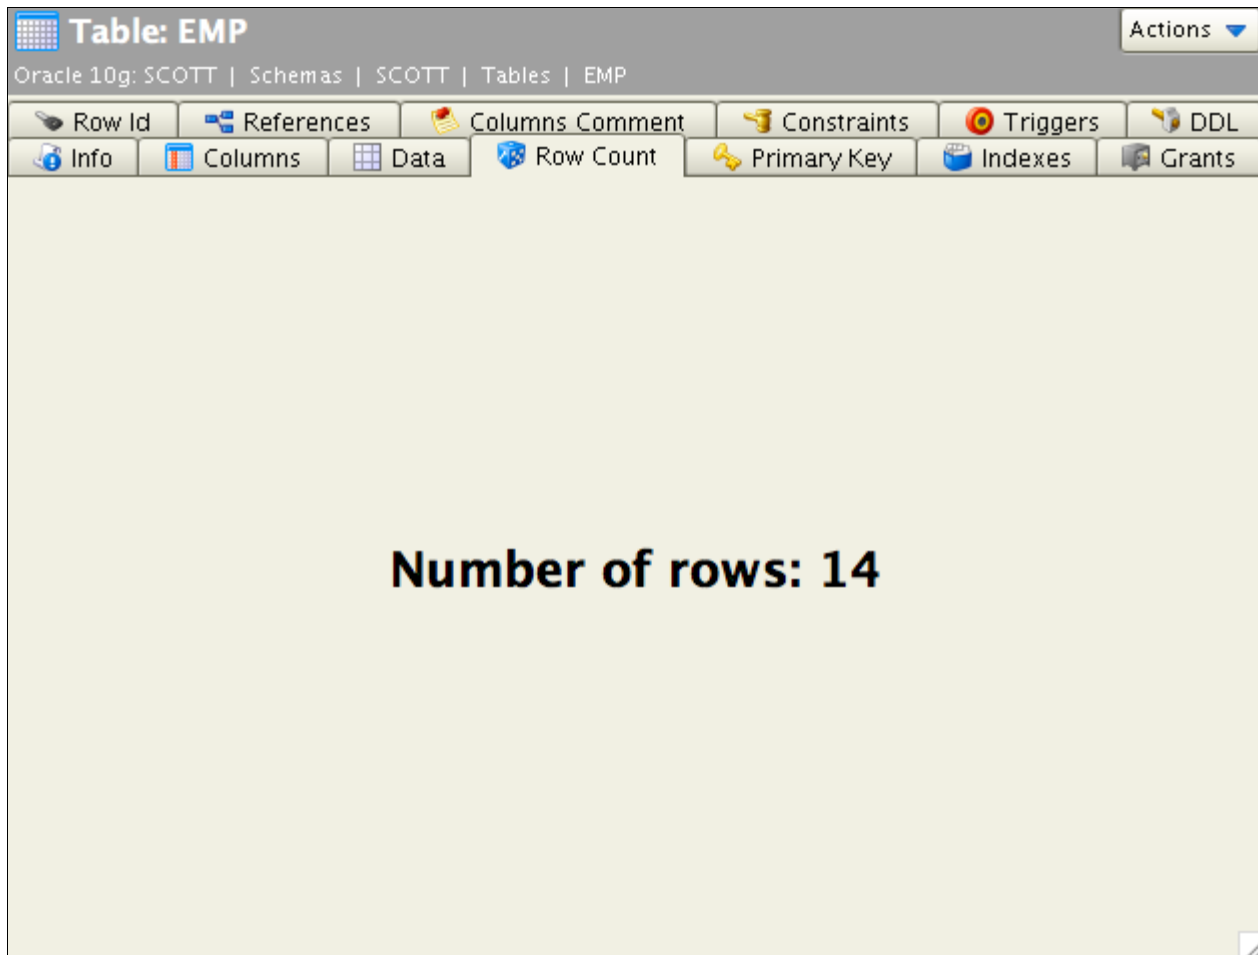


Figure: The table-rowcount viewer

<Command>

Please read the [Command](#) section above, as the capabilities of this element are the same when used with a data view.

<Message>

The Message element is very simple: it defines a message that should appear at the top of the viewer. The Message element is used to define the text for a description of the data presented in the viewer. The text in the message may contain common HTML tags such as `` (bold), `<i>` (italic), `
` (line break), etc.

Here is a sample of the XML for using the Message element in a grid viewer:

```
<ObjectView type="RecycleBin">
  <DataView type="RecycleBin" label="Recycle Bin" viewer="grid">
    <Command idref="oracle.getRecycleBin">
      <Input name="schema" value="{schema}"/>
      <Input name="login_schema" value="{dbvis-defaultCatalogOrSchema}"/>
    </Command>
    <Message>
```

```

<![CDATA[
<html>
These are the tables currently in the recycle bin for this schema. Right click on a bin
table in objects tree to restore or permanently purge it.<br>
<b>Note: The recycle bin is always empty if not looking at the bin for your
login schema (default).</b>
</html>
]]>
</Message>
</DataView>
</ObjectView>

```

And here is a screenshot of the Recycle Bin tab based on the previous definition.

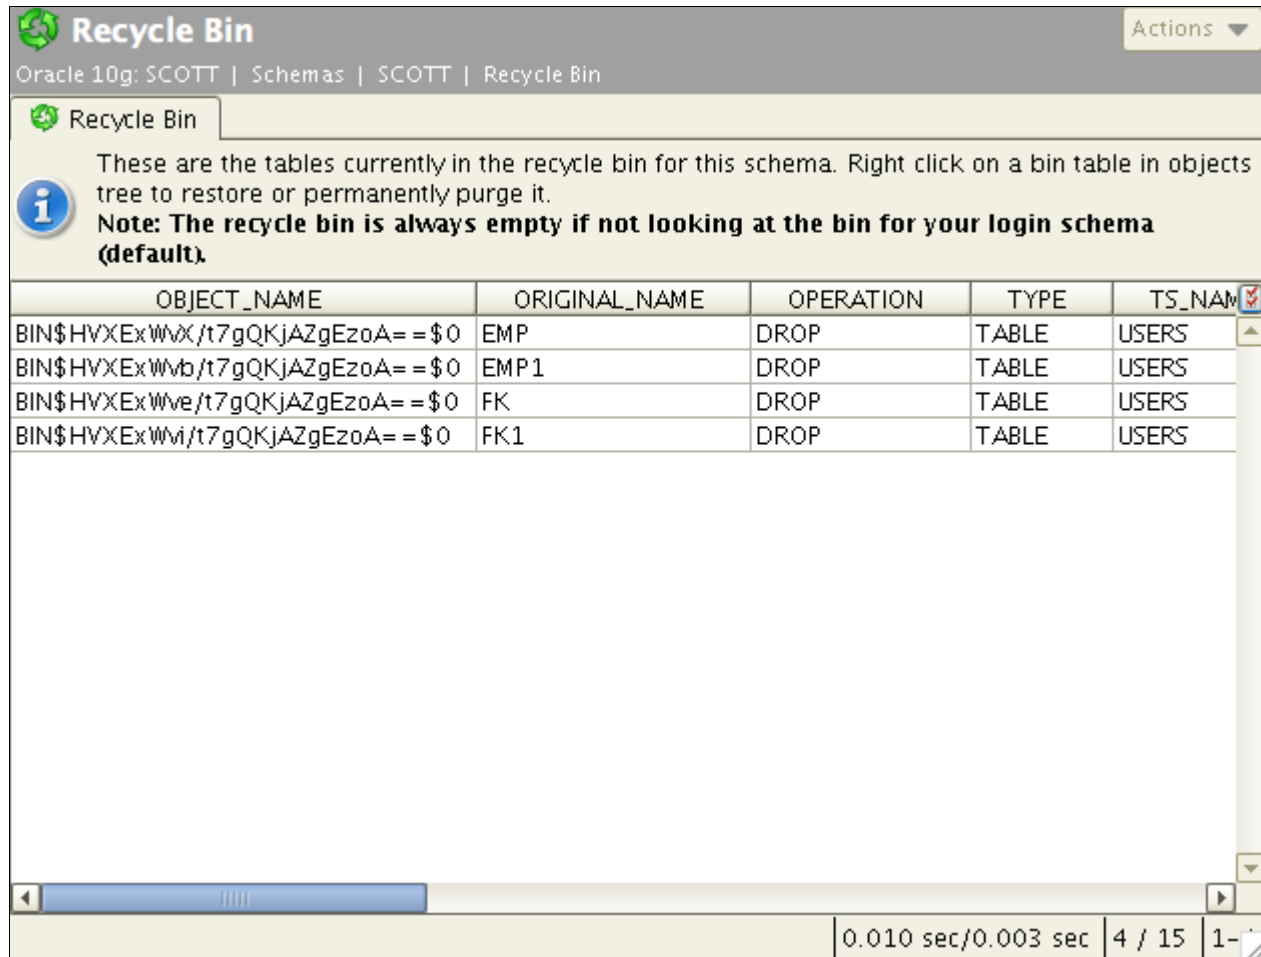


Figure: The appearance of a Message in a viewer

Extending ObjectView

An existing ObjectView definition in, for example, the generic-view.xml file can be extended in a database profile by using a few action attributes for each of the DataView elements. To extend a definition, the object type specified in the ObjectView type attribute must match the type in the parent profile. You have the following options when extending a definition:

- Adding a DataView

- Simply add the DataView definition and it will be added to the current list of DataView definitions
- Dropping an existing DataView
 - Add the `<DataView type="xxx" action="drop">` to drop the dataview type named "xxx"
- Replacing a DataView
 - Just add the DataView with the exact same type as in the parent DataView. All the settings of the new DataView will replace the old one

<ObjectsActionDef> - Definition of user actions

The previous sections describe how to define which objects should appear in the objects tree, and which views should be displayed when selecting an object in the tree. The ObjectsActionDef section in the profile defines which operations are available for the object types defined in the ObjectTreeDef. Object actions are very powerful, as they offer an extensive number of features to define actions for almost any type of object operation.

In DbVisualizer, the object type actions menu is accessed via the right-click menu in the objects tree or via the Actions button in the object view:

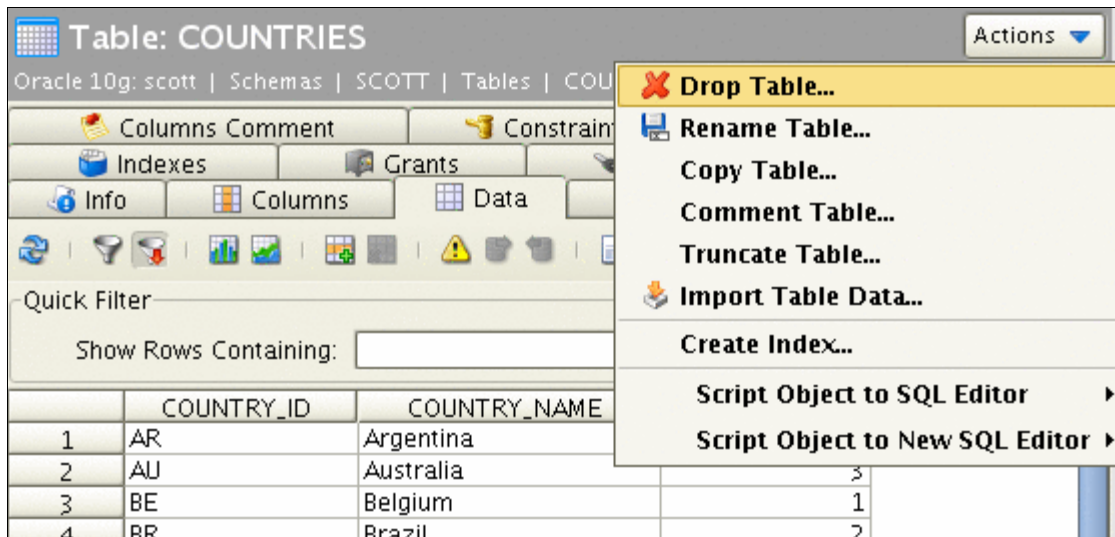


Figure: The Actions menu for the selected object

All of the operations for the selected Table object in the figure above are expressed in the ObjectsActionDef section. The implementation for these actions are either declared completely with XML elements via standard object actions, or via specialized action handlers. (The API for action handlers is not yet documented). The following screenshot shows the dialog appearing when executing an action via the default action handler:

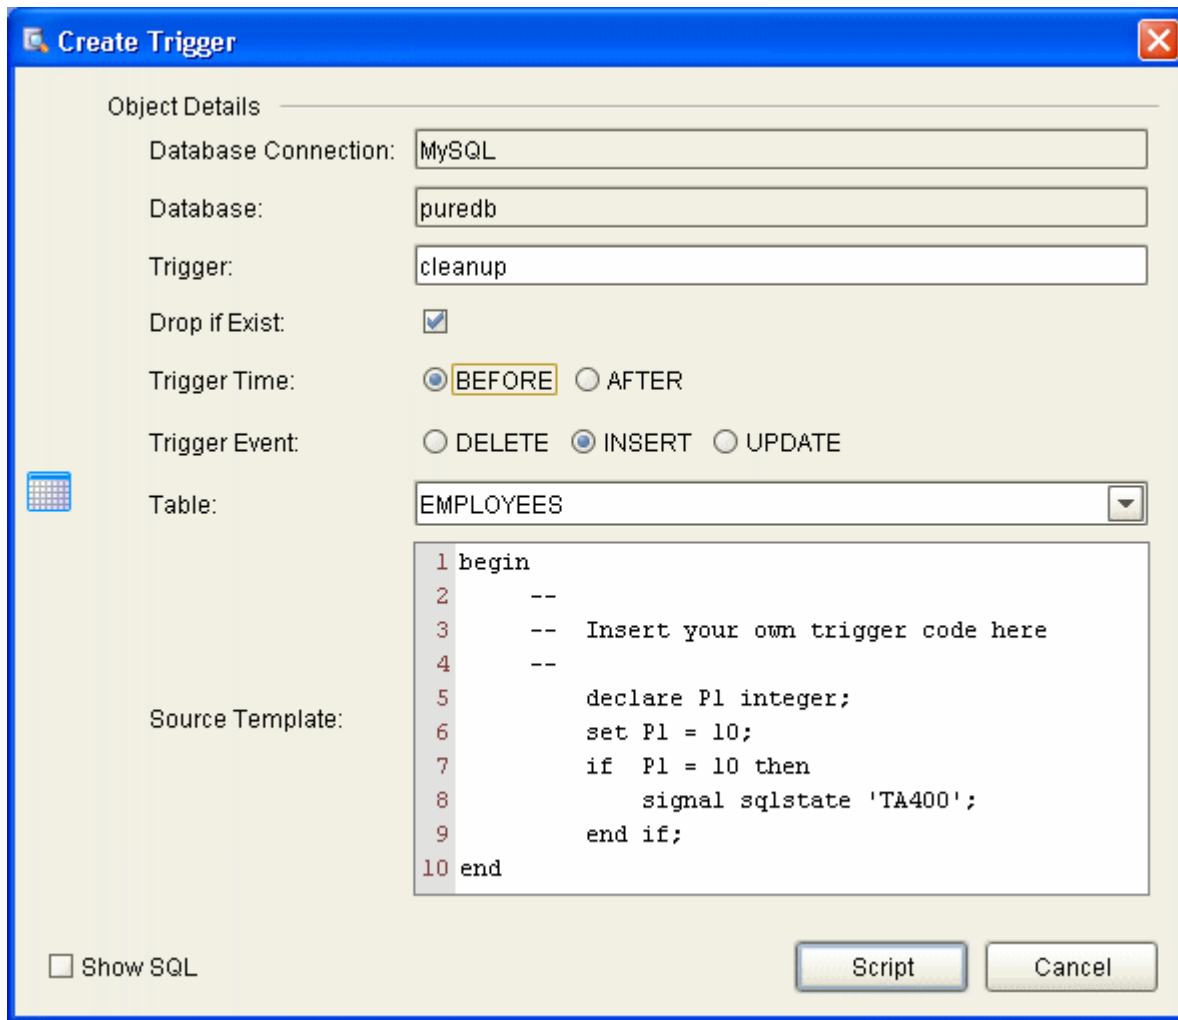


Figure: The default action handler

The first field in the dialog, Database Connection, is always present and shows the alias of the database connection the selected object is associated with. At the bottom, there is a Show SQL control that, when enabled, displays the final SQL for the action. The bottom right buttons are used to run the action (the label of the button may be Execute or Script based on the action mode), or to Cancel the action completely.

Variables

Variables are used to reference data for the object for which the action was launched, and the data for all its parent objects in the objects tree. Variables are also used to reference input data specified by the user in the actions dialog. Variables are typically used in the Command, Confirm, Result and SetVar elements.

Variables are specified in the following format:

```
#{variableName}
```

The following is an example for a Rename Table action. It first shows the name of the database connection (which is always present) along with information about the table being renamed. The last two input fields should be entered by the user and identify the new name of the table. The New Database control is a list from which the user should select the name of the new database. The new table name should be entered in the New Table Name field.

If the Show SQL control is enabled, you will see any edits in the dialog being reflected directly in the final SQL Preview.

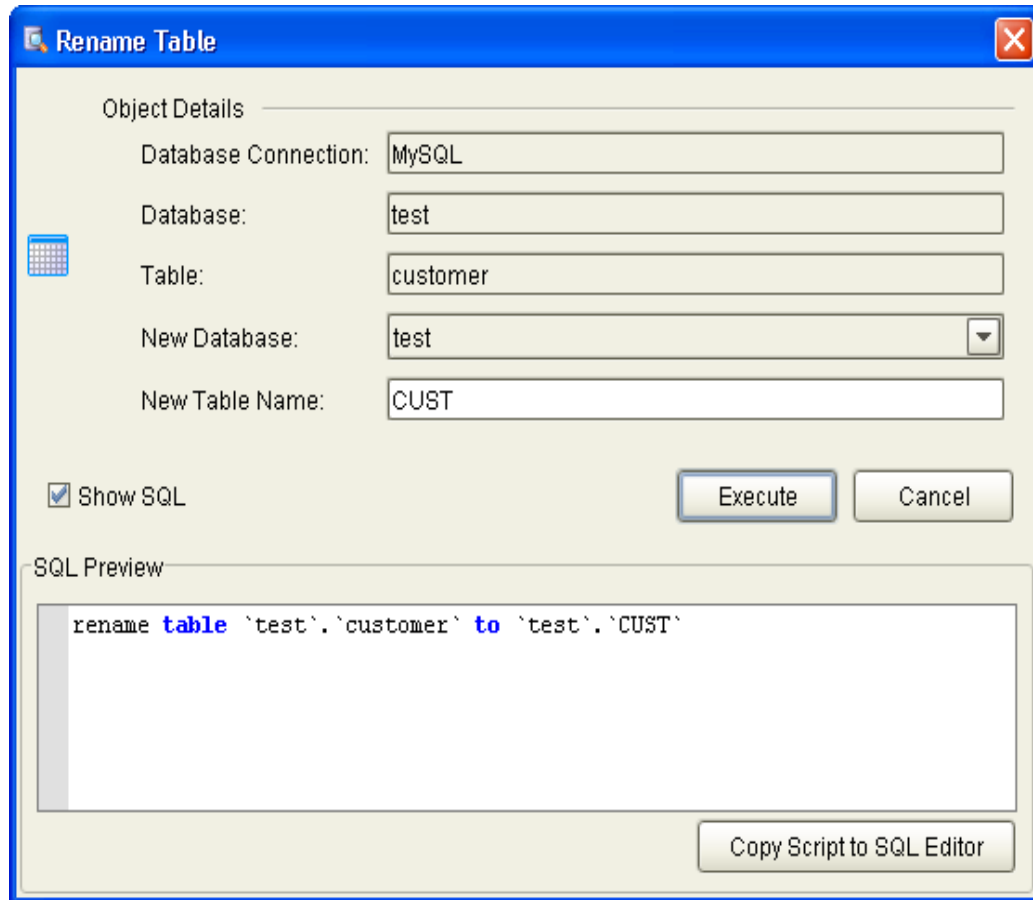


Figure: The default action handler

The complete action definition for the previous Rename Table action is as follows:

```
<Action id="mysql-table-rename" label="Rename Table" reload="true" icon="rename">
  <Input label="Database" style="text" editable="false">
    <Default>${catalog}</Default>
  </Input>

  <Input label="Table" style="text" editable="false">
    <Default>${objectname}</Default>
  </Input>

  <Input label="New Database" name="newCatalog" style="list">
    <Values>
      <Command><SQL><![CDATA[show databases]]></SQL></Command>
    </Values>
    <Default>${catalog}</Default>
  </Input>

  <Input label="New Table Name" name="newTable" style="text"/>

  <Command>
    <SQL>
```

```

        <![CDATA[
rename table `${catalog}`.`${objectname}`
to `${newCatalog}`.`${newTable}`
]]>
    </SQL>
</Command>

<Confirm>
    <![CDATA[
Confirm rename of ${catalog}.${objectname} to ${newCatalog}.${newTable}?
]]>
</Confirm>

<Result>
    <![CDATA[
Table ${catalog}.${objectname} renamed to ${newCatalog}.${newTable}!
]]>
</Result>
</Action>

```

First, there is the Action element with some attributes specifying the label of the action, icon and whether the objects tree (and the current object view) should be reloaded after the action has been executed.

The next block of elements are Input fields defining the data for the action. As you can see in the example, there is a `${catalog}` variable in the Default element for the Database input and an `${objectname}` variable in the Default element for the Tableinput. The values for these variables are fetched from the selected object in the objects tree. Variables are evaluated by first checking if the variable is in the scope of the action dialog (i.e., another input field), then if the variable is defined for the object for which the action was launched, and then if it is defined for any of the parent objects until the root object in the tree (Connections node) is reached. If a variable is not found, its value is set to (null).

In the previous XML sample, the value of the `${catalog}` variable is the name of the database in which the table object is stored. The `${objectname}` is the current name of the table (these variables are described in the ObjectsTreeDef section).

The New Database input field is a list component which shows a list of databases based on the result set of the specified SQL command. The Default setting for the database will be the database in which the table is currently stored based on the `${catalog}` variable.

The New Table Name input field is a simple text field in which the user may enter any text.

Both the New Database and New Table Name fields are editable and should be specified by the user. This data is then accessible via the variables specified in the name attribute, i.e., `newCatalog` and `newTable`.

The Commandelement declares the SQL statement that should be executed by the action. In this example, the SQL combines static text with variables.

<ActionGroup>

The ActionGroup element is a container and groups ActionGroup, Action and Separator elements. It is used to define what actions should be present for a particular object type. It also defines in what order the actions should appear in the menu and where any separators should be. ActionGroup elements can be nested to create sub menus.

```
<ActionGroup type="Table">
```

The attributes for an ActionGroup are:

- `type`
this defines what object type the ActionGroup represents. This attribute is valid only for top level action

groups. An example is the object of type Table, the corresponding ActionGroup will only be displayed when the selected object is a Table.

- label
this attribute is required for nested action groups. This label is displayed as the sub menu label for the nested action group. (The label attribute have no effect on top level action groups).

<Action>

The action element defines the action.

```
<Action      id="oracle-table-drop"
            icon="remove"
            label="Drop Table..."
            reload="true"
            mode="execute"
resultsetaction="ask">
```

The attributes for an action are:

- id
the id for the action. The recommended syntax for the id is "profileName-objectType-someGoodActionName"
- icon
specifies an optional icon that should be displayed next to the label in the menu
- label
the label for the action as it should appear in the menu in the action dialog
- reload
specifies if the parent node (in the objects tree) should be reloaded after successful execution. This is recommended for actions that change the visual appearance of the object, such as remove, add or name change
- mode attribute, can be set to any of these:
 - execute
(default) - show the action dialog, process user input and execute the final SQL within the scope of the action dialog
 - script
show the action dialog, process user input and send the final SQL to the SQL Commander
 - script-immediate
will not show the action dialog but instead pass the final SQL directly to the SQL Commander
- resultsetaction attribute, is only valid in combination with mode="execute". It can be set to any of:
 - ask
if the final SQL produced a result set, ask the user whether the result set should be displayed in a window or copied as text to the SQL Commander
 - show
if the final SQL produced a result set, show it in a window
 - script
if the final SQL produced a result set, copy it to the SQL Commander.
- class
used to launch a custom class. The execute attribute is obsolete if this attribute is set
- classargs
optional attribute used to specify arguments to the action handler defined by the class attribute

<Input>

An Input element specifies the characteristics of a visible field component for the actions dialog. The label attribute is recommended and is presented to the left of input field. If a label is not specified, the input field will occupy the complete width of the action dialog. All input fields are editable by default. The name attribute is required for editable fields and should specify the identity of the variable in which the user input is stored.

This is a minimal definition of an input field. It will show a read-only text field control labeled Size.

```
<Input label="Size" editable="false"/>
```

If the input field is changed to be editable, the name attribute must be used to specify the identifier for the variable name.

```
<Input label=Size" editable="true" name="theSize"/>
```

Any input element may contain the tip attribute. It is used to briefly document the purpose of the input field and is displayed as a tooltip when the user hovers the mouse pointer over it.

```
<Input label=Size" editable="true" name="theSize" tip="Please enter the size of the new xxx"/>
```

Specifying the default value as a result from an SQL statement is a trivial task:

```
<Input label=Size" editable="true" name="theSize">
  <Default>
    <Command>
      <SQL>
select size from systables where tablename = '${objectname}'
      </SQL>
    </Command>
  </Default>
</Input>
```

Since Default here will execute an SQL statement, it will automatically pick the value in the first row's first column and present it as the default. SQL may be specified in the Default and Values elements (as well as in the Labels element for list and radio styles). An alternative to embedding the SQL in the element body, as in the previous example, is to refer to a command via the standard idref attribute:

```
<Input label=Size" editable="true" name="theSize">
  <Default>
    <Command idref="getSize">
      <Input name="objectname" value="${objectname}"/>
    </Command>
  </Default>
</Input>
```

Instead of having duplicated SQLs in multiple actions, consider replacing these with Command elements referred via the idref attribute.

Referencing commands in actions via the idref attribute is recommended when the same SQL is used in several actions. Use Input elements to pass parameters to the command.

The following sections presents the supported styles that can be used in the Input element.

text (single line)

The text style is used to present single-line data in a text field.

```
<Input label="Enter your userid" name="userid" style="text">
  <Default>agneta</Default>
</Input>
```

- The optional Default element is used to define a default value for the field. Variables, static text and Command elements can be used to define the default value.
- A text input is editable by default. To make it read only just specify editable="false"

text-editor (multi line)

A text-editor field is the same as the text style except that it presents a multi-line field.

```
<Input label="Description" name="desc" style="text-editor" editable="true" args="height=50"/>
```

- The `args="height=50"` attribute defines the height (in DLU) for the text-editor. The default height is 30 DLU's.

number

A number style is the same as text except that it only accept number values.

```
<Input label="Size" name="size" style="number" editable="true"/>
```

password

A password field is the same as text except that it masks the value as "***".
Note that the password is visible in plain text in the SQL Preview.

```
<Input label="Password" name="pw" style="password" editable="true"/>
```

list (large number of choices)

The list style displays a list of choices in a drop-down component. The list can be editable, meaning that the field showing the selection may be editable by the user. Here is a sample XML for the list style.

```
<Input label="Select index type" name="type" style="list">  
  <Values>Pizza|Pasta|Burger</Values>  
  <Default>Pasta</Default>  
</Input>
```

The Values element should, for static entries, list all choices separated by a vertical bar (|) character. A Default value can either list the name of the default choice or the index number (first choice starts at 0). In the example above, setting Default to {2} would set Burger to the default selection.

It is also possible to use the Labels element. If present, this should list all choices as they will appear in the actions dialog. Consider these as being the labels shown to the user, while Values in this case should list the choices that will go into the final SQL via the variable. Here is an example:

```
<Input label="Select index type" name="type" style="list">  
  <Values>Pizza|Pasta|Burger</Values>  
  <Labels>Pizza the French style|Pasta Bolognese|Texas Burger</Labels>  
  <Default>Pasta</Default>  
</Input>
```

If the users selects Texas Burger then the value for variable type will be Burger.

The following shows how to use SQL to feed the list of values:

```
<Input label="New Database" name="newCatalog" style="list">  
  <Values>  
    <Command>  
      <SQL>  
        <![CDATA[
```

```

show databases
  ]]>
  </SQL>
</Command>
</Values>
<Default>${catalog}</Default>
</Input>

```

Here a Command element is specified as a sub element to Values. The result of the show databases SQL will be presented in the list component.

To make the list editable, specify the attribute `editable="true"`.

radio (limited number of choices)

The radio style displays a list of choices organized as button components. The only difference between the radio and list styles are:

- All choices for a radio style are displayed on the screen (better overview of choices but suitable only for a limited number of choices)
- The `args="vertical"` attribute can be specified for radio style to present the radio choices vertically

See the list style for complete capabilities of the radio style.

check (true/false, on/off, selected/unselected)

The check style is suitable for yes/no, true/false, here/there types of input. Its enabled state indicates that the Value for the input will be set in the final variable. If the check box is disabled, the variable value is blank

```

<Input label="Cascade Constraints" name="cascade" style="check">
  <Values>compact</Values>
</Input>

```

- This will create a check component with the label Cascade Constraints
- Enabling the check box will set the value of the variable identified by name (cascade) to the value of Value, which is compact.
- If the check box is unchecked, the variable value will be blank

separator (visual divider between input controls)

The separator style is not really an input element but is instead used to visually divide the fields in the in the actions dialog. If the label attribute is specified, it will be presented to the left of the separator line. If no label is specified, only the separator is displayed.

```

<Input label="Content" style="separator"/>

```

The separator is a useful substitute for the standard label presented to the left of every input field. Here is a sample:

Object Details

Database Connection: PostgreSQL 8.1.4

Schema: information_schema

Function: GenKeys

Options

Return Data Type: INTEGER

Language: SQL C INTERNAL

Volatility: VOLATILE STABLE IMMUTABLE

Returns Set:

Strict:

Security of Definer:

Parameters

Name	Data Type
id	integer
name	varchar

Source

```
select $1 + 100
```

Figure: Sample showing separators and wide fields

The previous figure shows the use of separators and two fields that extend to the full width of the action dialog. The separators for Parameters and Source are here used as alternatives to labels for the fields below them.

procedure-params

The procedure-params style is a specialized style used primarily to manage parameters for procedure and function object types.

```
<Input name="parameters" style="procedure-params"
  args="procedure,defaultType=VARCHAR(20),parmOrder=direction|name|type,directions=IN|INOUT|OUT/>
```

This example doesn't specify the label attribute as we want the field to extend the full width of the actions dialog. The procedure-params style use the args attribute extensively to customize the appearance and function of the field. The following content in the args attribute is handled by the procedure-params style:

- procedure
Defines what object type being handled: one of procedure or function
- defaultType=VARCHAR(20)

- Defines the default data type when the user adds a new row in the list
- parmOrder=directionInametype/default
Specifies the order of the parameters for each row as they will appear in the final variable. The identifiers above are static but can be ordered to comply with the desired order of the parameters
- directions=INIINOUTIOUT
This lists the available values for the Direction column (if it's present)

The resulting parameter list is created automatically by the control and is available in the variable name specified in the example to be parameters.

<SetVar>

The SetVar element is very powerful, as it is used to do conditional processing and create new variables based on the content of other variables.

Consider an SQL statement for creating new users in the database:

```
create user 'user' identified by 'password'
```

In this case it is quite easy to map the user field to an Input element for the action since it is a required field. The question arise for password which is optional. The identified by clause should only be part of the final SQL if the password is entered by the user. The solution for this scenario is to use the SetVar element. Here is the complete action definition:

```
<Action id="mydb-user-create" label="Create User" reload="true" icon="add">
  <Input label="Userid" name="userid" style="text"/>
  <Input label="Password" name="password" style="password"/>

  <SetVar name="_password" value='#password.equals("") ? "" : " identified by \"" + #password + "\"'/>

  <Command>
    <SQL>
      <![CDATA[
create user ${userid} ${_password}
      ]]>
    </SQL>
  </Command>
</Action>
```

The SetVar element accepts two attributes:

- name
should specify the name of the new variable
- value
this should contain the expression that will be evaluated. The expression is based on the OGNL toolkit provided by www.ognl.org. This is an expression library that mimics most of what is being supported by Java. Variables are referenced as #variableName.

The expression in the example above checks whether the password variable is empty. If it is empty, a blank value is being assigned to the _password variable. If it is not empty, the value for _password will be set to identified by "theEnteredPassword".

The SQL in the Command element now refer the new \$_password variable instead of the original #{password}.

Note: It is recommended that variables produced via SetVar elements are prefixed with an underline (_) to highlight were they come from.

<Confirm>

The Confirm element is displayed to the user when a request to Execute the action is made. If there are only read-only input fields in the action, this message is displayed in the body of the action dialog. The message is displayed in a confirmation dialog if there are editable fields.

Note that the message text can be composed of HTML tags such as , <i>,
, etc.

```
<Confirm>Really drop table ${table}?</Confirm>
```

<Result>

The Result element is optional and if specified, it is shown in a dialog after successful execution.

NOTE: Result elements are currently not displayed in DbVisualizer. It is however recommend that you specify these as they will most likely appear in some way or another in a future version. If you want to test the appearance of Result elements then open the DBVIS-HOME/resources/dbvis-custom.xml file in a text editor and make sure dbvis.showactionresult is set to true.

```
<Result>Table ${table} has been dropped!</Result>
```

- The Result message will be displayed in a dialog after successful execution.
- If the execution fails, a generic error dialog is displayed and the Result is not displayed.

<Command>

The Command element specifies the SQL code that is executed by the action.

```
<Command>
  <SQL>
    <![CDATA[
drop table ${table} mode ${mode} including constraints ${includeconstraints}
]]>
  </SQL>
</Command>
```

Conditional processing

Conditional processing means that a profile can adjust its content based on certain conditions. A few examples:

- Which version of the database it is
- The format of the database URL
- The client environment i.e Java versions, vendor, etc.
- User properties
- Database connection properties

Conditional processing is especially useful for adapting the profile for different versions of the database (and/or JDBC driver). Another use for the conditional processing is to replace generic error messages with more user friendly messages.

Programmers familiar with if, else if and else will easily learn the conditional elements.

Depending on in which of the two phases the conditions should be processed, some restrictions and rules apply. Please read the following sections for more information.

When are conditional expressions processed?

There are two phases when conditions are processed:

1. Conditional processing when database connection is established
<If>, <Elseif> and <Else> elements can be specified almost everywhere in the profile.
2. Conditional processing during command execution
The <OnError> element is used to define a message that will appear in DbVisualizer if a command fails.
Conditions are used to control what message should appear.

DbVisualizer uses the type attribute to determine which If elements should be executed in which phase. If this attribute has the value runtime, it will be processed in the second phase. If it is not specified or set to load, it will be processed in the first phase.

Conditional processing when database connection is established

The following example shows the use of conditions that are processed during connect of the database connection.

```
<Command id="sybase-ase.getLogins">
  <If test="#DatabaseMetaData.getDatabaseMajorVersion() lte 8">
    <SQL>
      <![CDATA[
select name from master.dbo.syslogins
      ]]>
    </SQL>
  </If>
  <ElseIf test="#DatabaseMetaData.getDatabaseMajorVersion() eq 9">
    <SQL>
      <![CDATA[
select name, suid from master.dbo.syslogins
      ]]>
    </SQL>
  </ElseIf>
  <Else>
    <SQL>
      <![CDATA[
select name, suid, dbname from master.dbo.syslogins
      ]]>
    </SQL>
  </Else>
</Command>
```

The above means that if the major version of the database being accessed is less than or equal to 8, the first SQL is used. If the version is equal to 9, the second SQL is used, and the last SQL is used for all other versions. The test attribute may contain conditions that are ANDed or ORed. Conditions can contain multiple evaluations, combined using parenthesis. The If, Elseif and Else elements may be placed anywhere in the XML file.

Here is another example that controls whether certain nodes will appear in the database objects tree or not.

```
<!-- Getting Table Engines was added in MySQL 4.1 -->
<If test="( #dm.getDatabaseMajorVersion() eq 4 and #dm.getDatabaseMinorVersion() gte 1)
  or #dm.getDatabaseMajorVersion() gte 5">
  <GroupNode type="TableEngines" label="Table Engines" isLeaf="true"/>

  <!-- "Errors" was added in MySQL 5 -->
  <If test="#dm.getDatabaseMajorVersion() gte 5">
    <GroupNode type="Errors" label="Errors" isLeaf="true"/>
```

```
</If>
</If>
```

As you can see, this example contains nested uses of If.

Conditional processing during command execution

Using conditional processing to evaluate any errors from a Command may be useful to rephrase error messages to be more user friendly.

```
<Commands>
  <OnError>
    <!-- The ORA-942 error means "the table or view doesn't exist" -->
    <!-- It is caught here since these errors typically indicates -->
    <!-- that the user don't have privileges to access the SYS and/or -->
    <!-- V$ tables. -->
    <If test="#result.getErrorCode() eq 942" context="runtime">
      <Message>
        <![CDATA[
You don't have the required privileges to view this object.
        ]]>
      </Message>
    </If>
    <ElseIf test="#result.getErrorCode() eq 17008" context="runtime">
      <Message>
        <![CDATA[
Your connection with the database server has been interrupted!
Please <a href="connect" action="connect">reconnect</a> to re-establish the connection.
        ]]>
      </Message>
    </ElseIf>
  </OnError>

  ...

</Commands>
```

The OnError element can be used in Commands and Command elements. If used in Commands element, its conditions are processed for all commands. If it is part of a specific Command, it is processed only for that command.

Current limitations

- The SQL statements in the profile must be statements that DbVisualizer can execute with JDBC. It can not contain any executables, scripts or OS specific calls
- It is not possible to specify conditions or compound commands, i.e., everything needed to execute a command must be expressed in a single SQL statement.

DB2 for Windows/Linux - Database Specific Support





















Introduction


In addition to the standard feature set in DbVisualizer this document briefly covers the database specific support in DbVisualizer for the DB2 on Windows/Linux RDBMS.

Database Objects Tree

The following table shows what objects DbVisualizer handles for a DB2 on Windows/Linux database. The first column list the structure and the object types. The actions column lists all actions that are associated with the object types and the last column shows all viewers that appear when selecting an object in the tree.

The screenshots illustrates the visual representation of the objects tree.

Object Type	Actions	Viewers
 Schemas	Create Schema	Schemas
 Schema	Comment Schema, Drop Schema	
 Tables	Create Table	References, Tables
 Table	Alter Table, Comment Table, Create Index, Drop Table, Empty Table, Import Table Data, Rename Table, Script: CREATE TABLE, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, Constraints, DDL, Data, Grants, Indexes, Info, Navigator, Primary Key, References, Row Count, Row Id, Triggers
 Views		Views
 View	Comment View, Drop View, Script: CREATE VIEW, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, DDL, Data, Info, Row Count
 Aliases	Create Alias	Aliases
 Alias	Comment Alias, Drop Alias, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, Data, Info, Row Count
 Nicknames		Nicknames
 Nickname	Comment Nickname, Drop Nickname, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, Data, Info, Row Count
 Indexes		Indexes
 Index	Comment Index, Drop Index, Rename Index	Columns, Index
 System Tables		SYSCAT, SYSFUN, SYSIBM, SYSSTAT
 Materialized Query Tables		Materialized Query Tables
 Materialized Query Table	Comment Materialized Query Table, Drop Materialized Query Table, Refresh Materialized Query Table, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, Data, Info, Row Count, Source
 Functions	Create Function	Functions
 Function	Drop Function	Info, Parameters, Result
 Procedures	Create Procedure	Stored Procedures
 Procedure	Drop Procedure	Info, Interface
 Package		Packages





Explain Plan

Explain Plan executes your query and records the plan that the database devises to execute it. By examining this plan, you can find out if the database is picking the right indexes and joining your tables in the most efficient manner. The explain plan feature works much the same as executing SQLs to present result sets, you may highlight statements, run a script or load from file. The explain plan results can easily be compared by using the pin feature in combination with window style presentation.

Informix IDS - Database Specific Support























Introduction

In addition to the standard feature set in DbVisualizer this document briefly covers the database specific support in DbVisualizer for the Informix IDS RDBMS.

Database Objects Tree

The following table shows what objects DbVisualizer handles for an Informix IDS database. The first column list the structure and the object types. The actions column lists all actions that are associated with the object types and the last column shows all viewers that appear when selecting an object in the tree.

The screenshots illustrates the visual representation of the objects tree.

Object Type	Actions	Viewers
 Database s	Create Database	Databases
 Catalog	Drop Database, Rename Database	
 Tables	Create Table	References, Tables
 Table	Alter Table, Create Index, Drop Table, Empty Table, Import Table Data, Rename Table, Script: CREATE TABLE, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, DDL, Data, Grants, Indexes, Info, Navigator, Primary Key, References, Row Count, Row Id, Triggers
 Column s		Columns
 Column		Column
 System Table s		References, System Tables
 System Table	Script: SELECT	Columns, Data, Info, Row Count
 Views		References, Views
 View	Drop View, Script: CREATE VIEW, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, DDL, Data, Indexes, Info, Row Count
 Synonym s		Synonyms
 Synonym		Info
 Index s		Indexes
 Index	Disable Index, Drop Index, Enable Index, Rename Index	Columns, Info
 Triggers		Triggers
 Trigger	Disable Trigger, Drop Trigger, Enable Trigger	Info, Source
 Procedure s		Procedures
 Procedure	Drop Procedure	Columns, Info, Source
 Functions 1		Functions
 Function	Drop Function	Columns, Info, Source
 Sequence s ¹	Create Sequence	Sequences
 Sequence	Drop Sequence, Rename Sequence	Info
 Data Types 1		Collection Types, Distinct Types, Opaque Types
 Users		Users



JavaDB/Derby - Database Specific Support















Introduction

In addition to the standard feature set in DbVisualizer this document briefly covers the database specific support in DbVisualizer for the JavaDB/Derby RDBMS.

Database Objects Tree

The following table shows what objects DbVisualizer handles for a Derby database. The first column list the structure and the object types. The actions column lists all actions that are associated with the object types and the last column shows all viewers that appear when selecting an object in the tree.

The screenshots illustrates the visual representation of the objects tree.

Object Type	Actions	Viewers
 Schemas	Create Schema	Schemas
 Schema	Drop Schema	
 Tables	Create Table	References, Tables
 Table	Alter Table, Create Index, Drop Table, Empty Table, Import Table Data, Rename Table, Script: CREATE TABLE, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, DDL, Data, Indexes, Info, Navigator, Primary Key, References, Row Count, Row Id
 Views		Views
 View	Drop View, Script: CREATE VIEW, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	DDL, Data, Info, Row Count
 Synonyms	Create Synonym	Synonyms
 Synonym	Drop Synonym, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Info
 Functions		Functions
 Function	Drop Function	Info
 Procedures		Procedures
 Procedure	Drop Procedure	Info
 Triggers		Triggers
 Trigger	Drop Trigger	Info

Mimer SQL - Database Specific Support











Introduction

In addition to the standard feature set in DbVisualizer this document briefly covers the database specific support in DbVisualizer for the Mimer SQL RDBMS.

Database Objects Tree

The following table shows what objects DbVisualizer handles for a Mimer database. The first column list the structure and the object types. The actions column lists all actions that are associated with the object types and the last column shows all viewers that appear when selecting an object in the tree.

The screenshots illustrates the visual representation of the objects tree.

Object Type	Actions	Viewers
 Schemas	Create Schema	Schemas
 Schema	Drop Schema	Info
 Tables	Create Table, Update Statistics	References, Table Facts, Tables
 Table	Add Check, Add Column, Add Foreign Key, Add Primary/Unique Key, Alter Table, Comment Table, Create Index, Create Trigger, Drop Column, Drop Constraint, Drop Table, Empty Table ³ , Grant Access, Import Table Data, Revoke Access, Update Statistics, Script: CREATE TABLE, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Access Rights, Columns, Comments, Constraints, DDL, Data, Indexes, Info, Navigator, Primary Key, References, Row Count, Table used by, Table uses, Triggers
 Column	Comment Column, Data Type, Drop Column, Drop Default, Grant Column Access, Revoke Column Access, Set Default	Column, Column Privileges, Column used by
 Views	Script: Create View	References, View Facts, Views
 View	Comment View, Create Trigger, Drop View, Grant Access, Revoke Access, Script: CREATE VIEW, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Access Rights, Columns, Comments, DDL, Data, Info, References, Row Count, Triggers, View used by, View uses
 ViewColumn	Comment Column, Grant Column Access, Revoke Column Access	Column, Column Privileges, Column used by
 Triggers	Create Table Trigger, Create View Trigger	Trigger Facts, Triggers
 Trigger	Comment Trigger, Drop Trigger	DDL, Info, Trigger uses
 Synonyms	Create Synonym	References, Synonym Facts, Synonyms
 Synonym	Comment Synonym, Drop Synonym, Script: Create Synonym	Access Rights, Columns, Constraints, DDL, Data, Indexes, Info, Primary Key, References, Row Count, Triggers
 Indexes		Index Facts, Indexes
 Index	Comment Index, Drop Index	Columns, DDL, Info
 Sequences	Create Sequence	Sequence Facts, Sequences
 Sequence	Alter Sequence, Comment Sequence, Drop Sequence, Grant Usage, Revoke Usage, Script: Create Sequence	DDL, Info, Sequence used by, Sequence uses, Usage Privilege
 Functions	Create Function	Function Facts, Functions
 Function	Comment Function, Drop Function, Grant Execute, Revoke Execute, Script: Test Function	DDL, Execute Privilege, Function used by, Function uses, Info, Parameters, Result
 Output		Result
 Input		Parameter

MySQL - Database Specific Support






















Introduction

In addition to the standard feature set in DbVisualizer this document briefly covers the database specific support in DbVisualizer for the MySQL RDBMS.

Database Objects Tree

The following table shows what objects DbVisualizer handles for a MySQL database. The first column list the structure and the object types. The actions column lists all actions that are associated with the object types and the last column shows all viewers that appear when selecting an object in the tree.

The screenshots illustrates the visual representation of the objects tree.

Object Type	Actions	Viewers
 Database s	Create Database	Databases
 Catalog	Drop Database	
 Tables	Create Table	References, Tables
 Table	Alter Table, Create Index, Drop Table, Empty Table, Import Table Data, Rename Table, Script: CREATE TABLE, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, DDL, Data, Grants, Indexes, Info, Navigator, Primary Key, References, Row Count, Row Id
 Views ¹		References, Views
 View	Drop View, Rename View, Script: CREATE VIEW ¹ , Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, DDL ¹ , Data, Grants, Info, Row Count
 Procedures 1	Create Procedure	Stored Procedures
 Procedur e	Drop Procedure	DDL, Info
 Functions 1	Create Function	User Functions
 Functio n	Drop Function	DDL, Info
 Triggers ¹		Triggers
 Trigger		DDL, Info
 DBA		DBA
 Users		Users
 Login		Grants
 Roles 2		Server Roles
 Process es		Processes
 Proces s	Kill Session	Process
 Status		Status, Variables
 TableEngine s ²		Table Engines
 Errors ¹		Errors

¹⁾ MySQL version 5 or later
²⁾ MySQL version 4.1 or later

Oracle - Database Specific Support

























Introduction

In addition to the standard feature set in DbVisualizer this document briefly covers the database specific support in DbVisualizer for the Oracle RDBMS.

Database Objects Tree

The following table shows what objects DbVisualizer handles for an Oracle database. The first column list the structure and the object types. The actions column lists all actions that are associated with the object types and the last column shows all viewers that appear when selecting an object in the tree.

The screenshots illustrates the visual representation of the objects tree.

Object Type	Actions	Viewers
 Schemas		Schemas
 Schema		
 Tables	Create Table	References, Table Comments, Tables
 Table	Alter Table, Comment Table, Copy Table, Create Index, Drop Table, Empty Table, Import Table Data, Rename Table, Script: CREATE TABLE, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, Columns Comment, Constraints, Data, Grants, Indexes, Info, Navigator, Primary Key, References, Row Count, Row Id, Triggers
 Columns		Columns
 Column		Column
 Constraints		Constraints
 Constraint		Columns, DDL ¹ , Info
 Column		Column
 Indexes		Indexes
 Index	Compute Statistics for Index, Drop Index, Rebuild Index, Rename Index	Columns, DDL ¹ , Info
 Column		Column
 Triggers		Triggers
 Trigger	Compile Trigger, Disable Trigger, Drop Trigger, Enable Trigger, Rename Trigger	DDL ¹ , Info, Source
 Views		View Comments, Views
 View	Comment View, Compile View, Drop View, Rename View, Script: CREATE VIEW, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, Columns Comment, DDL, Data, Grants, Info, Row Count, Triggers
 Columns		Columns
 Column		Column
 Synonyms	Create Synonym	Synonymed Objects, Synonyms
 Indexes		Indexes
 Index	Compute Statistics for Index, Drop Index, Rebuild Index, Rename Index	Columns, DDL ¹ , Info
 Sequences	Create Sequence	Sequences
 Sequence	Drop Sequence, Rename Sequence	DDL ¹ , Info
 MaterializedView		Materialized Views

Application Object Support (PL/SQL)

DbVisualizer includes special support to manage application objects such as package, package body, procedure and function. These objects can be created, edited and compiled. The editor reports errors and highlights the corresponding problems in the editor.



DBMS_OUTPUT Support

The DBMS Output tab (in SQL Commander) for Oracle is used to enable and disable capturing of messages produced by stored procedures, packages, and triggers. These messages are typically inserted in the code for debugging purposes. For SQL*Plus users the corresponding feature is enabled via the set serveroutput on command. To enable display of DBMS messages in DbVisualizer select the DBMS Output tab and press the Enable button.



Explain Plan

Explain Plan executes your query and records the plan that the database devises to execute it. By examining this plan, you can find out if the database is picking the right indexes and joining your tables in the most efficient manner. The explain plan feature works much the same as executing SQLs to present result sets, you may highlight statements, run a script or load from file. The explain plan results can easily be compared by using the pin feature in combination with window style presentation.



Data Type support

In addition to the standard data type support in DbVisualizer there is also support for the following Oracle specific data types:

- XMLTYPE
- TIMESTAMPTZ
- TIMESTAMPLTZ
- Automatic conversion of DATE data type to TIMESTAMP (configure in Connection Properties)

(For XMLTYPE you need to load the xdb.jar and xmlparserv2.jar in the DbVisualizer Driver Manager. These files are available in the Oracle XML Developers Kit available from Oracle Corp.).



PostgreSQL - Database Specific Support

Introduction

In addition to the standard feature set in DbVisualizer this document briefly covers the database specific support in DbVisualizer for the PostgreSQL RDBMS.

Database Objects Tree

The following table shows what objects DbVisualizer handles for a PostgreSQL database. The first column list the structure and the object types. The actions column lists all actions that are associated with the object types and the last column shows all viewers that appear when selecting an object in the tree.

The screenshots illustrates the visual representation of the objects tree.

Object Type	Actions	Viewers
 Database s		Databases
 Catalog		
 Schema s		Schemas
 Schem a	Comment Schema	Info
 Tables	Create Table	References, Tables
 Table	Alter Table, Comment Table, Create Index, Drop Table, Empty Table, Import Table Data, Rename Table, Script: CREATE TABLE, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, DDL, Data, Grants, Indexes, Info, Navigator, Primary Key, References, Row Count, Row Id
 Views		References, Views
 View	Comment View, Drop View, Rename View, Script: CREATE VIEW, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, DDL, Data, Indexes, Info, Row Count
 Indexe s		Indexes
 Index	Comment Index, Drop Index, Rebuild Index, Rename Index	DDL, Info
 Triggers		Triggers
 Trigger	Disable Trigger, Drop Trigger, Enable Trigger, Rename Trigger	Info, Source
 Function s	Create Function	Functions
 Functio n	Drop Function, Rename Function	Info, Source
 Sequenc es	Create Sequence	Sequences
 Sequenc e	Comment Function, Drop Sequence	Info
 Domain s		Domains
 Domain	Drop Domain	Info
 Aggregate s		Aggregates
 Aggregate	Drop Aggregate, Rename Aggregate	Info
 Operators		Operators
 Operator		Info
 Types		Types
 Type	Drop Type	Info
 DBA		DBA



SQL Server - Database Specific Support



Introduction


In addition to the standard feature set in DbVisualizer this document briefly covers the database specific support in DbVisualizer for the SQL Server RDBMS.

Database Objects Tree

The following table shows what objects DbVisualizer handles for a SQL Server database. The first column list the structure and the object types. The actions column lists all actions that are associated with the object types and the last column shows all viewers that appear when selecting an object in the tree.

The screenshots illustrates the visual representation of the objects tree.

Object Type	Actions	Viewers
 Database s		
 Catalog	Rename Database	
 Tables	Create Table	References, Tables
 Table	Alter Table, Create Index, Drop Table, Empty Table, Import Table Data, Rename Table, Script: CREATE TABLE, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, DDL, Data, Grants, Indexes, Info, Navigator, Primary Key, References, Row Count, Row Id, Triggers
 SystemTable s		References, SystemTables
 SystemTabl e	Script: SELECT	Columns, Data, Info, Row Count
 Views		References, Views
 View	Drop View, Rename View, Script: CREATE VIEW, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, DDL, Data, Indexes, Info, Row Count
 Indexe s		Indexes
 Index	Drop Index, Rebuild Index ¹ , Rename Index	Info
 Users		Users
 Groups		Groups
 Group		Info, Users
 Types		User Defined Data Types
 Triggers		Triggers
 Trigger	Drop Trigger	Info, Source
 Procedure s		Procedures
 Procedur e	Drop Procedure, Rename Procedure	Interface, Source
 Function s		Functions
 Functio n	Drop Function, Rename Function	Interface, Source
 DBA		DBA
 ServerInfo		Server Info
 Logins		Logins
 Device s		Database Devices
 RemoteServe rs		Remote Servers



Process

Processes



Explain Plan

Explain Plan executes your query and records the plan that the database devises to execute it. By examining this plan, you can find out if the database is picking the right indexes and joining your tables in the most efficient manner. The explain plan feature works much the same as executing SQLs to present result sets, you may highlight statements, run a script or load from file. The explain plan results can easily be compared by using the pin feature in combination with window style presentation.

Sybase ASE - Database Specific Support


























Introduction


In addition to the standard feature set in DbVisualizer this document briefly covers the database specific support in DbVisualizer for the Sybase Adaptive Server Enterprise RDBMS.

Database Objects Tree

The following table shows what objects DbVisualizer handles for a Mimer database. The first column list the structure and the object types. The actions column lists all actions that are associated with the object types and the last column shows all viewers that appear when selecting an object in the tree.

The screenshots illustrates the visual representation of the objects tree.

Object Type	Actions	Viewers
 Database s		
 Catalog		
 Tables	Create Table	References, Tables
 Table	Alter Table, Create Index, Drop Table, Empty Table, Import Table Data, Rename Table, Script: CREATE TABLE, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, DDL, Data, Grants, Indexes, Info, Navigator, Primary Key, References, Row Count, Row Id, Triggers
 SystemTable s		References, SystemTables
 SystemTabl e	Script: SELECT	Columns, Data, Info, Row Count
 Views		References, Views
 View	Drop View, Rename View, Script: CREATE VIEW, Script: DELETE, Script: INSERT, Script: SELECT, Script: UPDATE	Columns, DDL, Data, Indexes, Info, Row Count
 Users	Add User	Users
 User	Drop User	Info
 Groups	Add Group	Groups
 Group	Drop Group	Info, Users
 Types		User Defined Data Types
 Triggers		Triggers
 Trigger	Disable Trigger, Drop Trigger, Enable Trigger, Rename Trigger	Info, Source
 Procedure s		Procedures
 Procedur e	Drop Procedure, Rename Procedure	Interface, Source
 DBA		DBA
 ServerInfo		Character Set, Server Info
 Logins	Create Login	Logins
 Login	Change Password, Drop Login	Databases, Info, Roles
 Device s		Database Devices
 Devic e		Databases, Info
 RemoteServe rs		Remote Servers
 Process es		Processes



set showplan

DbVisualizer captures any output produced by the `set showplan` SQL statement and displays it in the SQL Commander log.