

DB Optimizer 19.1

User Guide

Exported on 04/14/2026

I D E R A

Table of Contents

| | |
|---|-----------|
| Welcome to DB Optimizer | 8 |
| Installation and Configuration | 9 |
| Installing DB Optimizer | 9 |
| Installation Issues | 9 |
| Installing DB Optimizer on Linux | 9 |
| Initial setup | 13 |
| Specify a workspace | 13 |
| Licensing DB Optimizer | 14 |
| Customizing DB Optimizer (Preferences) | 14 |
| Specify Data Sources preferences | 15 |
| Specify SQL Development preferences | 16 |
| Specify SQL Editor preferences..... | 16 |
| Specify Code Assist preferences | 18 |
| Specify Code Formatter preferences..... | 19 |
| Specify Code Quality preferences..... | 20 |
| Specify Results Viewer preferences | 21 |
| Specify SQL Templates preferences | 23 |
| Specify Syntax Coloring preferences..... | 24 |
| Specify SQL Execution preferences | 25 |
| Specify SQL Filters preferences | 26 |
| Specify Profile Alerts preferences..... | 27 |
| Specify Profile Repositories preferences..... | 27 |
| Specify Tuning Job Editor preferences | 27 |
| Specify VST Diagrams Tuning preferences..... | 27 |
| Specify File Encoding preferences..... | 28 |
| Adjusting Java heap size to optimize performance..... | 29 |
| Editing -Xms and -Xmx in the dboptimizer.ini..... | 29 |
| Introduction to database tuning..... | 31 |
| Introduction to DB Optimizer tuner..... | 31 |
| Tuning example | 31 |
| SQL tuning methodology | 32 |
| SQL tuner overview | 33 |

| | |
|---|-----------|
| What's happening on the databases?..... | 34 |
| Databases..... | 35 |
| Average Active Sessions (AAS) Load of selected database..... | 35 |
| Maximum CPU line..... | 35 |
| Top SQL, Top Bottlenecks, and Top Sessions..... | 35 |
| Tuning example | 37 |
| The database is hanging or the application has problems | 37 |
| The database caused the problem | 58 |
| The machine caused the problem | 58 |
| Finding and tuning problem SQL..... | 76 |
| Using DB Optimizer..... | 77 |
| Working with data sources..... | 77 |
| Register data sources | 78 |
| Add a new data source | 79 |
| Import and export data sources | 83 |
| Categorize data sources | 84 |
| Customizing data source categories | 85 |
| Browse a data source | 85 |
| View database object properties | 85 |
| Search for database objects..... | 87 |
| Filter database objects..... | 89 |
| Define Data Source-Specific Object Filters..... | 89 |
| Define Global Database Object Filters..... | 90 |
| Drop a database object | 90 |
| Working with SQL projects | 90 |
| Create a new SQL project..... | 91 |
| Open an existing project | 92 |
| Search a Project..... | 92 |
| Add files to a project..... | 92 |
| Delete a Project..... | 93 |
| Creating and editing SQL files (SQL Editor)..... | 93 |
| Semantic Validation | 95 |
| Create an SQL file | 95 |
| Open an existing SQL file..... | 95 |

| | |
|---|------------|
| Working in SQL Editor..... | 96 |
| Understanding Automatic Error Detection | 97 |
| Understanding Code Assist | 98 |
| Understanding Hyperlinks | 104 |
| Understanding Code Formatting..... | 104 |
| Understanding Code Folding | 121 |
| Understanding Code Quality Checks..... | 122 |
| View change history..... | 126 |
| Revert to an old version of a file..... | 127 |
| Delete an SQL file..... | 128 |
| Executing SQL files | 128 |
| Associate an SQL file with a data source | 129 |
| Configure an SQL session..... | 130 |
| Execute SQL code | 131 |
| View and save results | 131 |
| Troubleshooting | 132 |
| View log details | 134 |
| Maintain logs..... | 134 |
| Filter logs..... | 135 |
| Import and export error logs..... | 137 |
| Find and fix SQL code errors | 138 |
| Find and fix other problems..... | 138 |
| Using profiling..... | 140 |
| Understanding profiling | 140 |
| Understanding the interface | 141 |
| Running a profiling session | 142 |
| Execute a profiling session..... | 142 |
| Executing a session from the Command Line | 143 |
| Working with session results..... | 149 |
| Opening an existing profiling session..... | 149 |
| Filtering results | 150 |
| Analyze the load chart..... | 151 |
| Analyze the Top Activity section | 152 |
| Analyze profiling details | 159 |

| | |
|---|------------|
| Creating profiling reports | 179 |
| Saving profiling sessions | 181 |
| Work with the Profiling Repository | 182 |
| Import statements to tuning..... | 184 |
| Other profiling commands | 184 |
| Zooming in and out | 184 |
| Configuring profiling..... | 185 |
| Configuring DBMS properties and permissions | 185 |
| Configuring IBM DB/2 for Windows, Unix, and Linux | 185 |
| Configuring Microsoft SQL Server | 199 |
| Configuring Oracle..... | 199 |
| Configuring Sybase..... | 199 |
| Building profiling configurations | 200 |
| Specify profile alerts preferences during configuration | 203 |
| Specify profile repositories preferences during configuration | 206 |
| Register a profile repository..... | 207 |
| Unregister a profiling repository..... | 208 |
| Using SQL Load Editor/Tester | 209 |
| Methodology | 210 |
| Using tuning..... | 212 |
| Understanding the tuner interface | 212 |
| Understanding the Overview tab..... | 212 |
| Inputting SQL to tune | 212 |
| Running a tuning job | 215 |
| Creating tuning reports | 216 |
| Understanding the Analysis tab..... | 219 |
| Tuning SQL statements | 221 |
| Create a new tuning job | 221 |
| Specify a data source..... | 222 |
| Add SQL statements | 222 |
| Managing bind variable errors | 224 |
| Run a tuning job..... | 226 |
| Analyze tuning results | 228 |
| Compare cases..... | 231 |

| | |
|---|-----|
| Filter and delete cases | 231 |
| Create an outline | 232 |
| Modify tuning results | 232 |
| DB2, Oracle, SQL Server, and Sybase query rewrites | 233 |
| Using the Analysis tab | 234 |
| Implementing index analysis recommendations | 235 |
| Visual SQL tuning..... | 236 |
| Changing diagram detail display | 236 |
| Interpreting the VST diagram graphics | 245 |
| Using platform-specific features..... | 255 |
| Using the Table Statistics tab (Oracle and SQL Server) | 256 |
| Using the Column Statistics and Histograms tab (Oracle and SQL Server)..... | 256 |
| Oracle-specific Column Statistics and Histograms tab example | 257 |
| Using the Outlines tab (Oracle)..... | 258 |
| Using the Plan Guides tab (SQL Server)..... | 270 |
| Tuning SQL statements in the System Global Area (Oracle) | 277 |
| Additional tuning commands | 278 |
| View the source code of tuning candidates | 279 |
| View statement or case code in SQL Viewer..... | 279 |
| Open an explain plan for a statement or case | 279 |
| Executing a session from the command line | 280 |
| Saving a tuning job | 285 |
| Configuring tuning..... | 286 |
| Set roles and permissions on data sources..... | 286 |
| Specify Tuning Job Editor preference | 351 |
| General Preferences | 352 |
| Specify Case Generation preferences..... | 353 |
| About statement records | 354 |
| Specify VST Diagrams Tuning preference | 354 |
| Examples of transformations and SQL query rewrites | 355 |
| Cartesian product elimination..... | 355 |
| Expression transformation..... | 355 |
| Invalid outer join..... | 355 |
| Transitivity | 355 |
| Move expression to WHERE clause | 355 |

| | |
|--|------------|
| NULL column | 356 |
| Push subquery | 356 |
| Mismatched column types | 356 |
| DBMS hints | 356 |
| SQL Server hints..... | 357 |
| DB2 hints | 360 |
| Sybase hints | 361 |
| Reference..... | 365 |
| Database objects | 365 |
| DBMS connection parameters by platform | 374 |
| Using Multi-Factor Authentication | 374 |
| IBM DB2 LUW..... | 375 |
| Microsoft SQL Server | 375 |
| JDBC connection parameters | 377 |
| Oracle connection parameters | 377 |
| Sybase connection parameters | 378 |

Welcome to DB Optimizer

IDERA DB Optimizer simplifies SQL optimization and development for application developers with many features for improving productivity and reducing errors. A rich SQL IDE with statement tuning, data source profiling, code completion, real-time error checking, code formatting and sophisticated object validation tools helps streamline coding tasks. DB Optimizer's user interface helps improve overall productivity with integrated development, monitoring, and tuning components.

DB Optimizer has four components that when used together can optimize your database performance.

- **SQL Editor.** A developer can write Java in Eclipse that calls to the database with SQL. The SQL that calls to the database can be written in the SQL Editor with type ahead, code assist and quick fixes to show the users syntax and correct mistakes. For more information, see [Creating and editing SQL files \(SQL Editor\)](#).
- **Load Tester.** The SQL code can be run in the Load Tester to test execution by multiple concurrent users. User load testing is so often done by one single user and then problems don't appear until production with multiple concurrent users. Concurrent user testing is a breeze in DB Optimizer. For more information, see [Using SQL Load Editor/Tester](#).
- **Profiler.** You can run the Profiler while the Load Tester is executing to show clearly the impact on the database. The profiler can also be used by QA on load simulation. Finally, the Profiler can be run on any production database to clearly show load, bottlenecks, and sources of bottlenecks or resource consumption. For more information, see [Using profiling](#).
- **Tuner.** Finally, if a problem SQL is found on the system the Tuner will show if it's correctly optimized by the database or not, and if not it will show the best plan and what hints or optimizer directives can be included in the SQL to force the database to use the optimal plan. For more information, see [Using tuning](#).

DB Optimizer is part of a larger package of products known as DB PowerStudio, or DBPS. These tools work together to provide complete database management and development, from administration and change management to query tuning and development, for multiple database platforms.

Installation and Configuration

This section contains information about installing and configuring DB Optimizer. It includes information on setting up the system directory for project files, as well as licensing information. Additionally, this section contains information on setting preferences within the application for the customization of various features and functionality.

- [Installing DB Optimizer](#)
- [Initial setup](#)
- [Customizing DB Optimizer \(Preferences\)](#)

Installing DB Optimizer

To install DB Optimizer, run the installer and follow the prompts provided by the Installation Wizard. When the installation is complete, check the Readme file to see if it contains any known issues or advisory notes that will affect your installation of the product.

- [Installation Issues](#)
- [Installing DB Optimizer on Linux](#)

Installation Issues

Install DB Optimizer in a different directory than previous versions of DB Optimizer. If DB Optimizer is installed in the same directory, the following error appears at application startup: "Unable to read workbench state. Workbench UI layout will be reset."

If a previous version is already installed on your machine, either change the default Workspace or delete all files from the Workspace directory. The Workspace directory location appears whenever you start DB Optimizer, and can be redefined from this dialog box. If the same Workspace is indicated for the new version as was indicated previously, the following error appears: "Cannot restore workbench layout."

If DB Optimizer is already running, you can also modify the Workspace directory by choosing **File > Switch Workspace > Other** from the Main Menu.

Installing DB Optimizer on Linux

To install DB Optimizer in a Linux environment

1. Download the DB Optimizer installer and set the execution permissions as follows:



```
chmod
```

755

dbo_<ve

```
rsion> .
```

```
bin
```

2. Launch the installer in GUI mode from the console.

```
./
```

dbo_ < ve

rsion> .

bin

The DB Optimizer splash screen appears for a short time and afterwards the installer dialog appears.

3. On the Introduction dialog, click **Next**.
4. On the License Agreement dialog, read the License Agreement, select **I accept the terms of the License Agreement**, and then click **Next**.
5. On the Choose Install Set dialog, choose one of the following options, and then click **Next**.
 - **Install as Standalone Application.** This option installs DB Optimizer as a Standalone Application, without any prerequisites.
 - **Install as Eclipse Plug-in.** This option requires Eclipse 3.6 or higher or an Eclipse-based RCP application.

6. On the Choose Instal Folder dialog, depending on the previously selected option, you can select a destination folder where to install the Standalone version or you are required to browse to the folder that contains Eclipse 3.6 (or higher) or an RCP application based on Eclipse 3.6.
7. After selecting the desired folder, click **Next**.
8. On the Choose Link Folder dialog, you can change the location where the Application/ Uninstall links are created. Choose one of the options, and then click **Next**.
9. On the Pre-Installation Summary, review the installation details, and then click **Install**.
The Installing dialog shows the progress of the installation; it is closed when the installation is complete.
10. When the Install Complete dialog appears, you can select **Start IDERA DB Optimizer**, and then click **Done** in order to start the application for a first time use.

Initial setup

The following topics provide general help for configuring DB Optimizer:

- [Specify a workspace](#)
- [Licensing DB Optimizer](#)

Additionally, the following preferences are available to help you customize and tune functions within the application:


- [Specify Data Sources preferences](#)
- [Specify SQL Development preferences](#)
- [Specify SQL Editor preferences](#)
- [Specify SQL Execution preferences](#)
- [Specify SQL Filters preferences](#)
- [Specify Profile Alerts preferences](#)
- [Specify Profile Repositories preferences](#)
- [Specify Tuning Job Editor preferences](#)
- [Specify VST Diagrams Tuning preferences](#)
- [Specify File Encoding preferences](#)

Specify a workspace

When you install DB Optimizer, you are prompted to create a workspace. Then, when you launch DB Optimizer, you have an opportunity to choose your workspace. At any time while running DB Optimizer, you can change your workspace.

To switch your workspace

1. In DB Optimizer, select **File > Switch Workspace**. The Workspace Launcher appears.



2. Browse for the appropriate workspace or create a copy of an existing workspace using the available **Copy Settings**.
3. Click **OK** after making your selection. DB Optimizer opens or creates the appropriate workspace.

Licensing DB Optimizer

The first time you launch DB Optimizer, you are prompted to activate the product. Choose to activate by Internet, and then follow the prompts. During the activation process you will receive an email with an activation key; after you enter that key into the License Setup dialog, you will receive a free 14-day evaluation license.

If due to firewall or other restrictions you cannot use Internet activation, select the email alternative.

To continue using DB Optimizer after the evaluation period

1. Select **Help > IDERA Licensing > License Registration**.
2. Follow the prompts.

Customizing DB Optimizer (Preferences)

To customize various aspects of DB Optimizer, select the aspect you want to customize from the **Preferences** menu.

This section is comprised of the following topics:

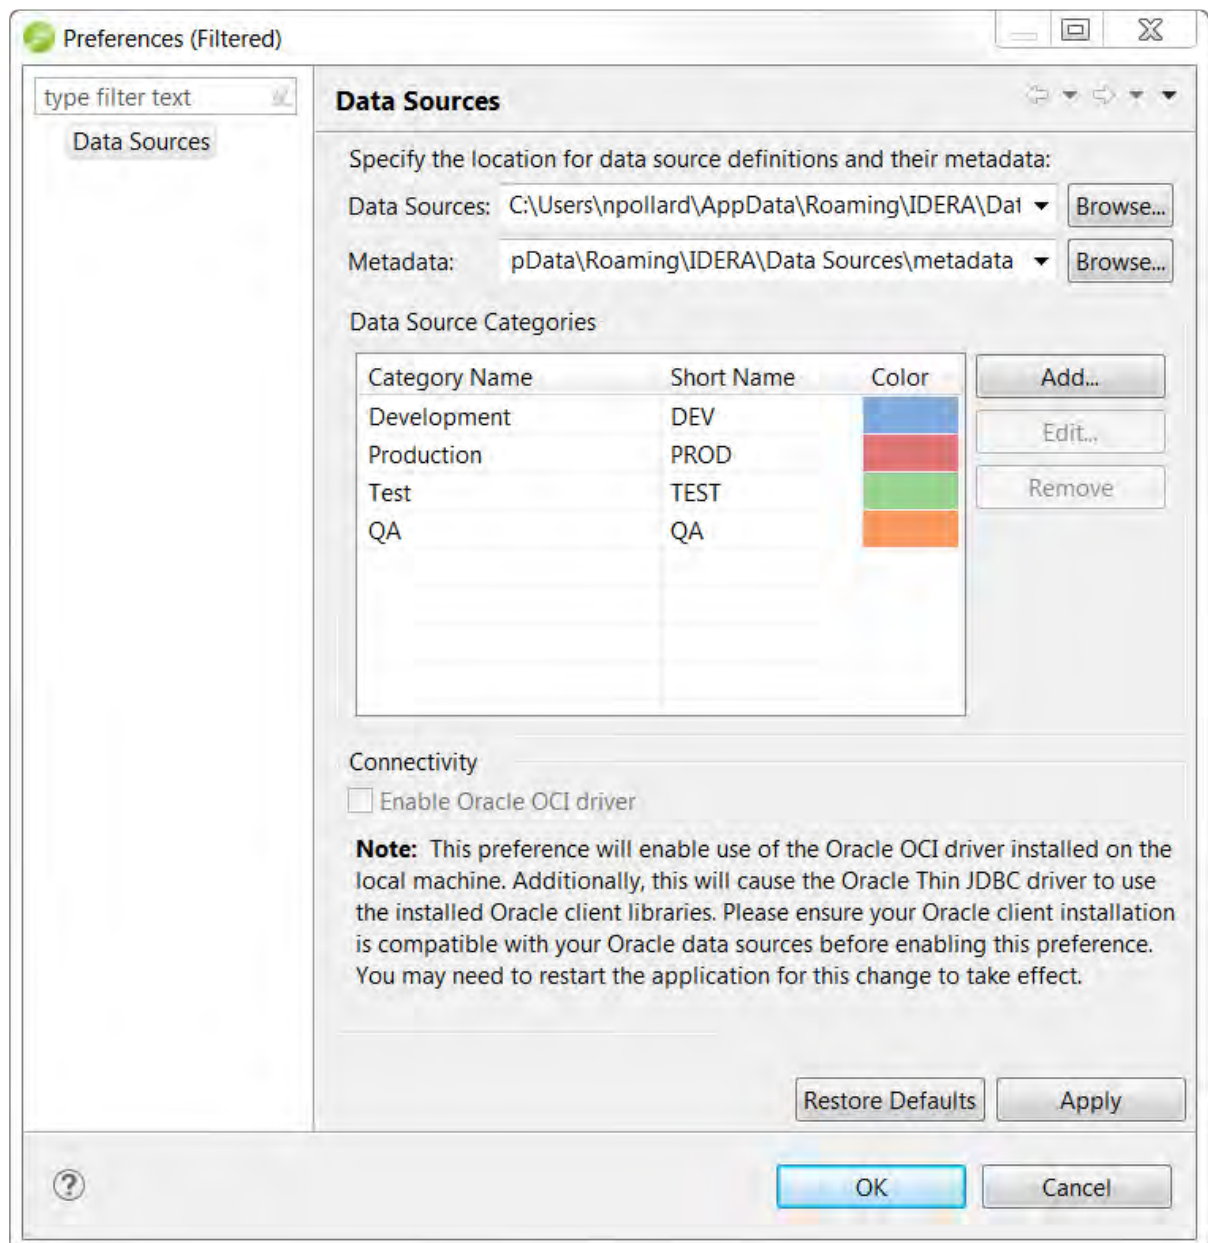
- [Specify Data Sources preferences](#)
- [Specify SQL Development preferences](#)
- [Specify SQL Editor preferences](#)
- [Specify SQL Execution preferences](#)
- [Specify SQL Filters preferences](#)
- [Specify Profile Alerts preferences](#)
- [Specify Profile Repositories preferences](#)
- [Specify Tuning Job Editor preferences](#)
- [Specify VST Diagrams Tuning preferences](#)
- [Specify File Encoding preferences](#)
- [Adjusting Java heap size to optimize performance](#)

Specify Data Sources preferences

When you add a data source to your list of Managed Data Sources in the Data Source Explorer, DB Optimizer stores the definition and metadata for the data source in the location you specify here. For information on adding data sources, see [Register data sources](#).

To specify Data Sources preferences

1. From the **Preferences** menu, select **Data Sources**. The **DataSources** pane appears.



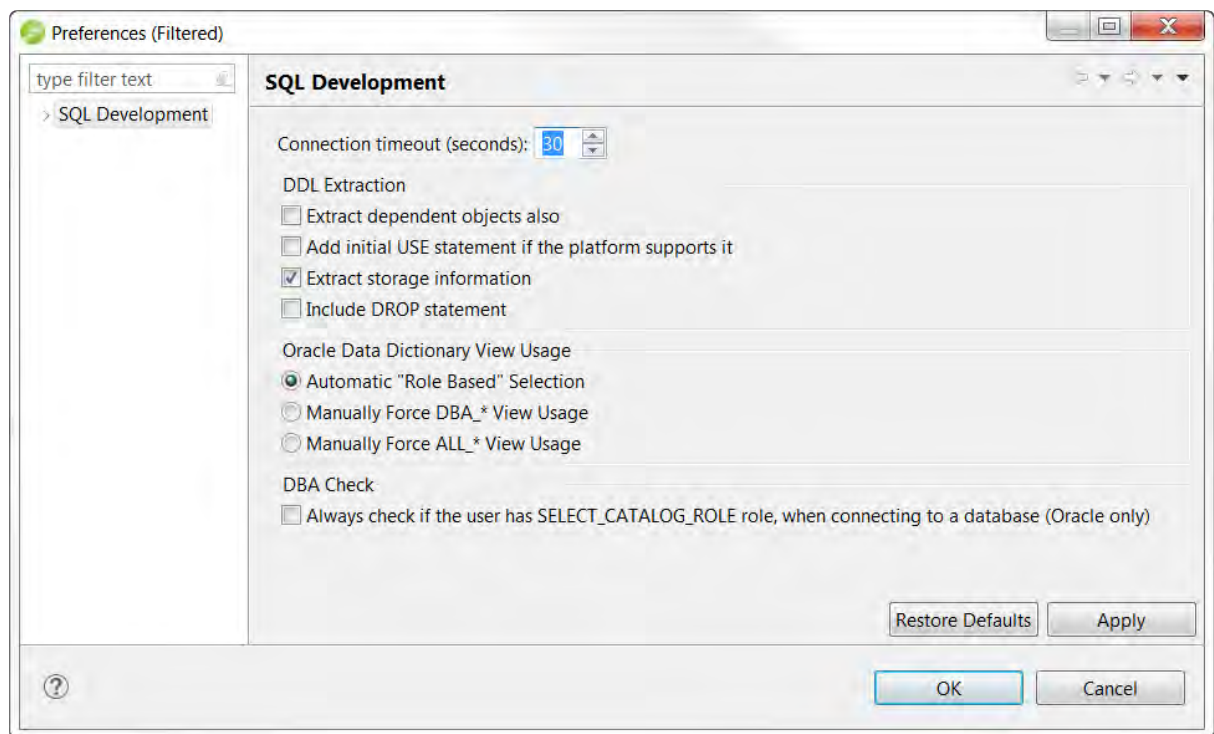
2. Specify the location for data source definitions and their metadata.
3. Click **Apply**.

 For information on adding custom categories, see [Customizing data source categories](#).

Specify SQL Development preferences

The SQL Development Preferences specified on the first page of the SQL Development Preferences determines DB Optimizer's behavior when connecting to and extracting DDL from a data source. For information on preferences accessible by expanding SQL Development, see [Customizing DB Optimizer \(Preferences\)](#).

1. From the **Preferences** menu, select **SQL Development**.



2. Choose your preferences, and then click **Apply**.

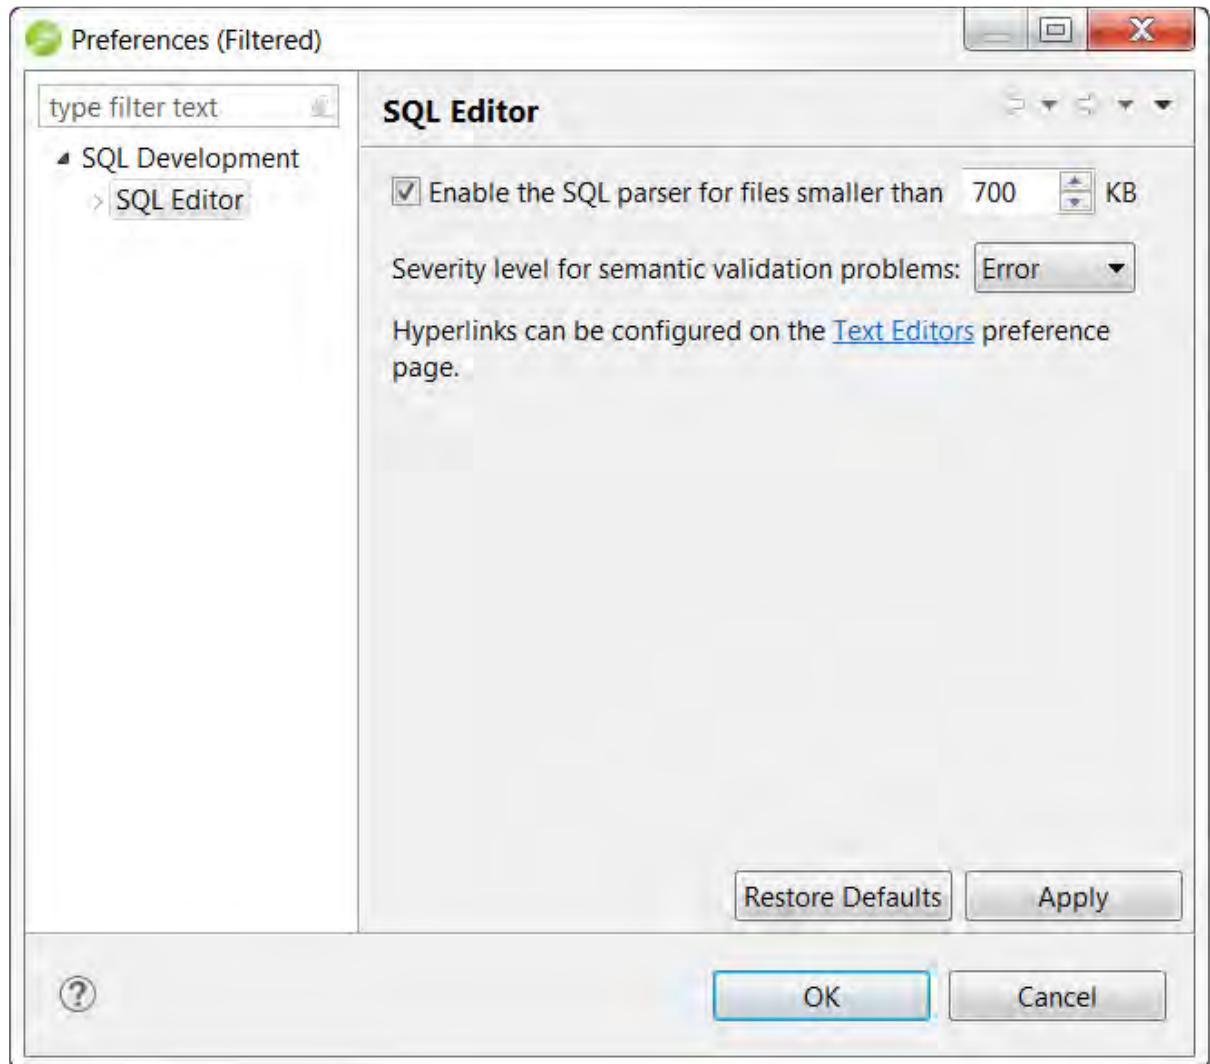
The following describes the SQL Development preferences:

- **Connection timeout (seconds)**. Specify the connection timeout before the connection to the database fails.
- **Extract dependent objects also**. If selected, when extracting DDL dependent objects such as indexes are also extracted.
- **Add initial USE statement if the platform supports it**. If selected, a USE statement is added to the DDL extracted. Adding the USE statement ensures that when you run the DDL, you are using the correct database context.
- **Extract storage information**. If selected, when extracting DDL object storage information is also extracted.
- **Include DROP statement**. If selected, the DROP statement will be added to the DDL so you can easily execute the statement.

Specify SQL Editor preferences

To specify SQL Editor preferences

1. Select **Preferences > SQL Editor**.



2. Change the settings as appropriate in each section and then click **Apply**. the following options are available:

- **Enable the SQL parser.** For performance reasons, you may want to enable the SQL parser only if a SQL file is smaller than the size you specify here.
- **Severity level for semantic validation problems.** Choose a security level from this list. This option determines how semantic code errors are flagged in the editor and the Problems view.
- The link to specify hyperlinks takes you to the Text Editors preference page.

i Clearing **Enable SQL Parser** disables many of the "smart" SQL editor features, including code formatting, auto completion, semantic validation, and hyperlinks. For better performance, you may disable the parser for files above a specified size.

See the following topics to configure other SQL Editor preferences.

- [Specify Code Assist preferences](#)
- [Specify Code Formatter preferences](#)
- [Specify Code Quality preferences](#)
- [Specify Results Viewer preferences](#)

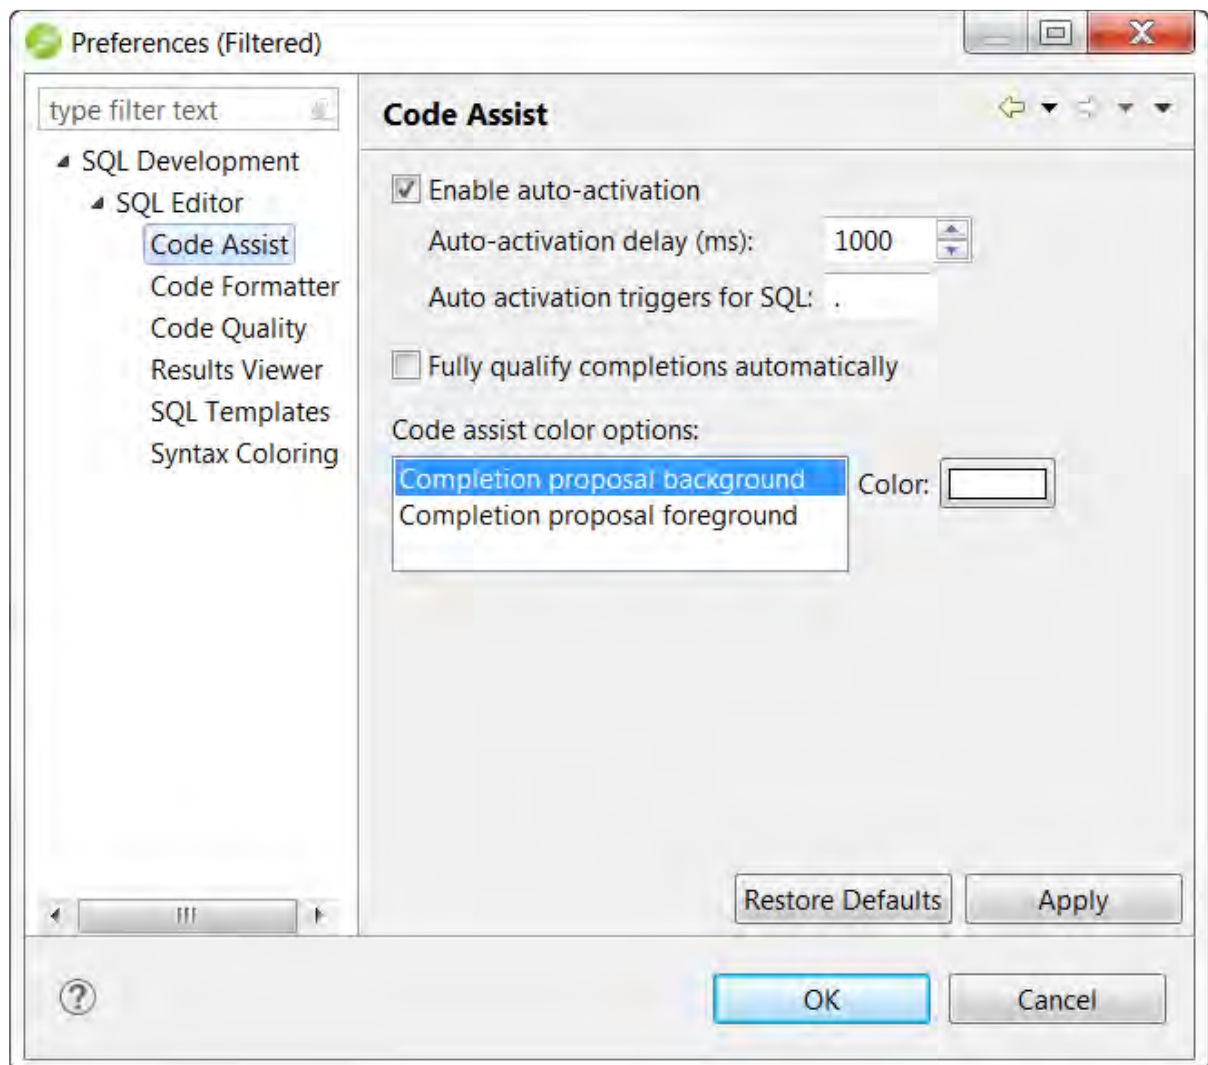
- [Specify SQL Templates preferences](#)
- [Specify Syntax Coloring preferences](#)

Specify Code Assist preferences

The Code Assist panel is used to specify configuration parameters that determine how code completion features in SQL Editor behave.

To specify Code Assist preferences

1. Select **Preferences > SQL Editor**.



2. In the **Preferences** dialog, expand **SQL Editor**, and then click **Code Assist**.

The following options are available on the Code Assist Preferences page:

- **Enable auto-activation.** When selected enables code assist functionality with the Ctrl + Space command. If this option is selected, the code assist window automatically appears when you stop typing.
- **Auto-activation delay (ms).** Specify the amount of time in milliseconds that the window automatically appears.

- **Auto activation triggers for SQL.** Enter a trigger character or trigger characters. When you enter an activation trigger in the SQL Editor, you will see the code assist options available.
- **Fully qualify completions automatically.** When selected, specifies if code completion results are returned specific (fully qualified), rather than the minimum required to identify the object.
- **Code assist color options.** Specifies the color formatting of code completion proposals.

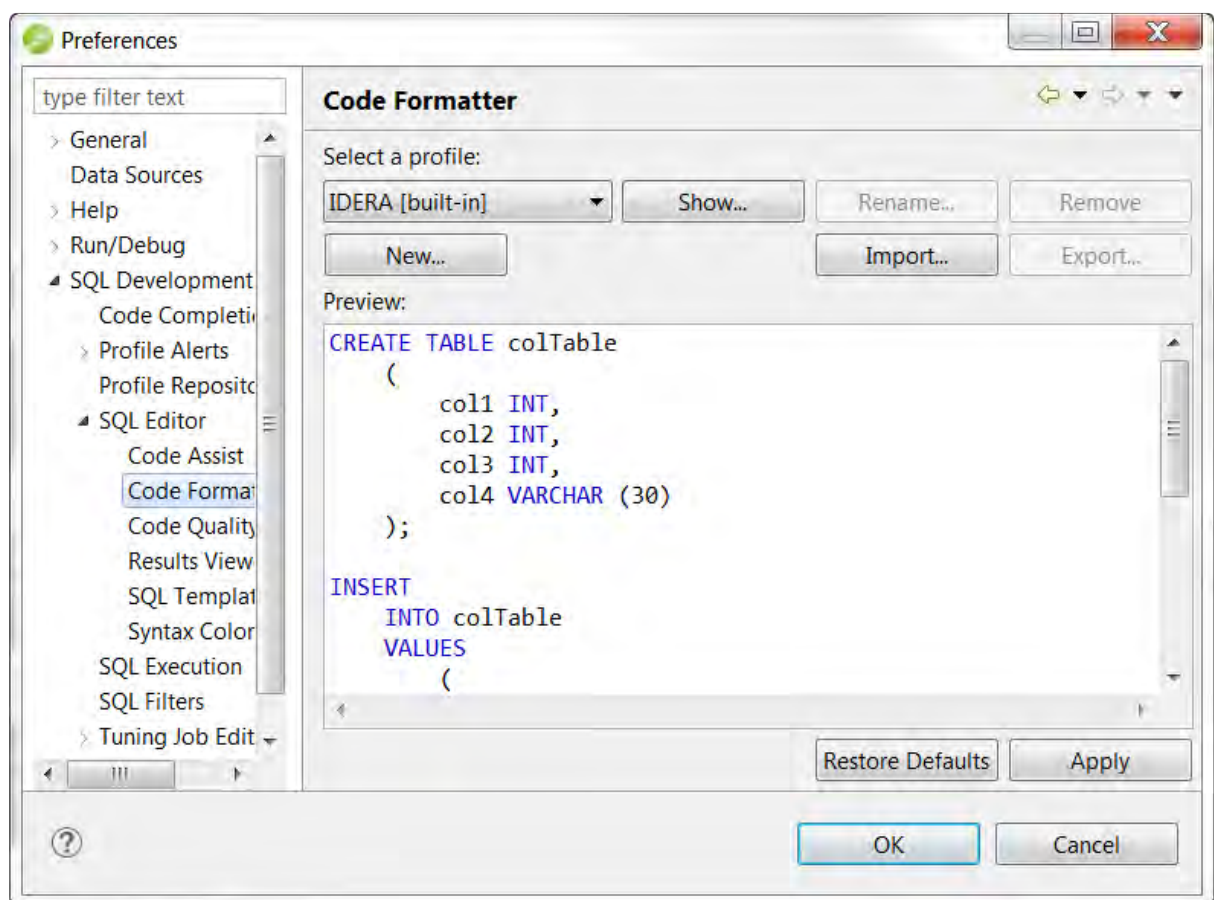
Select background or foreground options from the menu and modify them as appropriate.

Specify Code Formatter preferences

The Code Formatter pane provides configuration options for code formatting functionality in SQL Editor.

To specify Code Formatter preferences

1. Select **Preferences > SQL Editor**.



2. In the **Preferences** dialog, expand **SQL Editor**, and then click **Code Formatter**.

The panel provides a drop-down list of formatting profiles and a preview window that displays how each profile formats code.

- **Select a profile.** From the list, choose the profile you want to view.
- Click **Show** to view the details of the Profile. On the **Show Profiles** dialog that appears, you can edit the profile and save the changes.
- Click **New** to define additional code formatting profiles.

- Click **Edit** to modify existing profiles. You can modify how code characters appear in the interface and how SQL Editor determines line breaks.
- Click **Rename** to change the name of an existing profile. The new name cannot be the same as another existing profile.

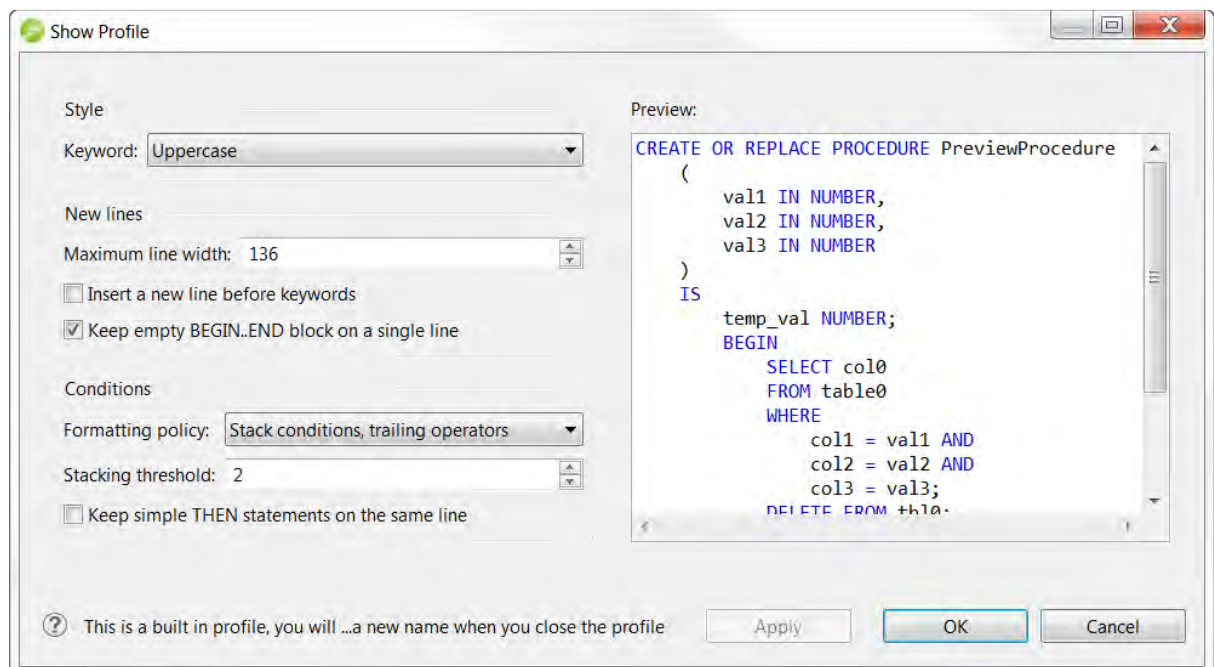
i If you create a new profile with a name that already exists in the system, a prompt will appear asking you to change the name of the new code formatting template.

Create and edit code formatting profiles

You can create your own code formatting profiles that define how your SQL code is formatted.

To create a code formatting profile

1. Select **Preferences > SQL Editor**.
2. In the **Preferences** dialog, expand SQL Editor, and then click **Code Formatter**.
3. On the **Code Formatter** pane, click **Show** or **New** and a dialog similar to the following appears.



4. Make your changes, and then click **Apply** to preview your changes.
5. Click **OK** to create the new profile or to exit the **Show Profile** dialog.
6. Use the preview pane to preview the changes you make to your code formatting preferences. Changes are not implemented until you click **Apply** or **OK**.

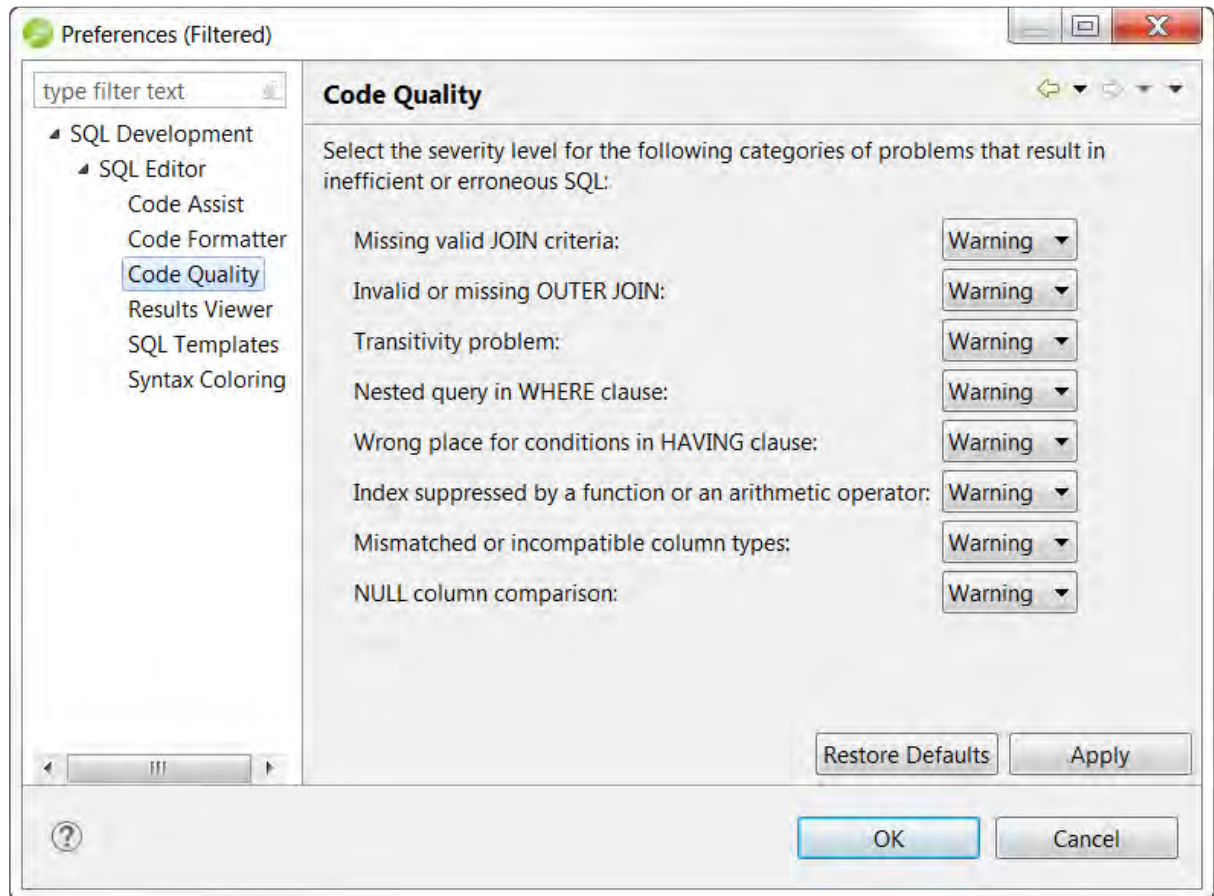
Specify Code Quality preferences

The Code Quality preferences allow you to specify the severity level for several categories of problems that result in inefficient or erroneous SQL.

To specify Code Quality preferences

1. Select **Preferences > SQL Editor**.

2. In the **Preferences** dialog, expand **SQL Editor**, then and click **Code Quality**.



3. Change the severity levels by clicking the list next to the category, and then choosing the level.

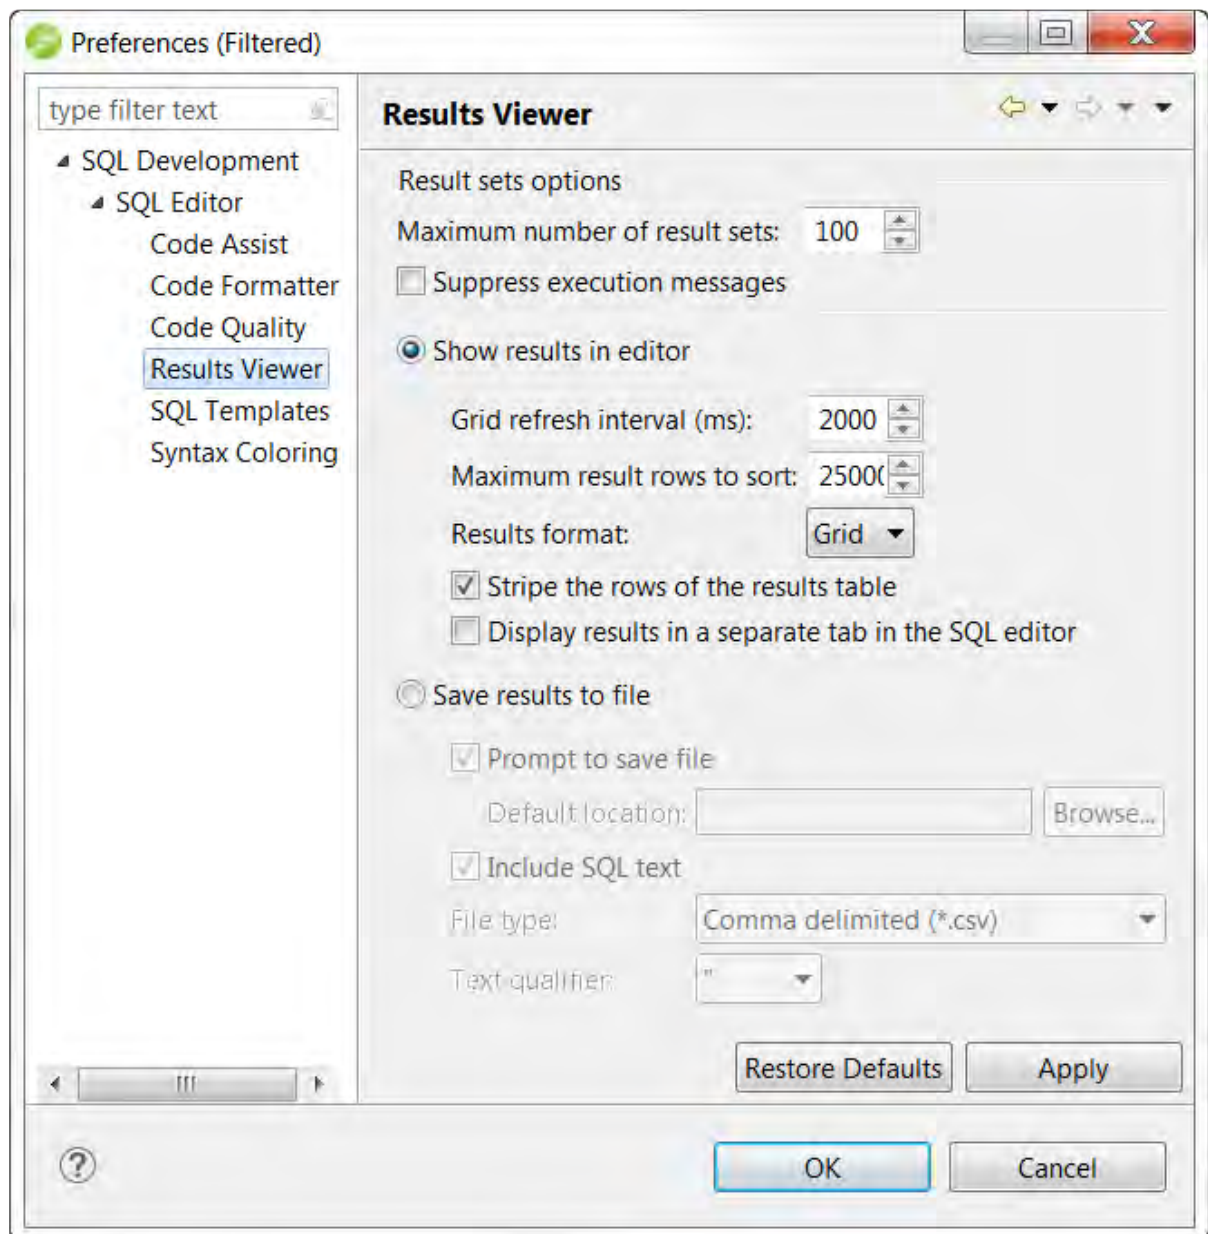
4. Click **Apply** to save your changes.

Specify Results Viewer preferences

The Results Viewer pane provides configuration options that specify how the Results view displays or saves results.

To specify Results Viewer preferences

1. Select **Preferences > SQL Editor**.
2. In the **Preferences** dialog, expand **SQL Editor**, and then click **Results Viewer**.



3. Make your preference changes, and then save your changes by clicking **Apply**. The following options are available:

- **Maximum number of result sets:** If selected and the executed SQL returns more results sets than the maximum specified, result sets in excess of the maximum specified are not displayed.
- **Suppress execution messages:** If selected and the SQL you execute returns informational messages, they are not displayed.
- **Show results in editor:** Execution results can be either shown in the editor or sent to a file.
- **Grid refresh interval:** Indicates the speed (in milliseconds) that the Results view refreshes.
- **Maximum result rows to sort:** If the number of rows and results exceed this number, the column sorting in the result set is disabled.
- **Results format:** These are the different formats that can be used to display the results in the editor.
- **Stripe the rows of the results table:** Adds intermittent highlighted bars in the Results view.

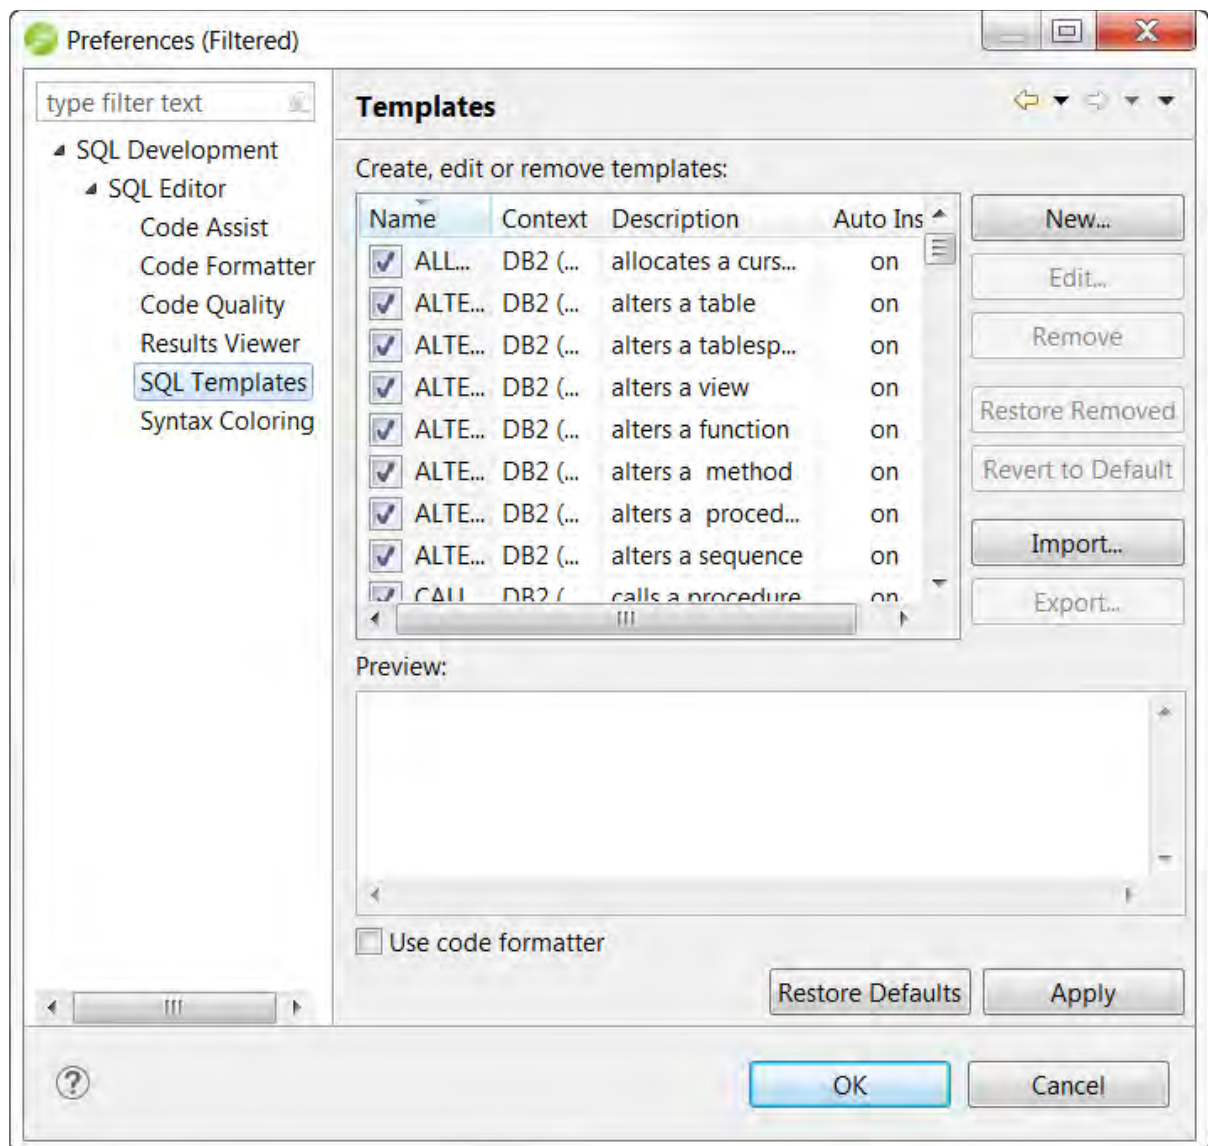
- **Display results in separate tab in SQL Editor:** Opens the Results view in a separate window on the Workbench.
- **Save results to file:** Provides options that let you save the contents of the result sets to a file. You can also specify the file type, delimiter and text qualifier.

Specify SQL Templates preferences

The SQL Templates panel provides customization options for creating and modifying SQL code templates.

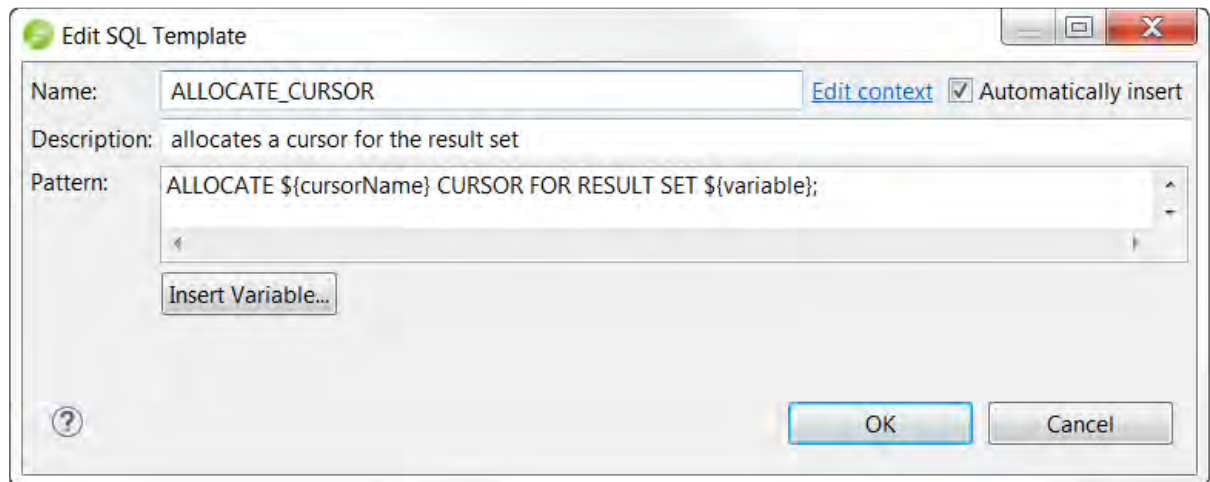
To specify SQL Templates preferences

1. Select **Preferences > SQL Editor**.
2. In the **Preferences** dialog, expand **SQL Editor**, and then click **SQL Templates**. The SQL Templates panel displays a list of all SQL code templates currently available. Additionally, when you select a template from the list, the **Preview** section displays the code block as it will appear when the template is selected in SQL Editor.



3. Click the check box beside each template to specify whether it is included in the code assist check within SQL Editor. Use the buttons on the right-hand side of the panel to create, edit, or delete SQL templates, as needed.

When you create or edit a template, the **Edit SQL Template** dialog appears.



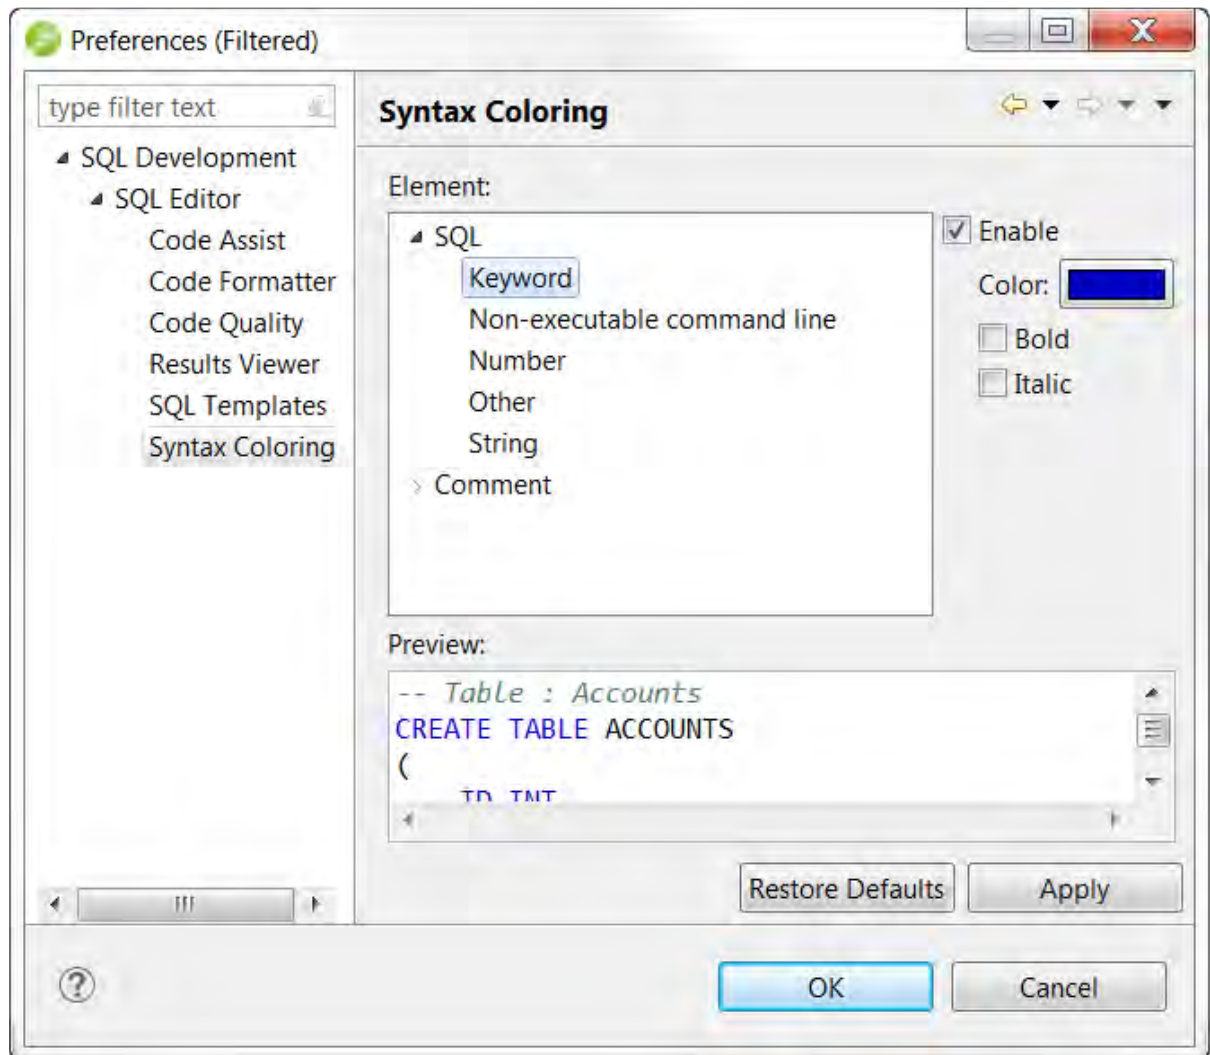
4. Enter a **Name**, **Description**, and **Pattern** in the fields provided, and then click **OK**. If the template name doesn't match an existing SQL code template, your new template is added to the list, and will automatically be considered when the code assist function is executed in SQL Editor.
5. Select the **Use Code Formatter** check box to apply code formatting preferences to the specified template. For more information about setting code formatter preferences, see [Specify Code Formatter Preferences](#).

Specify Syntax Coloring preferences

The Syntax Coloring panel provides configuration options that change the look and feel of code syntax in SQL Editor.

To specify Syntax Coloring preferences

1. Select **Preferences > SQL Editor**.
2. In the **Preferences** dialog, expand **SQL Editor**, and then click **Syntax Coloring**.



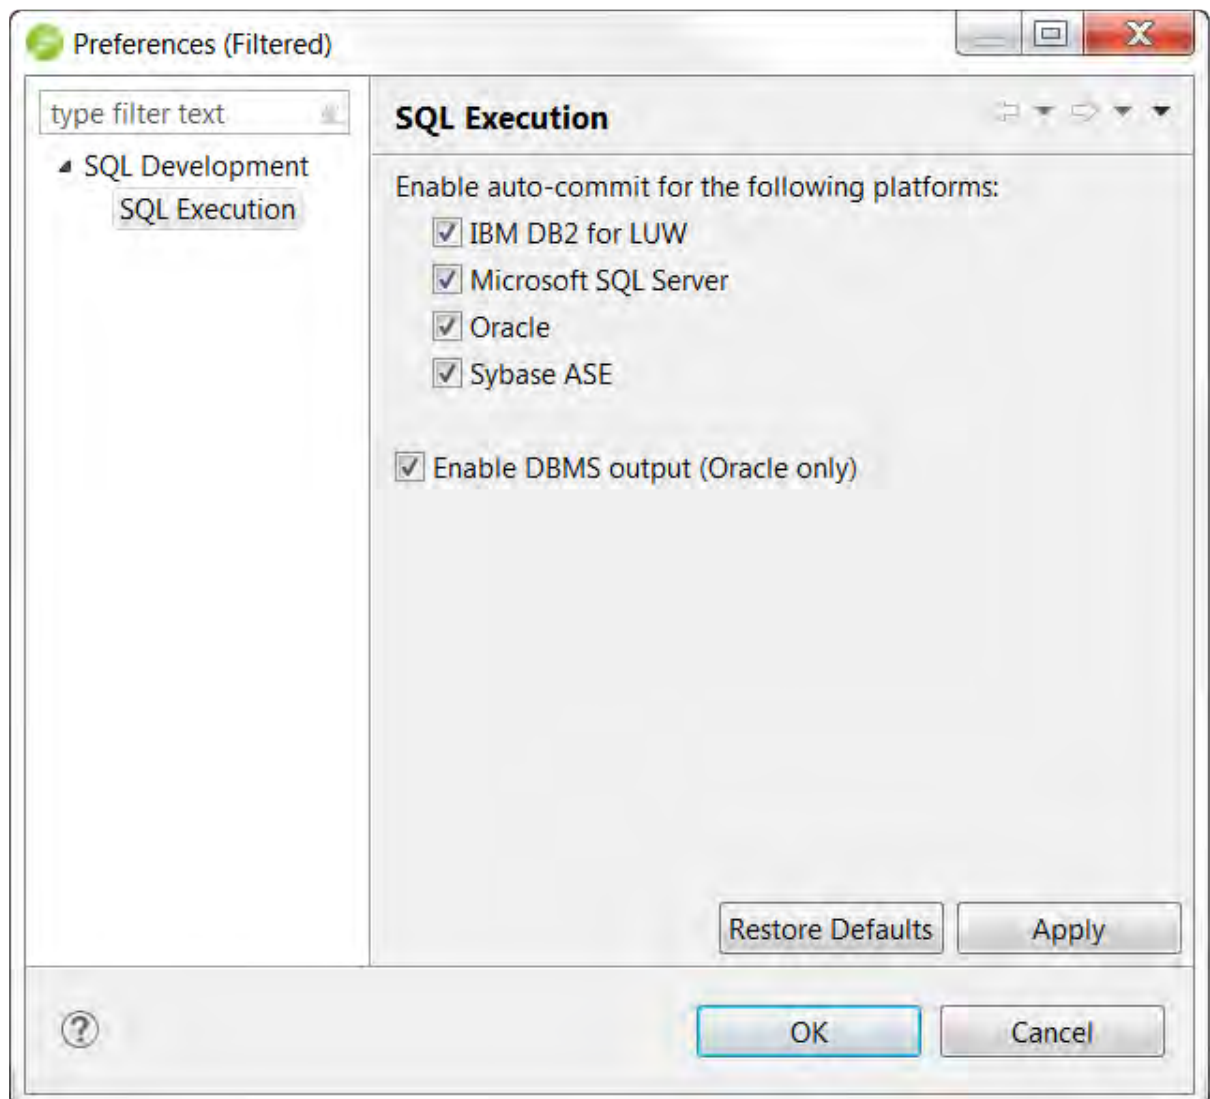
3. Use the tree view provided in the Element window to select the comment type or code element you want to modify. Select the options to the right-hand side of the window to modify it. The **Preview** window shows a piece of sample code that updates according to the changes you made.

Specify SQL Execution preferences

The SQL Execution preferences you set determine how SQL is executed.

To specify SQL Execution preferences

1. Select **Preferences > SQL Execution**.



2. Select from the following options, and then click **OK**.

- **Enable auto-commit for the following platforms:** When disabled, the SQL is executed within transactions that must be manually committed.
- **Enable DBMS output (Oracle only):** When disabled, this omits the output statements that Oracle would otherwise display.

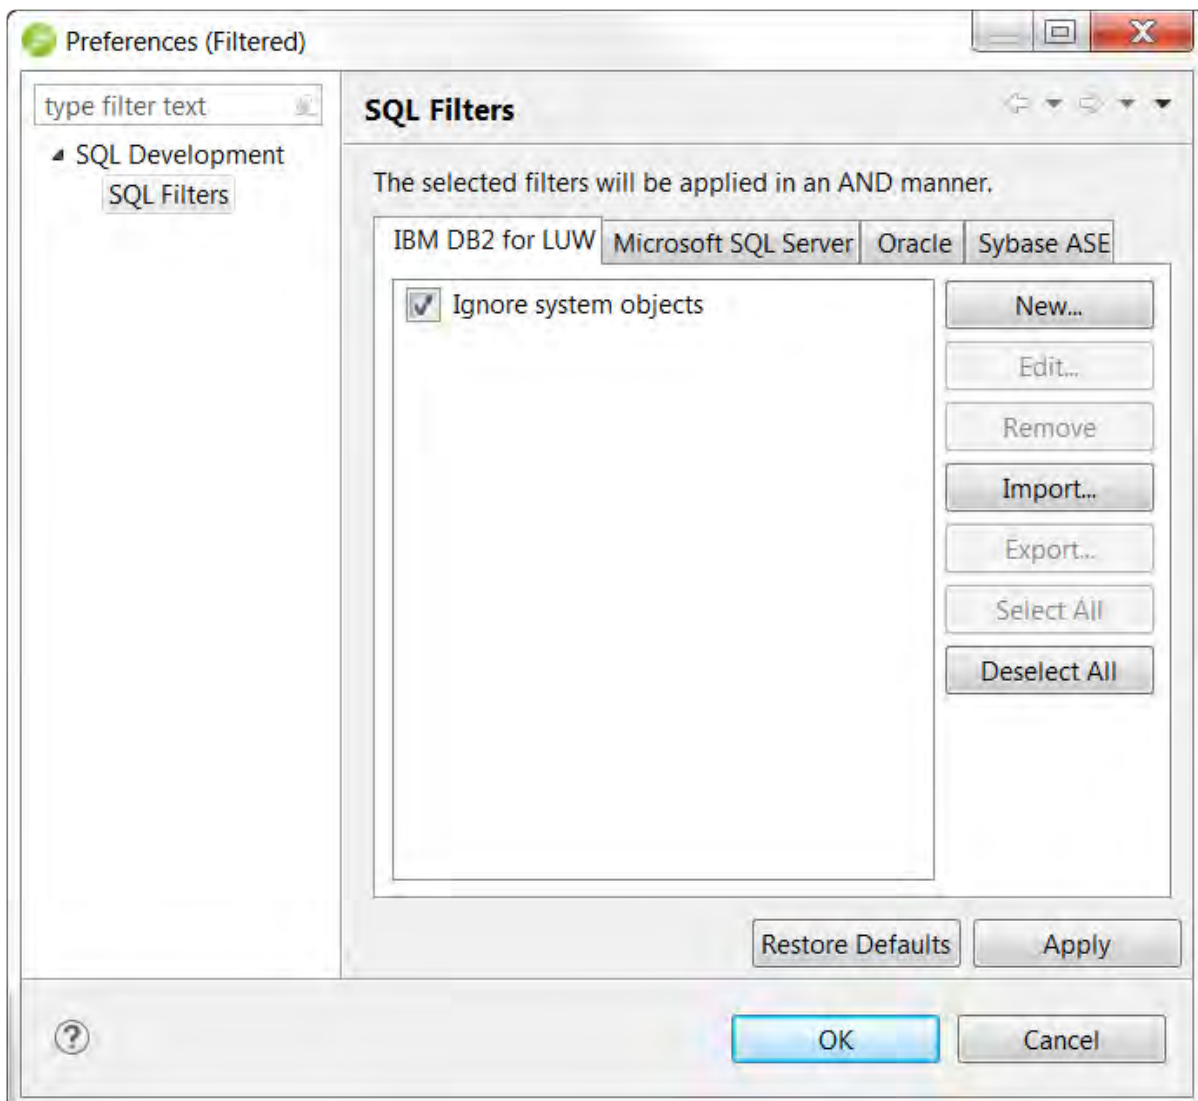
i If you disable auto-commit for a platform, you must use SQL Editor's transaction features to execute code on that platform.

Specify SQL Filters preferences

These are the set of controls that determine what objects are shown in the Data Source Explorer.

To specify SQL Filters preferences

Select **Preferences > SQL Filters**.



Specify Profile Alerts preferences

For information about specifying the Profile Alert preferences, see [Configuring profiling](#).

Specify Profile Repositories preferences

For information about specifying the Profile Repositories preferences, see [Configuring profiling](#).

Specify Tuning Job Editor preferences

For information about specifying the Tuning Job Editor preferences, see [Configuring tuning](#).

Specify VST Diagrams Tuning preferences

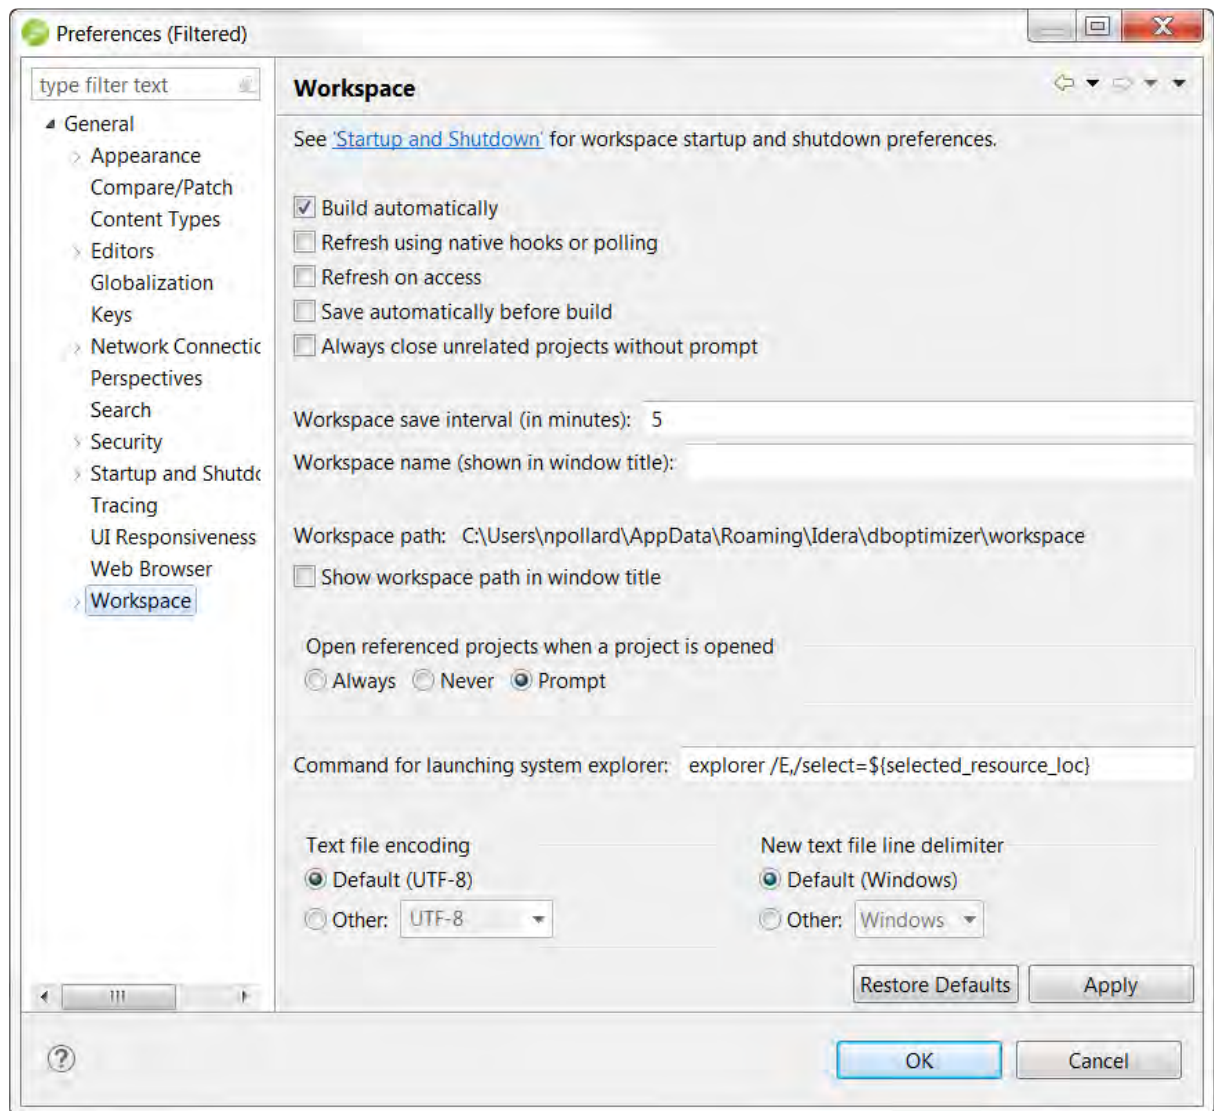
For information about specifying the VST Diagrams Tuning preferences, see [Configuring tuning](#).

Specify File Encoding preferences

The Workspace panel provides options for Unicode support in SQL files.

To specify File Encoding preferences

1. Select **Preferences > General**.
2. Expand **General**, and then click **Workspace** in the tree.



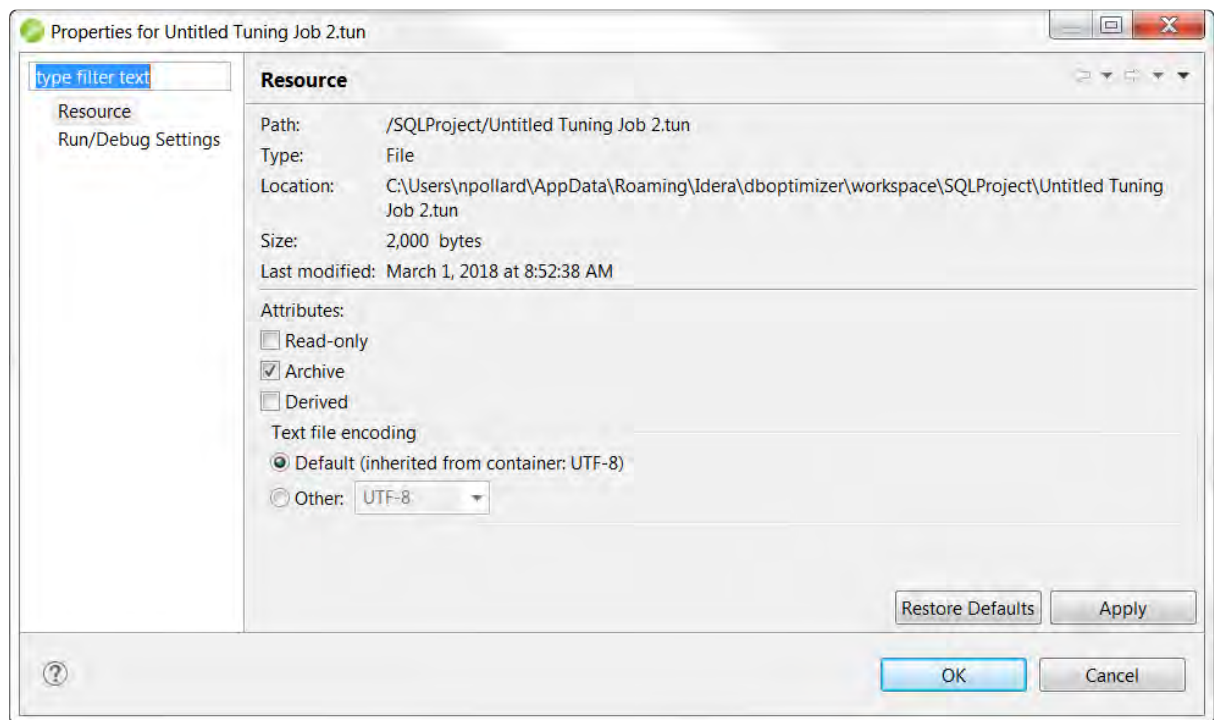
The default encoding for text files on Windows platforms is Cp1252. You can change Unicode support in from file to file using the Text File Encoding options available on the Workspace panel.

To change text file encoding in the development environment

1. Select **Preferences > General > Workspace**, and then click the **Other** option under **Text File Encoding**.
2. Use the drop-down menu to select an encoding mode from the list provided.
3. Click **Apply** to save your changes.

To change text file encoding on a specific, folder, or project

1. Right-click on the file, folder, or project that you want to modify, and then choose **Properties**.
2. Modify the **Text file encoding** selection on the **Resource** properties page that appears.



Adjusting Java heap size to optimize performance

In Java applications, the parameters **-Xms** and **-Xmx** are integral to Java Virtual Machine (JVM) memory management.

- **-Xms**. This parameter dictates the initial heap size allocated by the JVM when the application starts.
- **-Xmx**. This parameter sets the maximum heap size that the JVM can use for the application during its execution.

Properly configuring these parameters is crucial for optimizing the performance of Java applications, especially when handling extensive datasets or executing complex operations that demand significant memory resources.

Editing -Xms and -Xmx in the dboptimizer.ini

To adjust the Java heap size of DB Optimizer

1. Open File Explorer, and then navigate to **C:\Program Files\IDERA\DBOxxxx** where **xxxx** is your installed version number.
2. Open the Files with Administrator Rights. Right-click your chosen text editor (e.g., Notepad), and then select **Run as administrator** to ensure necessary privileges for editing files in the Program Files directory.
3. Look for the following lines in **ini/dboptimizer.ini** specifying JVM options. These lines are:
-Xms128M
-Xmx768M
4. Adjust the values of **-Xms** and **-Xmx** based on the available system resources. For example:
-Xms1024M
-Xmx4096M
 This configuration sets the initial heap size to 1024 megabytes and the maximum heap size to 4 gigabytes. Adjust these values according to your specific needs.

5. Save your changes made to the *ini/dboptimizer.ini* file.
6. After editing the file, restart DB Optimizer to apply the changes effectively.

 You can view the Heap status by checking the **Show Heap Status** checkbox from the application preferences (**Tools > Preferences** menu).

Introduction to database tuning

This discussion will help you understand the methodology behind DB Optimizer's tuning functionality and how you can use it to optimize database performance. This discussion is comprised of the following topics:

- [Introduction to DB Optimizer tuner](#)
- [SQL tuning methodology](#)
- [SQL tuner overview](#)
- [What's happening on the databases?](#)
- [Tuning example](#)
- [Finding and tuning problem SQL](#)

Introduction to DB Optimizer tuner

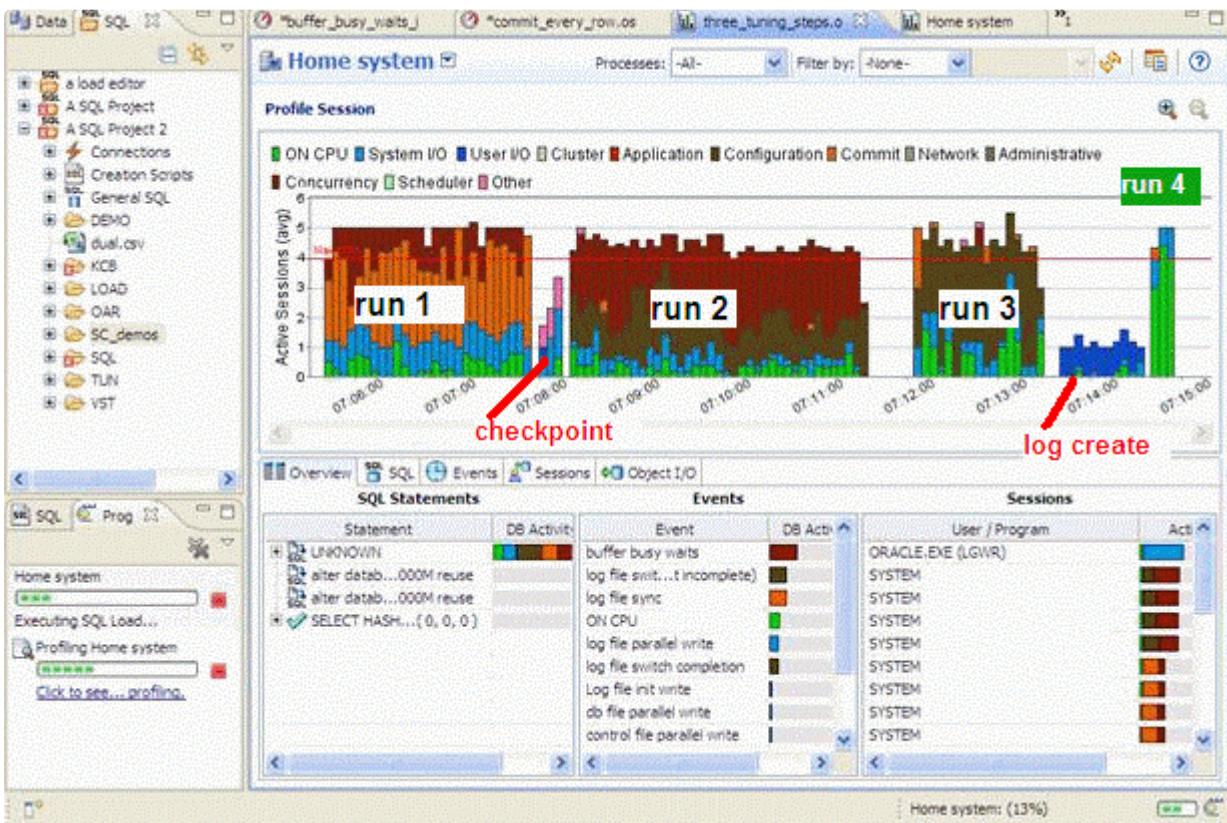
DB Optimizer's methodology grew out of the impossible predicament presented by the defacto method of database tuning. The standard method was trying to collect 100% of the statistics 100% of the time. Trying to collect all the statistics as fast as possible ends up putting load on the monitored database and creating problems. Stories of problems created by database monitoring products abound in the industry. In order to avoid putting load on the target database, performance monitoring tools have to collect less often as a compromise. Oracle compromised in 10g with AWR (their automated performance data collector), only running it once an hour because of the performance impact. Not only is the impact on the monitored target high, but the amount of data collected is staggering, but the worst problem of all though, is the impossibility of correlating statistics with the sessions and SQL that created the problems or suffered the consequences.

The solution to collecting performance data required letting go of the old problematic paradigm of trying to collect as many performance counters possible as often as we could and instead freeing ourselves with the simple approach of sampling session state. Session state includes what the session is, what its state is (active, waiting, and if waiting, what it is waiting on) and what SQL it is running. The session state method was officially packaged by Oracle in 10g when they introduced Active Session History (ASH). ASH is an automated collection of session state sampling. The rich robust data from ASH in its raw form is difficult to read and interpret. The solution for this was Average Active Sessions (AAS). AAS is a single powerful metric which measures the load on the database based on the ASH data. AAS data provided the perfect road map for what data to drill into. The main drill downs are "top SQL," "top session," "top event," and "top objects."

Other aggregations are possible based on the different dimensions in the ASH data.

Tuning example

Here is an example screen shot of the same batch job being run four times. Between each run performance modifications are made based on what was seen in the in the profiling load chart:



Run examples:

1. In run 1, the "log file sync" event is the primary bottleneck. To correct this, we moved the log files to a faster device. (You can see the checkpoint activity just after run 1 where we moved the log files.)
2. In run 2, the "buffer busy wait" event is the primary bottleneck. To correct this, we moved the table from a normal tablespace to an Automatic Segment Space Managed tablespace.
3. In run 3 the "log file switch" (checkpoint incomplete) event is the primary bottleneck. To correct this, we increased the size of the log files. (You can see the IO time spent creating the new redo logs just after run 3.)
4. The run time of run4 is the shortest and all the time is spent on the CPU which was our goal, take advantage of all the processors and run the batch job as quickly as possible.

i To view an explanation of the event, hover over the even name in the Event section.

Conclusion:

With the load chart we can quickly and easily identify the bottlenecks in the database, take corrective actions, and see the results. In the first run, almost all the time is spent waiting, in the second run we eliminated a bottleneck but we actually spent more time as the bottleneck was worse. Sometime this happens as eliminating one bottleneck causes great contention on the next bottleneck. You can see that the width of the run or the time it ran is wider in run 2. By the third run, we are using more CPU and the run time is faster and finally by the 4th run all the time spent is on CPU, no waiting, and the run is the fastest by far.

SQL tuning methodology

A common SQL tuning methodology is to:

1. Verify that the execution path is the optimal for the query. If not, either use the tuning directives (such as hints on Oracle) or identify why the native optimizer failed to pick the optimal path.

2. If the query is still slow, then look at adding indexes.
3. If the query is still slow, then you know you are going to have to look at the architecture.

Ask yourself:

- What information is the query trying to get?
- Is this information necessary?
- Are there alternative ways to get this information?

DB Optimizer's SQL Tuner can help with steps 1 and 2. Step 3 must be done by a developer or DBA, but knowing that steps 1 and 2 are already validated can indicate to management that step 3 is necessary and therefore allocate sufficient resources for step 3.

How do we know if the native database optimizer chose the optimal path? How long would it take to check this by hand?

DB Optimizer's SQL Tuner is a solid fast sanity test to verify the plan chosen by the native database SQL optimizer. Tuner quickly generates as many alternative paths as possible and allows the user to execute them to see if there are more efficient execution paths. DB Optimizer's SQL Tuner is successful at tuning queries that have a sub-optimal execution path.

A query has a sub-optimal execution path when the database optimizer has miscalculated the cost of the various possible access paths and mistakenly chosen a bad path. The access path calculations are sometimes miscalculated because of the following reasons:

- **The table/index statistics are missing or wrong.** For example, the number of rows is missing or way off.
- **The data is skewed, for example, the number of orders with an open status is usually low compared to all the orders that have a closed status because the work is complete.** For example, orders get filled every day, but only a few are open and needing to be processed. Looking for open orders should probably use an index and return fewer rows than looking for closed orders which should probably just do a full table scan.
- **The predicates used are correlated.** The optimizer treats two predicate filters on a table as more selective than just one, but this is not always the case such is the case in the query, how many Swedes speak Swedish which basically returns the same number of results as just asking for the number of Swedes alone. Another example is how many Swedes speak Swahili, which is probably more selective than the optimizer would guess.
- **There is a bug in the optimizer.**

DB Optimizer's SQL tuner takes a query and tries to produce as many execution paths as possible. These alternative execution paths can then be run to see if there is a faster or less resource expensive execution path. The execution of each alternative case is timed and if the execution exceeds 1.5 X the original case then its execution is stopped and we move on to the next case. This avoids wasting time and resources on execution plans that are clearly sub-optimal.

SQL tuner overview

Tuning provides an easy and optimal way to discover efficient paths for queries that may not be performing as quickly or as efficiently as they could be.

The application enables the optimization of poorly-performing SQL code through the detection and modification of execution paths used in data retrieval. This process is performed through the following functions:

- Hint Injection
- Index Analysis
- Statistic Analysis (Oracle only)
- Query re-writes such as suggesting joins to eliminate Cartesian joins, adding transitivity predicates, and unnesting subqueries in the WHERE clause.

Tuning analyzes an SQL statement and supplies execution path directives to the application that encourage the database to use different paths.

For example, if tuning is selecting from two tables (A and B), it will enable the joining of A to B, or B to A as well as the join form. Additionally, different joining methods such as nested loops or hash joins can be used and will be tested, as appropriate.

Tuning selects alternate paths and enables you to change the original path to one of the alternates. Execution paths slower than the original are eliminated, which enables you to select the quickest of the returned selections and improve query times, overall.

This enables the DBA to correctly optimize queries in the cases where the native optimizer failed.

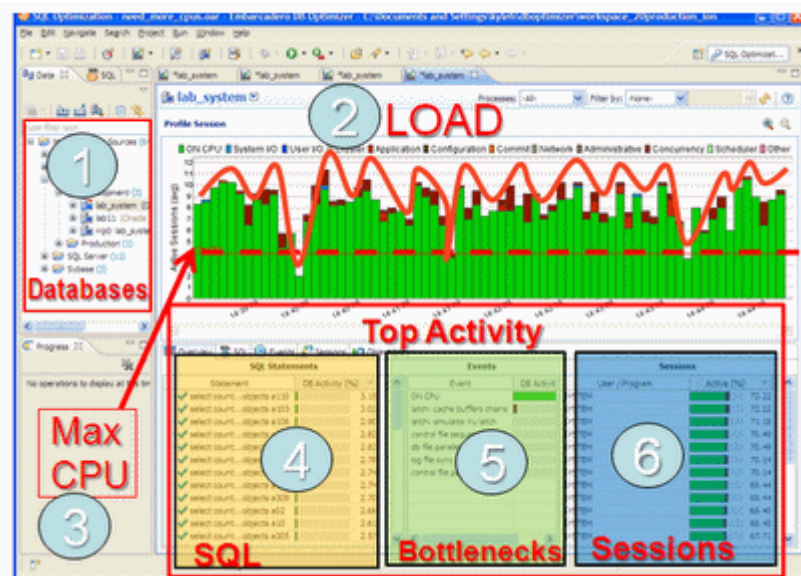
What's happening on the databases?

Is the database idle, working, or bottlenecked?

When a bottleneck happens how can you know which of these problems are causing the problem? A bottleneck could be caused by:

- An application problem
- An undersized machine
- SQL requiring optimization
- A misconfigured database

All of these can be easily identified from DB Optimizer's performance profiling screen. Let's look at the components of the performance profiling screen.



The screen has the following important areas.

1. **Databases.** For more information, see [Databases](#).
2. **Average Active Sessions (AAS) Load of selected database.** For more information, see [Average Active Sessions \(AAS\) Load of selected database](#).
3. **Maximum CPU line.** For more information, see [Maximum CPU line](#).
4. **Top SQL.** For more information, see [Top SQL](#), [Top Bottlenecks](#), and [Top Sessions](#).
5. **Top Bottlenecks.** For more information, see [Top SQL](#), [Top Bottlenecks](#), and [Top Sessions](#).
6. **Top Sessions.** For more information, see [Top SQL](#), [Top Bottlenecks](#), and [Top Sessions](#).

Databases

First, on top left, is a list of our databases we have registered.

Average Active Sessions (AAS) Load of selected database

The most important part of the screen is the Average Active Sessions (AAS) graph. AAS shows the performance of the database measured in the single powerful unified metric AAS. AAS easily and quickly shows any performance bottlenecks on the database when compared to the Maximum CPU line. The Max CPU line is a yardstick for performance on the database. When AAS is larger than the Max CPU line there is a bottleneck on the database. Bottleneck identification is that easy.

AAS or the average number of sessions active, shows how many sessions are active on average (over a 5 second range in DB Optimizer) and what the breakdown of their activity was. If all the users were running on CPU then the AAS bar is all green. If some users were running on CPU and some were doing IO, represented by blue, then the AAS bar will be partly green and partly blue.

Maximum CPU line

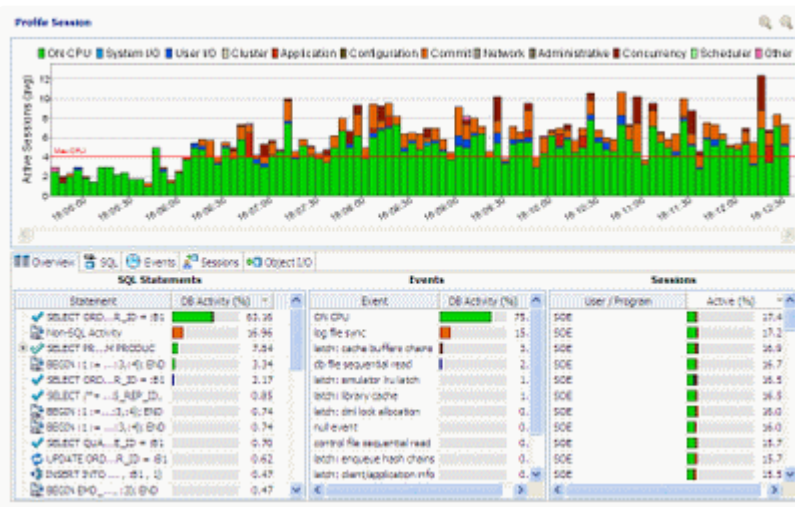
The line "Max CPU" represents the number of CPU processors on the machine. If we have one CPU then only one user can be running on the CPU at a time. If we have two CPUs then only 2 users can be on CPU at any instant in time. Of course users can go on and off the CPU extremely rapidly. When we talk about sessions on the CPU we are talking about the average number of sessions on CPU. A load of one session on the CPU thus would be an average which could represent one user who is consistently on the CPU or many users who are on the CPU for short time periods. When a CPU becomes a resource bottleneck on the database we will see the average active sessions in CPU state go over the Max CPU line. The number of sessions above the max CPU line is the average number of sessions waiting for CPU.

The Max CPU is a yardstick for performance on the database.

From looking at the previous chart the problem is a machine resource problem.

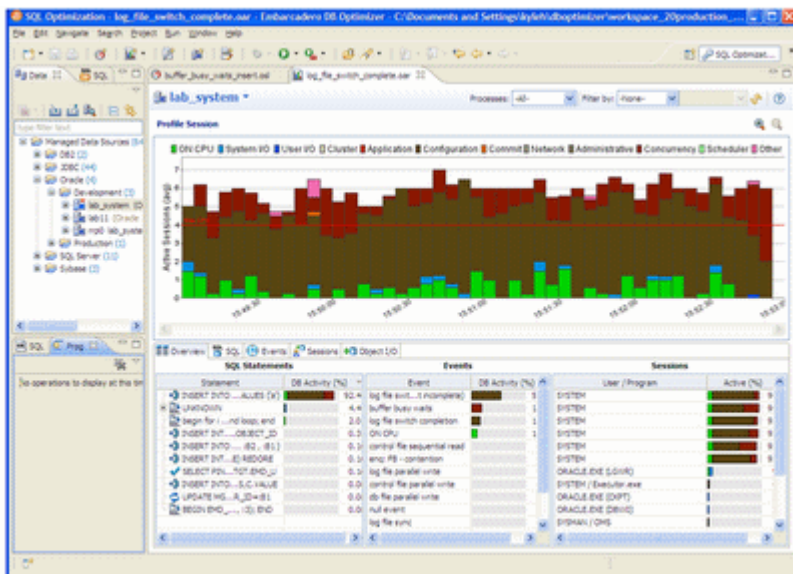
Top SQL, Top Bottlenecks, and Top Sessions

In order to know what the problem is, we have to find out where that demand is coming from. To find out where the demand is coming from we can look at Top SQL and Top Session tables below the load chart. In our case shown here the load is well distributed over all SQL in Top SQL and all sessions in Top Session. There is no outlier or resource hog. In this case it's the machine that's underpowered. What does a case look like where we should tune the application? The following screenshot depicts such a problem.



In this case, again the CPU demand is more than the machine can supply but if we look at "Top SQL" we can see that the first SQL statement (with the large green bar) uses up much more CPU than any of the rest, actually 60%! If we could get it down to 10% CPU then we'd save 50% of the CPU usage on the machine! Thus in this case it's worth our while to spend a day or week or even a couple weeks trying to tune that one SQL statement instead of buying a bigger machine.

Finally, how do we know when the database configuration is a problem? We know it's a configuration problem when we are seeing something other than CPU as the bottleneck in Top Bottleneck section. Here's an example



In this case we can see the load is higher than the Max CPU line but the load is coming from brown colored bars and the green CPU colored bars. If we look at Top SQL we see that there is only one SQL taking up almost all the load, but it's not because of CPU which would be a green bar, but some other color. What does this other color represent? We can look at the Top Bottleneck section and see that it is "log file switch (incomplete)" which basically means the log files are too small, the database is not correctly configured. This bottleneck can be resolved simply by increasing the log size.

Tuning example

This example is comprised of the following parts:

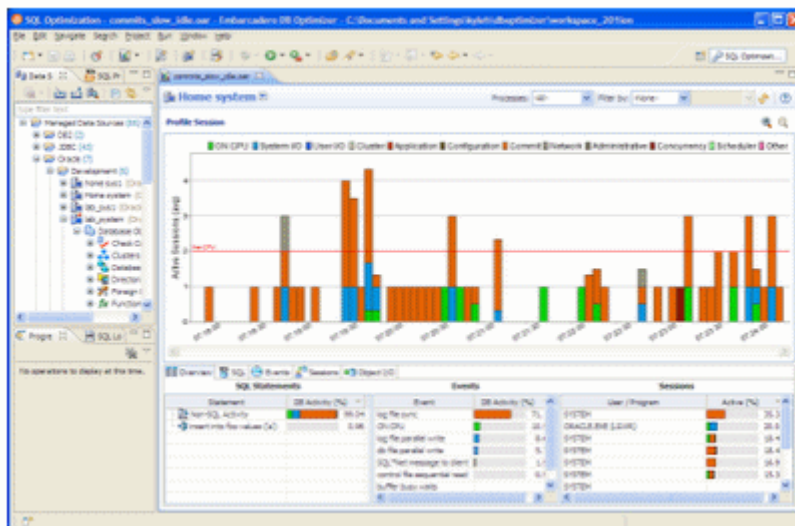
- The database is hanging or the application has problems
- The database caused the problem
- The machine caused the problem

The database is hanging or the application has problems

I wonder if you can imagine, or have had the experience of the application people calling with anger and panic in their voices saying, "The database is so slow, you've got to speed it up."

What's your first reaction? What tools do you use? How long does it take to figure out what's going on?

Let's take a look at how it would work with DB Optimizer.



We can clearly see that the database is not bottlenecked and there must be a problem on the application.

Why do we think it's the application and not the database? The database is showing plenty of free CPU in the load chart, the largest chart, on the top in the image above. In the load chart, there is a horizontal red line. The red line represents the number of CPUs on the system, which in this case is two CPUs. The CPU line is rarely crossed by bars which represent the load on the database, measured in average number of sessions. The session activity is averaged over five samples over five seconds, thus bars are five seconds wide. The bars above fall mostly about one average active session and the bars are rarely green. Green represents CPU load. Any other color bar indicates a sessions waiting. The main wait in this case is orange, which is log file sync, waits for commits. Why is the database more or less idle and why are most of the waits we do see for "commit"? When we look at the code coming to the database we see something like this:

insert

into

foo

```
values
```

```
('a');
```

```
commit;
```

insert

into

foo

```
values
```

```
('a');
```

```
commit;
```

insert

into

foo

```
values
```

```
('a');
```

```
commit;
```

insert

into

foo

```
values
```

```
('a');
```

```
commit;
```

insert

into

foo

```
values
```

```
('a');
```

```
commit;
```

insert

into

foo

```
values
```

```
('a');
```

```
commit;
```

insert

into

foo

```
values
```

```
('a');
```

```
commit;
```

Doing single row inserts and committing after each is very inefficient. There is a lot of time wasted on network communication which is why the database is mainly idle. When the application thinks it's running full speed ahead, it is actually waiting mainly on network communication and commits. If we commit less and batch the work we send to the database, reducing network communications, we will run much more efficiently. Change the code to:

begin

for i

in

1 . . 10000

loop

insert

into

foo

values

```
('a');
```

```
-
```

```
commit;
```

```
end
```

```
loop;
```

```
end;
```

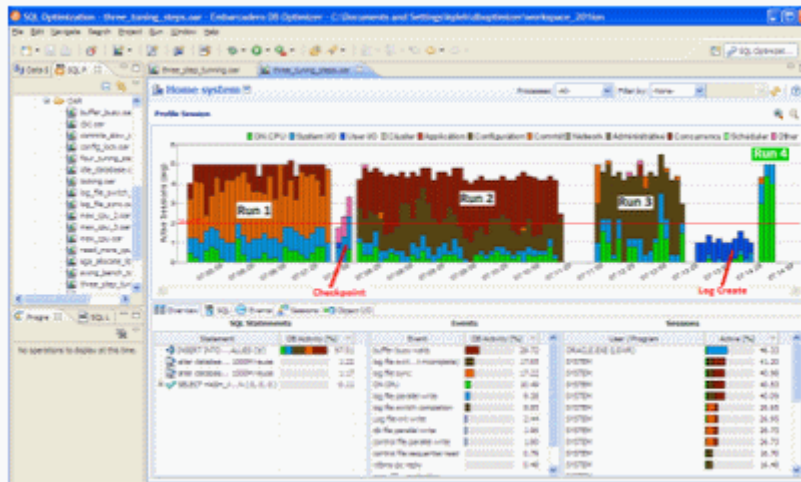


commit



This change improves the communication delay and now we get a fully loaded database but we run into database configuration issues.

The database caused the problem



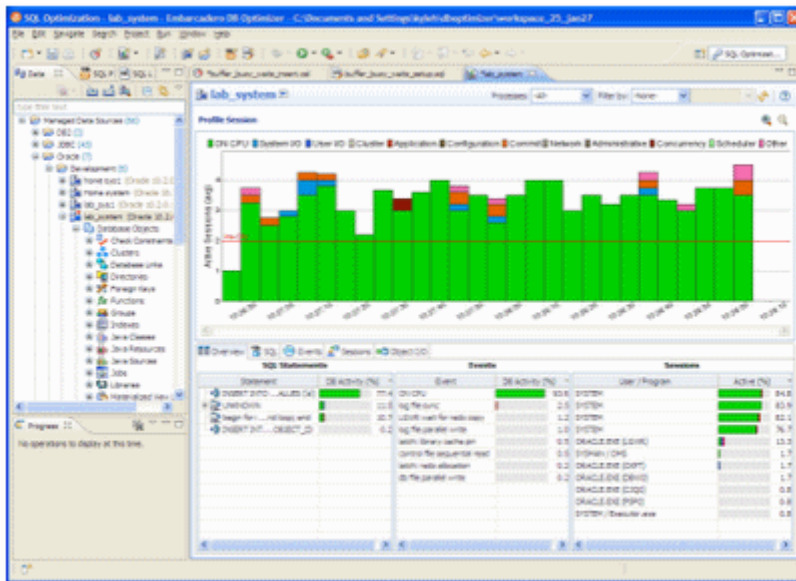
In the above DB Optimizer screen, the same workload ran four times. We can see that the time (width of the load) reduced, and the percent of activity on CPU increased.

Runs:

1. In Run 1, the orange "log file sync" area is the largest, which means users are waiting on commits even though we are committing less in the code. In this case we moved the log files to a faster device. You can see the checkpoint activity just after run 1 where we moved the log files.
2. In Run 2, the dark red "buffer busy wait" area is the largest. We drilled down on the buffer busy wait event in the Top Event section and the details tell us to move the table from a normal tablespace to an Automatic Segment Space Managed tablespace.
3. In Run 3, the dark brown "log file sync (checkpoint incomplete)" area is the largest, so we increased the size of the log files. You can see the IO time spent creating the new redo logs just after Run 3. The run time is the shortest and all the time is spent on the CPU which was our goal, to take advantage of all the processors and run the batch job as quickly as possible.
4. The run time is the shortest and all the time is spent on the CPU, which was our goal to take advantage of all the processors and run the batch job as quickly as possible.

The machine caused the problem

Now that the application is tuned and the database is tuned, let's run a bigger load:



We can see that the CPU load is constantly over the max CPU line. How can we have a bigger CPU load than there are actually CPUs on the machine? Because the demand for CPU is higher than the CPU available on the machine.

In the image above, there are two CPUs on the machine but an average of three users think they are on the CPU, which means that on average one user is not really on the CPU but ready to run on the CPU and waiting for the CPU.

At this point we have two options. In this case we are only running one kind of load, the insert. For inserts we can actually go even further tuning this insert and use Oracle's bulk load commands:

declar

e

TYPE

IDX IS

TABLE

OF

Integer

INDEX

BY

BINARY_

INTEGER

;

MY_IDX

IDX;

begin

for i

in

1 . . 8 0 0

0 2 0 0 p

MY_IDX

(i) := 1

;

end

loop;

FORALL

indx

IN

MY_IDX

• FIRST

• •

MY_IDX

. LAST

INSERT

INTO

foo

(dumm

y)

VALUES

(MY_I

DX(ind

x)) ;

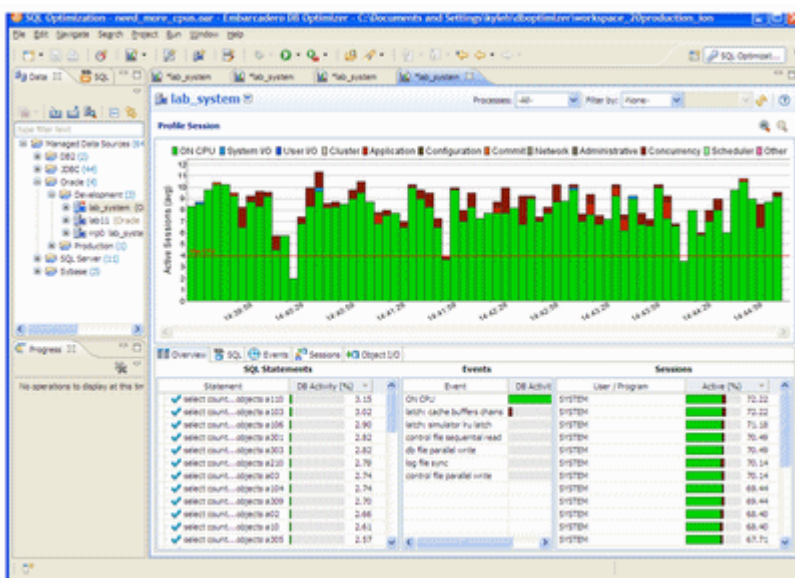
COMMIT

;

end;

/

But if this was an application that had a lot of different SQL and the SQL load was well distributed across the system then we'd have a case for adding more hardware to the system. Making the decision to add more hardware can be a difficult decision because in general the information to make the decision is unknown, unclear or just plain confusing, but DB Optimizer makes it easy and clear, which can save weeks and months of wasteful meetings and debates. For example:

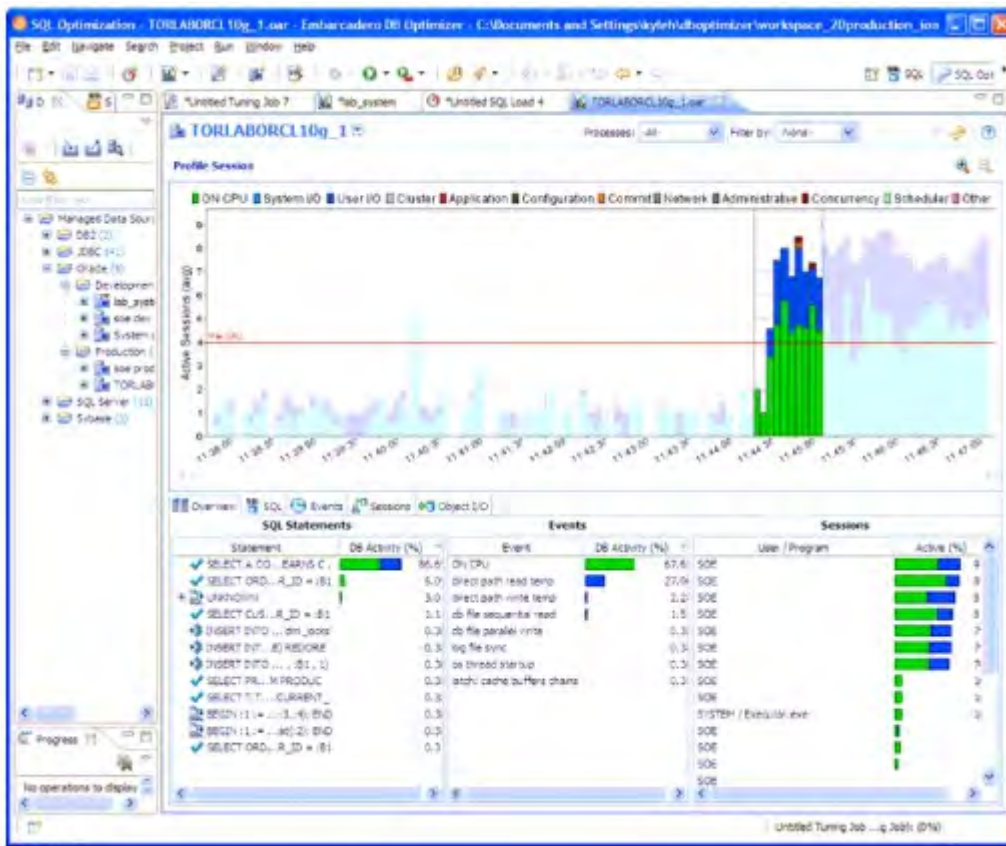


If we look in the bottom left, there is no SQL that takes up a significant amount of load, there is no outlier SQL that we could tune and gain back a lot of wasted CPU. We'd have to tune many SQL and make improvements on most of them to gain back enough CPU to get our load down below the max CPU line. In this case, adding CPUs to the machine might be the easiest and most cost effective solution.

Conclusion:

- With the load chart we can quickly and easily identify the bottlenecks in the database, take corrective actions, and see the results. In Part 1, we had an application problem, in Part 2 we had three database configuration issues, and in Part 3 we had a hardware sizing issue. In all three instances DB Optimizer provides a clear and easy presentation of the data and issues making solutions clear.

Finding and tuning problem SQL



DB Optimizer is targeted at finding problem SQL in a running load with the profiler and then tuning that (or those) specific queries with the tuner.

It's not efficient just to dump a bunch of procedure code into the tuner and then try and see if any of the SQL in the package or procedure is tunable. Most queries should, by default, run optimally on a database, so the goal of DBO is to tune those queries that for one reason or another are not optimally tuned by the database by default. The easiest way to find those queries is to identify them on a running system. They can be identified on a running system because they take up a lot of resources. If we find a resource intensive query, then it's worth the time to generate cases and analyze it for missing indexes to see if there is a way to tune it.

Using DB Optimizer

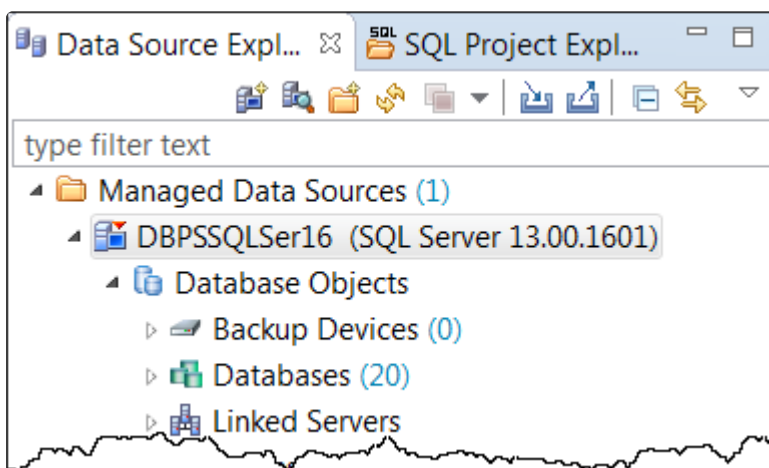
This section describes how to use the features of DB Optimizer to optimize your database operations. This section contains an overview of DB Optimizer functionality and also contains detailed instructions for:

- [Working with data sources](#)
- [Working with SQL projects](#)
- [Creating and editing SQL files \(SQL Editor\)](#)
- [Executing SQL files](#)
- [Troubleshooting](#)

Working with data sources

The Data Source Explorer provides a tree view of all registered data sources and associated database objects. When you first start DB Optimizer, a prompt appears and offers to populate Data Source Explorer from multiple sources on the system. This includes previously-registered data sources on other IDERA products, and third-party DBMS clients such as TOAD. If DB Optimizer cannot detect a data source, you can register it manually.

Additionally, you can initiate this feature by clicking the Auto-Discovery button on the Toolbar or via the **File > Import > IDERA > Data Sources > Previously Registered IDERA Data Sources (Registry)** command from the Main Menu.



The Profiling Repository entries in the Data Source Explorer are available only when configured in the **Profile Configuration Dialog** for Oracle data sources only. These are saved profiling sessions that you can share with other DB Optimizer users. For information on configuring the data source profiles, see [Building profiling configurations](#).


- [Register data sources](#)
- [Add a new data source](#)
- [Import and export data sources](#)
- [Categorize data sources](#)
- [Customizing data source categories](#)
- [Browse a data source](#)
- [View database object properties](#)
- [Search for database objects](#)
- [Filter database objects](#)
- [Drop a database object](#)

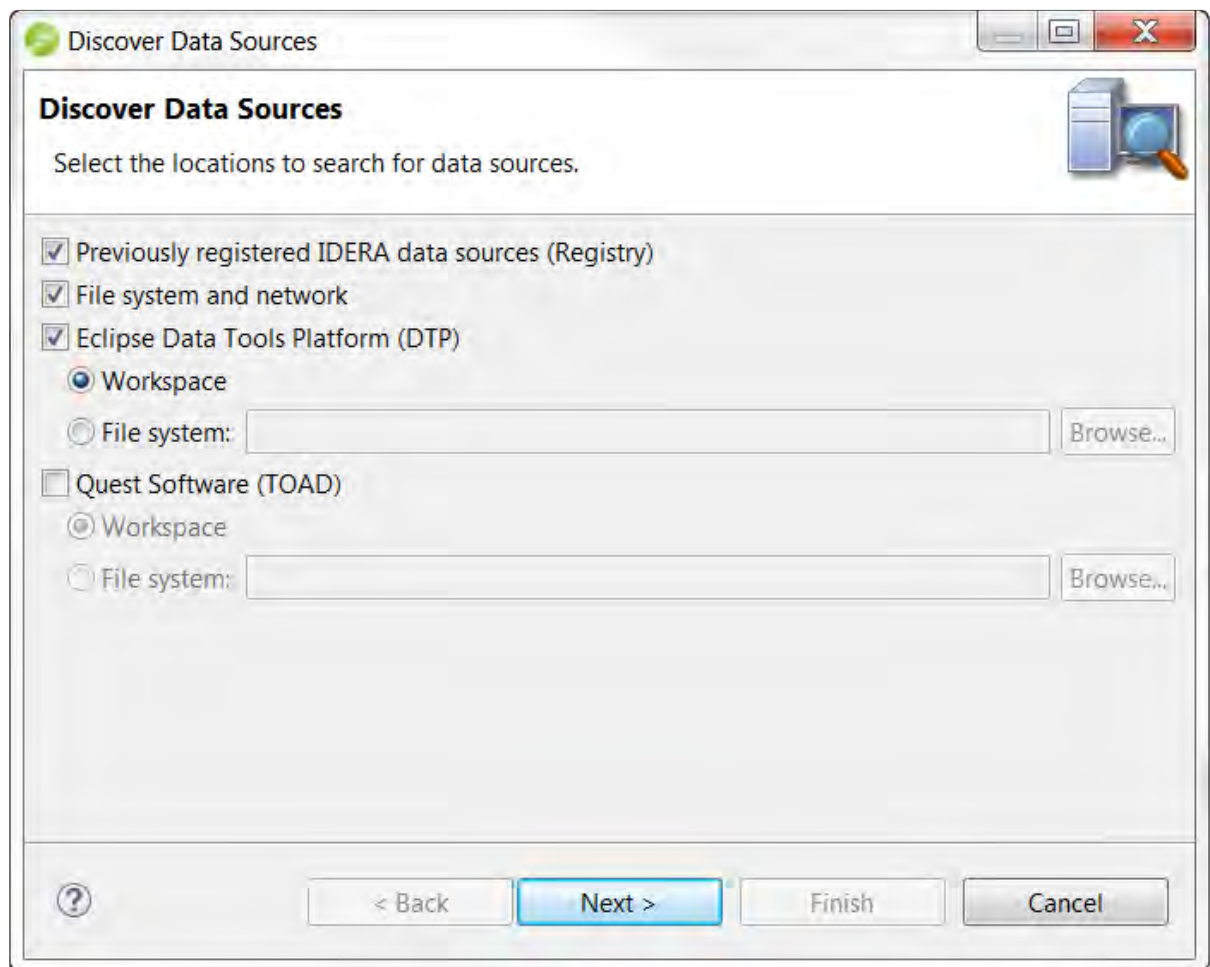
Register data sources

When DB Optimizer is started, it prompts you to discover data source catalogs that have been created by any previously-installed IDERA products (DBArtisan, Rapid SQL, etc.), or other instances of DB Optimizer.

Additionally, the system scans your machine for the client software of all supported third-party DBMS platforms (TOAD, Eclipse Data Tools Platform, etc.). These data sources are automatically added to the data source catalog.

To manually initiate the scan later

1. Click the Discover Data Sources icon  at the top of Data Source Explorer. The Discover Data Sources dialog appears.



2. Choose the type of data sources for which you want to scan, and then click **Next**. The wizard automatically returns all data sources it finds on your machine based on the criteria you specified.
3. Choose the data sources you want to add to the DB Optimizer environment, and then click **Finish**.


Data Source Explorer automatically populates with the new data source selections.

i To add data sources manually, right-click **Managed Data Sources** in the Data Source Explorer tree, select **New > Data Source**, and then enter the connectivity parameters as prompted.

For additional information on data source connection parameters, see [DBMS connection parameters by platform](#).

Once registered, the data source appears in the Data Source Explorer view. If you have created more than one workspace, they all share the same data source catalog.

Once a data source is registered, the connection parameters are stored locally. In some cases, a user ID and password are required to connect to a registered data source. DB Optimizer can encrypt and save user IDs and passwords to connect automatically.

 In some cases, older versions of DB Optimizer and DBArtisan/Rapid SQL are not compatible with this version of DB Optimizer, and the methods listed above will not import these older data source catalogs. If you are experiencing difficulties, you can import the old data sources via the Windows registry by selecting **File > Import > IDERA > Data Sources > Previously Registered IDERA Data Sources (Windows Registry)**.

Add a new data source

You can also add a new data source manually.

To manually add a new data source

1. Click **File > New > Data Source**.

New Data Source: New data source

Register a new data source

Choose the server type and location for your new data source.

Generate an unique name based on the alias or the host name

Data source name:

Select a server type:

- Generic JDBC
- IBM DB2 for LUW
- Microsoft SQL Server
- Oracle**
- Sybase ASE

Select a data source group:

- Managed Data Sources

Create a new server type-specific subgroup


Select a category:



[Configure data source categories](#)


? < Back Next > Finish Cancel

- If you want the system to assign a name to the data sources, select **Generate a unique name based on the alias or the host name**.
If you want to enter the data source name, deselect **Generate a unique name based on the alias or the host name**, and then in the **Data source name** field, enter the data source name.
- In the **Select a server type** area, select the type you want to add.
- In the **Select a data source group** area, select the data source group where you want the new data source to appear in the Data Source Explorer.


- If you want to assign the new data source to a category (useful if you have a large number of data sources to manage), click the **Select a category** list, and then choose one of the categories.



Categorized data sources appear with the color for the designated category on the bottom left of the data source icon in the bottom of such as  for the **Test** category and  for the **Production** category.

 For information on adding custom categories, see [Customizing data source categories](#).

When you open a tuning job or SQL Editor window to create SQL for the categorized data source you will see that the category color is applied to the top of the window, as follows.



- Click **Next**.

New Data Source: New data source

Configure a new Oracle data source

Enter Oracle-specific connection and security credential information for the new data source.

Oracle

Use a TNS name alias

Server alias:

Use connection descriptor

Use a direct connection

Host: *

Port: 1521

Type: Service name SID

Service/SID name:*

Instance name:

Security credentials

User name:

Password:

Connect as:

Auto-connect (Saves and encrypts password)

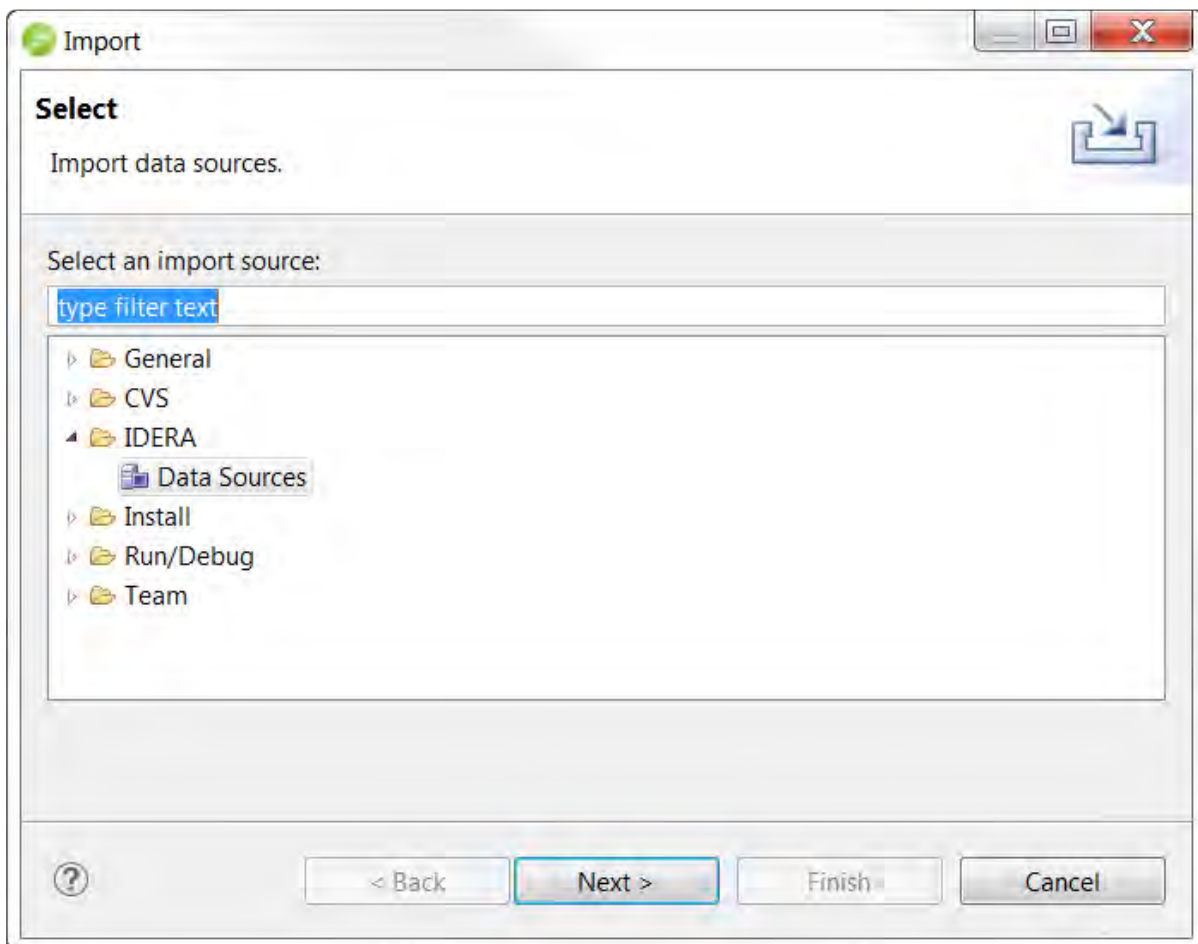
Connect using Windows Authentication

7. Complete the data source configuration and then click **Test Connection**. This will ensure your configuration is correct. If the connection test fails, make the necessary corrections, and then click **Finish**. The new data source appears in the **Data Source Explorer**.

Import and export data sources

Some IDERA products contain data source catalogs that are shared with DB Optimizer. In other words, instead of manually adding data sources to the environment, you can import an existing data source catalog from other IDERA products or third-party DBMS sources.

You can import Data Sources via the **File > Import** command in the Main Menu, expanding the IDERA folder in the Data Source Selection tree, and then choosing **Data Sources**.



The following types of sources can be imported to Data Source Explorer:

- Eclipse Data Tools Platform (DTP)
- Previously Registered IDERA Data Sources (File)
- Previously Registered IDERA Data Sources (Registry)
- Quest Software (TOAD)

Once a data source is registered, it automatically appears in Data Source Explorer. Connection parameters are stored locally, and DB Optimizer can be set to connect automatically each time you select the data source from the tree.

Conversely, you can also Export your current data source catalog to a file, which can then be imported into other instances of DB Optimizer via the **Previously Registered IDERA Data Sources (File)** option. This is performed using

the **File > Export** command in the Main Menu, and then selecting **IDERA > Data Sources** from the tree view in the **Export** dialog.



To import data sources

1. Select **File > Import**. The **Import** dialog appears.
2. Choose **IDERA > Data Sources** from the tree, and then click **Next**.
3. Choose a source from which you want to import the data sources. You can choose to import data sources from the DTP, TOAD, or an existing IDERA data source catalog stored in the Windows registry or as a file (created via the **Export** command).
4. Click **Next**.
5. Specify the location of the import source, and then click **Finish**. Data Source Explorer is automatically populated with the new data sources.

To export data sources

1. Select **File > Export**. The **Export** dialog appears.
2. Choose **IDERA > Data Sources** from the export tree and click **Next**.
3. Use the check boxes beside each listed data source to indicate which data sources you want to export., and then click **Next**.
4. Click **Finish**. The data sources are automatically exported in the form of an XML file. You can import this file to other instances of DB Optimizer via the **Import** command.

Categorize data sources

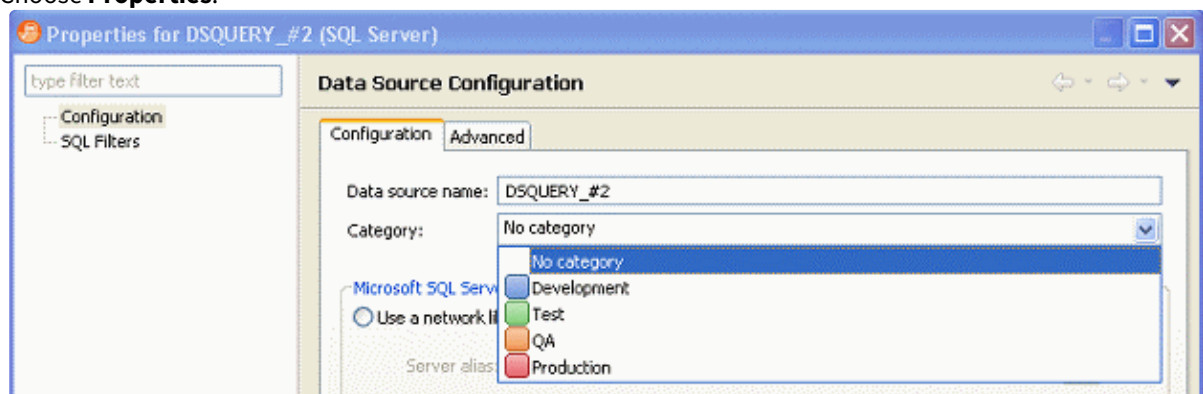
To make managing a large number of databases easier, you can assign a category to a data source. Categorized data sources appear with the color for the designated category on the bottom left of the data source icon in the bottom of such as  for the **Test** category and  for the **Production** category.

 For information on adding custom categories, see [Customizing data source categories](#).

You can categorize a data source when you add a new data source (see [Add a new data source](#)) or by editing the properties of an existing data source.

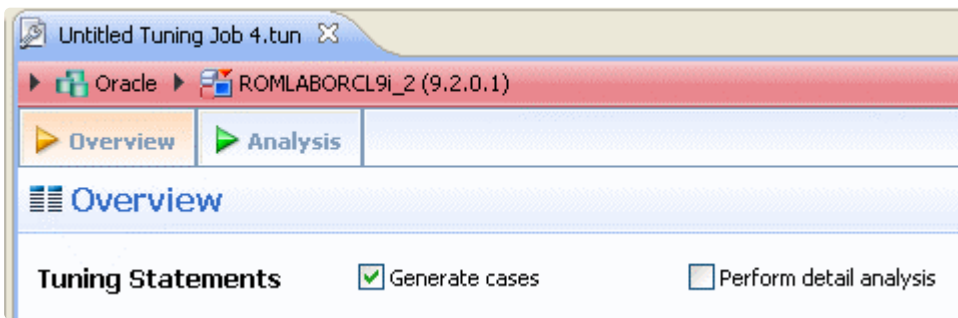
To edit properties for an existing data source

1. In the **Data Source Explorer**, locate and then right-click the data source you want to add to a category.
2. Choose **Properties**.



3. From the **Category** list, choose the category you want, and then click **OK**.

When you create a tuning job for the categorized data source you will see that the category color is applied to the top of the tuning job data source details.




Customizing data source categories

DB Optimizer lets you customize your data source category scheme. A data source category has the following configurable components:

- **Category name.** The name displayed as a selection when selecting a category
- **Short Name.** An abbreviation shown in window components and icons
- **Color.** The color used to denote a categorized data source in the Explorer tree icons and window tabs.

To customize your datasource categories

1. Select **Preferences > Data Source**, and then select the **Datasource Group** panel.
2. Take one of the following actions:
 - Create a new category by clicking **Add** and selecting or providing a **Full name**, **Short name**, and **Color** combination.
 - Edit an existing category by selecting the category, clicking **Edit** and modifying the name and color combination.

 The short name for a category cannot be edited.

- Delete an existing category by selecting the category, clicking **Delete** and verifying the deletion at the prompt.

Browse a data source

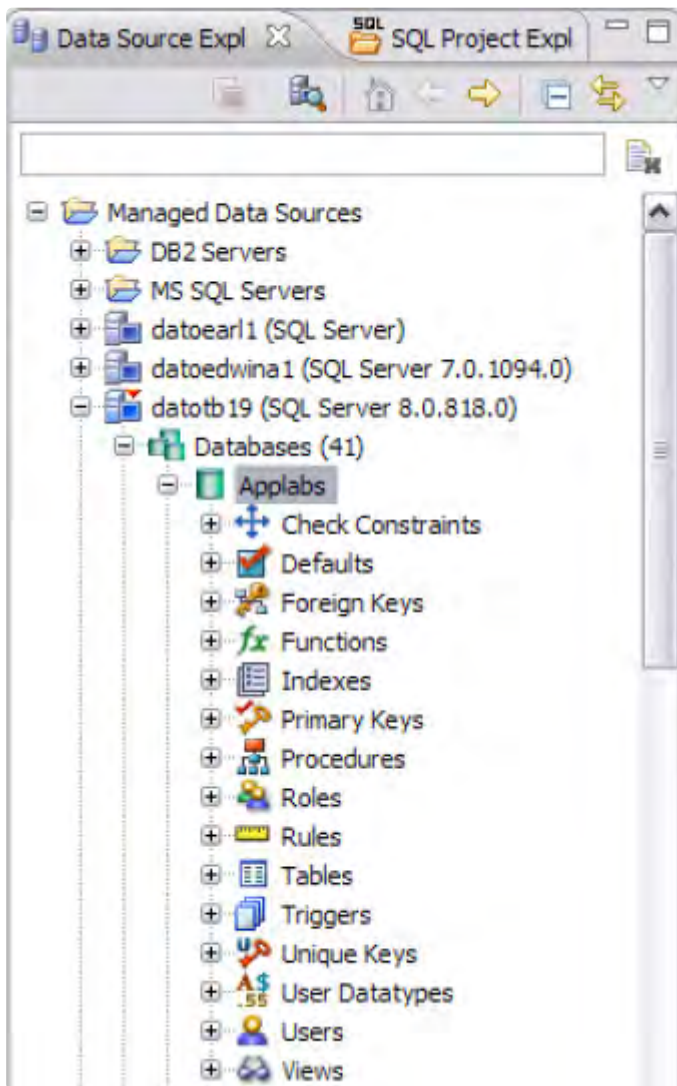
You can drill down in the Data Source Explorer tree to view registered databases on a server, and view tables, and other objects in a database. Additionally, you can view the structure of individual objects such as the columns and indexes of a table. Right-click the object for a menu of available commands, such as Extract to Project, which creates a new SQL file containing the object's DDL.

In most cases, whenever you browse a data source, DB Optimizer requires login information in order to connect with the data source. Enter a valid user name and password in the fields provided. The Auto Connect option retains your login credentials for future connections to the same data source.

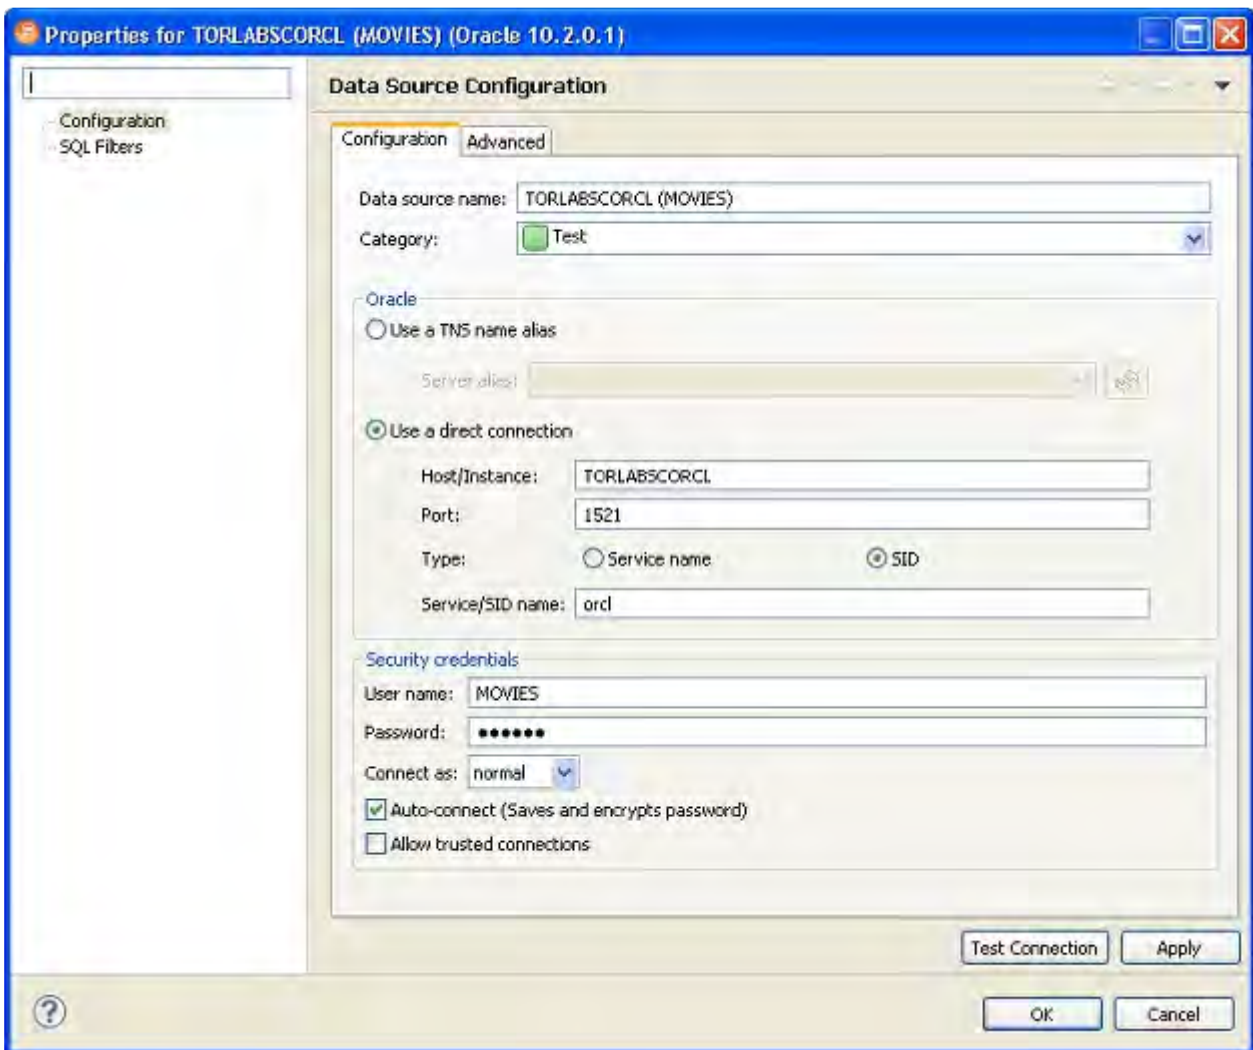
You can turn off the Auto Connect feature by right-clicking on a specified data source and toggling the Connect on Expand option. By default, when Connect on Expand is active, DB Optimizer automatically attempts to connect to the server each time you browse a data source.

View database object properties

All objects in Data Source Explorer contain properties as they relate to the DB Optimizer application.



DB Optimizer Object Properties are viewed via the Properties dialog. The dialog is accessed by right-clicking the object in Data Source Explorer.



To view Data Source Explorer object properties

Right-click a data source in Data Source Explorer. The dialog displays properties with regards to Configuration and SQL Filters.

The Configuration node provides information about the parameters used to initially define the data source during the data source registration process. For more information on these values and how to modify them, see [Register data sources](#).

The SQL Filter node enables a developer to place filters on data source objects that appear in the Database Explorer. For more information, see [Filter database objects](#).

Search for database objects

Database object searches rely on the Object Index when returning results. By default, caching is set to configure only parts of a database. To configure the Index to expand object searches, see [Specify Data Source Indexing Preferences](#).

To search for database objects

1. Select **Search > Database**. By default, the search scope is all currently connected databases. Under Specify the scope for the search, clear any databases or server check boxes you do not want to search.
2. Specify the search criteria:
 - Type the value to search for in the **Search String** field. Use the * character to indicate wildcard string values and the ? character to indicate wildcard character values.
 - Select **Case Sensitive** to indicate to the search function that you want case sensitivity to be a factor when searching for appropriate string matches.
 - Select **Search Indexed Data** to indicate that the search function should read the Index. This increases the performance of the search function and will typically result in faster returns on any hits the search might make.
 - Select **Apply SQL Filters** to apply any relevant database or vendor filters to the search.
 - Choose **Declarations, References, or All Occurrences** to specify what the search is restricted to in terms of database objects.
 - A **Declaration** is an instance where an object is declared. For example, an object is declared in a CREATE table.
 - A **Reference** is an instance where an object is used or referred to. For example, an object is referred to in a procedure or as a foreign key in a table.
 - Choose **All Occurrences** to return both declarations and references in the search results.
 - Use the check boxes beside the database object panel to select and deselect the specific database objects that you want to be included in the search process.
- Click **Search**.

The results of your search are generated in the **Search** view. When you open a matched file, references to the keyword are flagged with yellow arrow icons that appear in the left-hand column of the editor.

```

DELETE FROM Shape_Edge_Display_Ver
WHERE Shape_Edge_Display_ID IN
(SELECT do.Object_Table_Row_ID
FROM Diagram_Object do,
Meta_Table mt
WHERE mt.Name = 'Shape_Edge_Display' and
do.Meta_Table_ID = mt.Meta_Table_ID and
do.Diagram_ID = @DelDiagramID)

DELETE FROM Shape_Edge_Display
WHERE Object_GUID IN
(SELECT do.Object_GUID
FROM Diagram_Object do,
Meta_Table mt
WHERE mt.Name = 'Shape_Edge_Display' and
do.Meta_Table_ID = mt.Meta_Table_ID and
do.Diagram_ID = @DelDiagramID)
  
```

You can navigate between keywords within all returned files using the yellow "up" and "down" arrows that appear at the top of the **Search** view.

Filter database objects

You can place filters on data sources and corresponding data source objects to restrict their display in Data Source Explorer. This feature is useful if you have data sources that contain large numbers of database objects. You can apply filters to view only the schema objects you need for the development process.

The following types of data source filters are available:

- **Global filters** that affect all registered data sources in the DB Optimizer development environment.
- **Data Source specific filters** affect only the specified data source for which they are defined.
- *On Sybase and SQL Server platforms*, you can apply **database filters**, which enables you to set different filters on different databases within the same source.

In all of the above cases, data source object filters are defined via the Object Filter Manager, through the development of filter templates. Once defined, filter templates can be activated and deactivated as you need them.

Several filter templates can be combined at a global level or applied to a specific data source.

See also:

- [Define Data Source-Specific Object Filters](#)
- [Define Global Database Object Filters](#)

Define Data Source-Specific Object Filters

Data source-specific object filters affect only the specified data source.

To define data source-specific filters

1. In **Data Source Explorer**, right-click the data source and select **Properties**. The **Properties** dialog appears.
2. Select the **SQL Filters** node and select **Enable data source-specific settings**. The other controls on the dialog become enabled.
3. Click **New**. The **Filter Template** dialog appears.
4. Specify the parameters of the filter.
 - In the **Name** field, enter the name of the filter as you want it to appear in the selection window on the SQL Filter node.
 - The **Database Type** pane provides a list of data source objects. Deselect the data source objects that this template filters so that they do not appear in Database Explorer when displaying data source objects for the data source.
 - Click **New** to add filter parameters for data source object properties. The **New SQL Filter Predicate** dialog appears.
 - Use the **Property** and **Operator** fields to supply the filter criteria. **Property** specifies whether the value is a **Name** or **Schema**, and **Operator** specifies the matching type of the filter syntax. (Equals, Not Equals, Like, Not Like, In, Not In)
 - In the **Value** field, enter the full or partial syntax of the property or properties you want to filter in Data Source Explorer.
 - Click **OK**. The filter property specification is added to the Filter Template.
 - When you have finished defining the filter template, click **OK**. The template name is added to the Properties dialog. It can be enabled and disabled by selecting or deselecting the check box beside its name, respectively.


Define Global Database Object Filters

Global filters affect all registered data sources in the DB Optimizer development environment. When you create and apply a global filter to a platform vendor in DB Optimizer, all databases associated with that vendor are affected by the filter, as defined.

Individual global filter templates are separated, by supported data source platform, on tabs in the SQL Filter window. Select the appropriate tab to view existing filter templates or add new ones, as needed.

To define a global filter

1. Select **Window > Preferences** from the Main Menu. The **Preferences** dialog appears.
2. Expand the **SQL Development** node and select the **SQL Filter** subnode. The **SQL Filter** pane appears.
3. Click **New**. The **Filter Template** dialog appears.
4. Specify the parameters of the filter template:
 - In the **Name** field, enter the name of the filter as you want it to appear in the selection window on the SQL Filter node.
 - The **Database Type** pane provides a list of data source objects. Deselect the data source objects that this template filters so that they do not appear in Database Explorer when displaying data source objects for the data source.
 - Click **New** to add filter parameters for data source objects properties. The **New SQL Filter Predicate** dialog appears.
 - Use the **Property** and **Operator** fields to supply the filter criteria. Property specifies whether the value is a Name or Schema, and Operator specifies the matching type of the filter syntax. (Equals, Not Equals, Like, Not Like, In, Not In)
 - In the **Value** field, enter the full or partial syntax of the property or properties you want the template to filter in data source Explorer.
 - Click **OK**. The filter property specification is added to the Filter Template.
 - When you have finished defining the filter template, click **OK**. The template name is added to the Properties dialog. It can be enabled and disabled by selecting or de-selecting the check box beside its name, respectively.

 Data Source object filters are added and removed from the development environment by selecting and de-selecting the checkboxes associated with each filter template on both the global and data source-specific dialogs.


Drop a database object

To delete an object permanently from a database, right-click the object in Data Source Explorer and choose **Drop** from the menu. The **Drop Wizard** prompts you to confirm removal of the object and provides a DDL preview of the deletion code.

Working with SQL projects

You create projects to organize and store SQL development files. The purpose of projects is to keep your work-in-progress files organized, as well as maintain a common directory structure when developing code and executing files on registered data sources. Once a file has been developed and is ready for deployment, that file can then be executed on a registered data source.

SQL Project Explorer is used to view and access files. It uses a tree view to display the project as a series of folder directories with a folder labeled with the project name as the parent directory, and with project categories, and associated project files as its children.



All files in an SQL Project project are organized under the following categories:

- **Connections.** List the connections of any given SQL file of a data source associated with the project.
- **Creation Scripts.** Provide DDL statements and statements that define database objects.
- **General SQL.** Provide a category for all other SQL files that are not used in database object creation. This includes DML files, and so on.
- **Large Scripts.** Contain all files larger than the currently set SQL Editor preference. The file size limit can be modified on the Preferences panel by selecting **Window > Preferences** in the Main Menu.

Physically, the projects and files you create as you work in DB Optimizer are stored under the Workspace directory you specified at the prompt when DB Optimizer was started. The directory and files can be shared, and other tools can be used to work on the files, outside the DB Optimizer development environment.

You can move existing files within a project by clicking and dragging the file you want to move in the Project Explorer from one node to another, or via the **File > Move** command.

- [Create a new SQL project](#)
- [Open an existing project](#)
- [Search a Project](#)
- [Add files to a project](#)
- [Delete a Project](#)

Create a new SQL project

To create a new SQL project

1. Select **File > New > SQL Project** from the DB Optimizer Main Menu. The New Project Wizard appears.
2. Enter the appropriate information in the fields provided:
 - **Name.** Enter the name of the project as you want it to display in the Project Explorer view.
 - **DBMS Platform.** Select the data source platform to which the new project will be associated. This enables DB Optimizer to properly parse SQL development code for project files.
 - **Location.** When selected, the **Use Default Location** check box indicates the project is to be created under the currently selected Workspace. Deselect the check box and specify a new folder path if you do not want to create the project in the currently selected Workspace.
3. Click **Finish**. The new project icon appears in the Project Explorer view under the name that you specified. If you did not select **Use Default Location**, the project will appear in the appropriate Workspace when you open it in DB Optimizer.


i Alternatively, you can select **New > SQL Project** from the Main Menu or click the New Project icon in the Tool Bar to create a new project.

Open an existing project

You can open projects by navigating to SQL Project Explorer and expanding the node of the project that contains the files you want to access.

Below each project name are a series of nodes that categorize any existing SQL files by development type:

- **Connections.** Lists the connections of any given SQL file of a data source associated with the project.
- **Creation Scripts.** General data source object development scripts. This node contains DDL statements and statements that define database objects.
- **General SQL.** Provides a category for all other SQL files that are not used in database object creation. DML files, etc.
- **Large Scripts.** Contains all files larger than the currently set SQL Editor preference. The file size limit can be modified on the Preferences panel. (Choose **Window > Preferences** in the Main Menu to access the panel.)

 Physically, the projects and files you create as you work in DB Optimizer are stored under the project directory that you specified at the prompt when the project was created. The directory and files can be shared, and other tools may be used to work on the files, completely exempt from the DB Optimizer development environment.

Search a Project

To search an existing project

1. Select **Search > File**.
2. Specify the search criteria:
 - Type the value to search in the **Containing Text** field. Use the * character to indicate wildcard string values, the ? character to indicate wildcard character values, and the \ character to indicate an escape character for literals (* ? /).
 - Select **Case Sensitive** and indicate to the search function that it should take into account case when searching for appropriate string matches.
 - Select **Regular Expression** to indicate to the search function that the string is a regular function.
 - In the **File Name Pattern** field, specify the extension name of the files to search for explicitly. If the value in this field is a * character, the search function searches all files regardless of extension. Manually type in the extensions to indicate file type (separate multiple file types with commas), or click **Choose** and use the Select Types dialog to select the file extensions the process will search for the string by.
 - Select **Consider Derived Resources** to include derived resources in the search.
 - Select **Workspace** or **Working Set** to choose the scope of the search. If you choose **Working Set**, specify the name of the defined working set manually, or click **Choose** and navigate to the working set you want to search for in the provided string.
3. Click **Search**. The results of your search are generated in the Search view on the Workbench.


Add files to a project

Existing files that reside in directories outside of the workspace can be added to a project via the following methods:

- Dragging and dropping the file set from a system directory to SQL Project Explorer.
- Copying and pasting the file set from a system directory to SQL Project Explorer.
- Executing the Import command.


To drag/drop or copy/paste files from a system directory to SQL Project Explorer

1. With the SQL Project Explorer view open, navigate to the directory where the files you want to add to the project are located on the system.
2. Drag and drop the files you need from Windows Explorer into SQL Project Explorer. The files appear in the tree view under the appropriate categories.

 Alternatively, you can use the **Copy** command on the files you want to add in Windows Explorer and then right-click the **Project Explorer** and select **Paste** from the menu. The files appear in the tree view under the appropriate categories.

To use the Import command

1. Right-click anywhere on the Project Explorer and select **Import**. The **Import** dialog appears.
2. Expand the **General** node and double-click **File System**. A dialog containing the import specification parameters appears.
 - In the **From directory** field, manually type the directory location of the files you want to import to Project Explorer, or click **Browse** and navigate to the appropriate folder. The panels below the field populate with the folder selection and a list of suitable files contained in that folder. Use the check boxes beside each folder and file to specify what folders/files you want the import function to add in Project Explorer.
 - In the **Into folder** field, manually type the name of the folder within Project Explorer where you want to import the files specified in the panels above, or click **Browse** and navigate to the appropriate folder.
 - Select the **Overwrite existing resources without warning** check box if you do not want to be prompted when the import process overwrites Project Explorer files that contain the same name as the imported files.
 - Choose **Create complete folder structure** or **Create selected folders only**, depending on whether you want the import process to build the folder structure of the imported directory automatically, or only create those folders you selected in the panels above, respectively.
 - Click **Finish**. The import process moves all selected folders and files into Project Explorer and thus into the DB Optimizer development environment.

 In addition to accessing the Import command via the shortcut menu, you can also access the Import dialog by choosing **File > Import** from the Main Menu.

Delete a Project

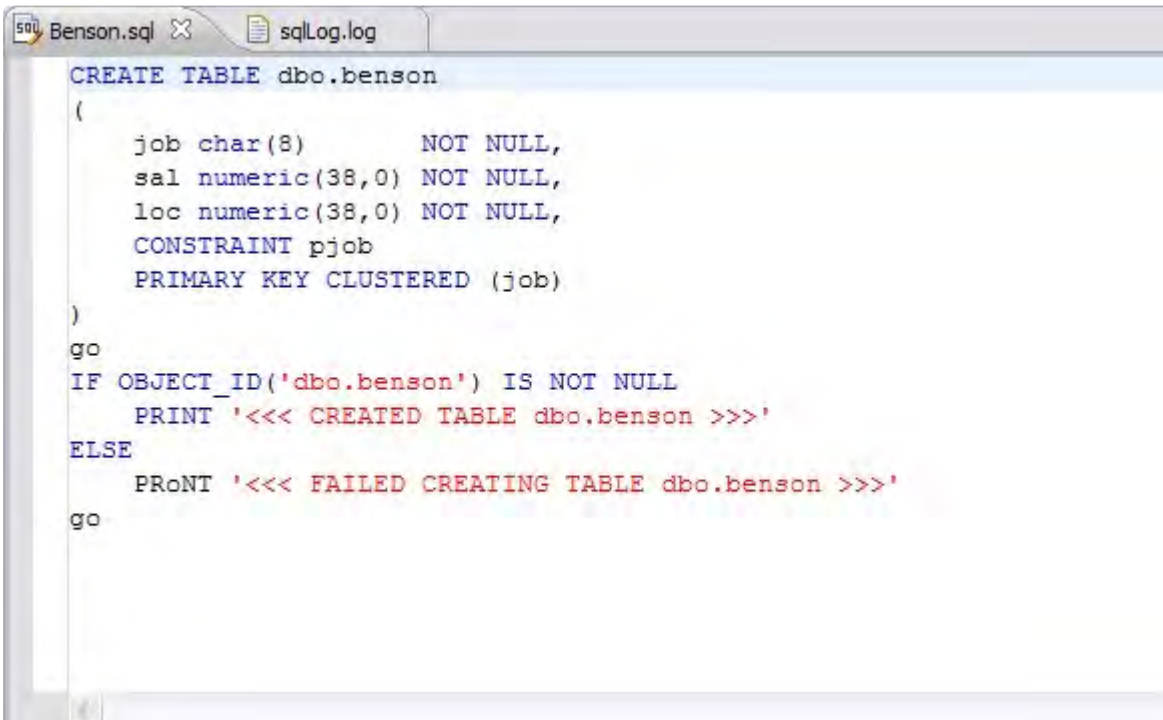
You can delete a project by right-clicking its folder in the SQL Project Explorer and selecting **Delete**.

When you delete a project, DB Optimizer will prompt you with a Confirm Project Delete dialog that asks you to confirm the deletion of the project, and offers you the option of deleting the project from the DB Optimizer interface, or deleting the project from the system.

- If you select **Do not delete contents**, the files and directory structure are removed from SQL Project Explorer, but they still exist on your machine.
- If you select **Also delete contents**, the files and directory structure are removed from SQL Project Explorer and deleted from your machine.

Creating and editing SQL files (SQL Editor)

The SQL Editor is a Workbench interface component that enables the development, viewing, and formatting of SQL code.



```

CREATE TABLE dbo.benson
(
    job char(8)          NOT NULL,
    sal numeric(38,0)   NOT NULL,
    loc numeric(38,0)   NOT NULL,
    CONSTRAINT pjob
    PRIMARY KEY CLUSTERED (job)
)
go
IF OBJECT_ID('dbo.benson') IS NOT NULL
    PRINT '<<< CREATED TABLE dbo.benson >>>'
ELSE
    PRINT '<<< FAILED CREATING TABLE dbo.benson >>>'
go

```

The Editor supports the following the functionality:

- **Code assist:**
 - **Code complete.** Type Ahead and Name completion. For more information, see [Understanding Code Assist](#).
 - **Code templates.** Templates for creation of tables, procedures, etc. For more information, see [Understanding SQL Templates](#).
 - **Hyperlinks.** For more information, see [Understanding Hyperlinks](#).
 - **Semantic validation.** For more information, see [Semantic Validation](#).
 - **Object hovering.** Hover over an error found and an explanation of the cause of the error appears.
- **Code formatter.** For more information, see [Understanding Code Formatting](#).
- **Code correction and transformations.** For more information, see [Examples of transformations and SQL query rewrites](#).
- **Object indexing.** For more information, see [Specify Data Source Indexing Preferences](#).
- **SQL Project Explorer.** For more information, see [Working with SQL projects](#).

SQL Editor contains context-sensitive command menus that are tailored with pertinent functionality for the specified file format.

If SQL Editor does not recognize a selected file format, DB Optimizer automatically launches the file externally in the system default application. External editors are not embedded in the Workbench. For example, on most machines, the default editor for HTML files is the system Web browser. SQL Editor does not, by default, recognize HTML files, and opening an HTML file from the Workbench launches the file in an instance of the Web browser instead of the Editor.

Any number of instances of SQL Editor can be open on the Workbench at the same time. Multiple instances of SQL Editor displaying different content may be open in the same Workbench. These instances will be stacked by default, but can also be tiled side-by-side so the content of various files can be viewed simultaneously for comparison or multi-tasking purposes. When an instance of SQL Editor is active, the Workbench Main Menu automatically contains commands applicable to the file format. If a view is active, SQL Editor commands are disabled automatically, except when commands are still valid between the selected view and the file displayed in the interface.

Semantic Validation

When working with code in SQL Editor, the window contains a number of features that provide an increase in the efficiency and accuracy of code development. The following syntax highlighting changes are automatically applied to code as a user adds lines in the interface.

| Code | Formatting |
|-----------------------------------|-------------------------|
| Comments | Green font, italics SQL |
| Commands | Dark blue font Coding |
| Errors | Red underline |
| Strings | Red font |
| Non-Executable Command Line Comma | Aqua font |

Single-line and multiple-line comments appear in different colors.

Furthermore, SQL Editor provides two column bars, one on either side of the code window. The purple change bar in the left-hand column indicates that the line of code has been modified. Hover over the change bar to display the original code text. The red square in the right-hand column indicates that there are errors in the code window. Hover the mouse over the square to view the error count. Click the red bar in this column to navigate directly to the line in which the SQL Editor detects the error. SQL Editor automatically highlights the appropriate code. Non-executable command line commands are displayed in a different formatting style than SQL commands. Syntactic and semantic errors are also highlighted.

SQL Editor also features dynamic error detection, object lookup and suggestion features, code folding, and auto-formatting. SQL Editor is able to identify different areas in a statement, and enables users to retrieve subclauses, resolve table aliases, and dynamically return lists of tables, views, and columns, as needed.

See also [Working in SQL Editor](#).

Create an SQL file

To create an SQL file

1. Create or open an SQL project.
2. Select **File > New > SQL File**. A blank instance of **SQL Editor** appears.



If you are not in a SQL project when you create a new SQL file, it does not open in SQL Editor.

Open an existing SQL file

To open an existing SQL file

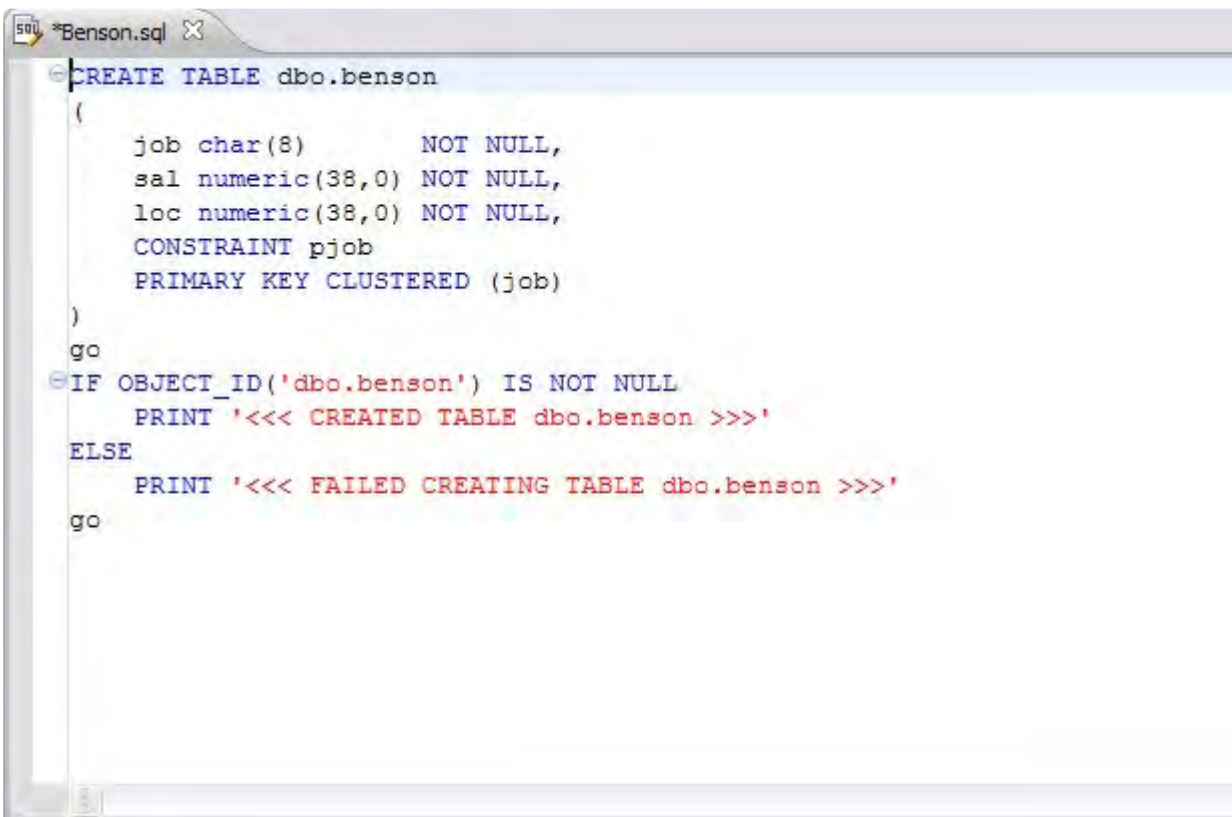
1. Open the SQL project containing the file, or that you want to contain the file.
2. If necessary, add the file to the project. See [Add files to a project](#).
3. In the SQL Project Explorer, double-click the file to open it in SQL Editor.

Working in SQL Editor

SQL Editor handles SQL code formats and contains context-sensitive command menus, tailored with pertinent functionality for development purposes. Other files may be opened in DB Optimizer, as well, but these are handled by other editors.

For example, if a text file is opened in the Workbench, DB Optimizer detects and opens the contents of that file in a text editor viewer with pertinent commands for that file type.

Any number of instances of SQL Editor can be active on the Workbench at the same time. Multiple instances of SQL Editor displaying different content may be active on the same Workbench. These instances will be stacked, by default, but can also be tiled side-by-side, so the content of various files can be view simultaneously for comparison or multi-tasking purposes. When an instance of SQL Editor is active, the Main Menu contains commands applicable to the file format. If a view is active, SQL Editor commands are disabled automatically, except when commands are still valid between the selected view and the file displayed in the interface.




```

CREATE TABLE dbo.benson
(
    job char(8)          NOT NULL,
    sal numeric(38,0) NOT NULL,
    loc numeric(38,0) NOT NULL,
    CONSTRAINT pjob
    PRIMARY KEY CLUSTERED (job)
)
go
IF OBJECT_ID('dbo.benson') IS NOT NULL
    PRINT '<<< CREATED TABLE dbo.benson >>>'
ELSE
    PRINT '<<< FAILED CREATING TABLE dbo.benson >>>'
go
  
```

Among the commands SQL Editor supports via the right-click menu:

- **Revert File.** Automatically restores the working file to the original text as it appeared the last time the Save command was issued.
- **Shift Right/Shift Left.** Indents the line of code in the working file to the right or left, respectively.
- **Toggle Comments.** Hides or displays comments in the code of the working file, depending on the current hide/show state.
- **Add Block Comment/Remove Block Comment.** A block comment is used to insert a comment into SQL code that spans multiple lines and begins with a forward slash and asterisk. While block comments are typically used to insert a command that spans multiple lines, some developers find them more useful than line comments, especially if a development team is using different text editors on an individual basis. Moving code from one text editor to another often breaks line comments in the middle of a line and causes errors. Block comments can be broken without causing errors.

 In addition to editing commands, some commands such as extract, drop, and execute can be accessed by right-clicking over statements in SQL code that are performed on specific tables, views, and columns. These commands will appear automatically in the appropriate menu when the code is highlighted. Full information on using these commands is found elsewhere in this documentation, based on the task each executable performs.

- **Explain Plan.** An explain plan details the steps that occur in SELECT, UPDATE, INSERT, and DELETE statements and is primarily used to determine the execution path followed by the database in its SQL execution.

See also:

- [Understanding Automatic Error Detection](#)
- [Understanding Code Assist](#)
- [Understanding Hyperlinks](#)
- [Understanding Code Formatting](#)
- [Understanding Code Folding](#)
- [Understanding Code Quality Checks](#)
- [Understanding SQL Templates](#)

Understanding Automatic Error Detection

SQL Editor orders and classifies SQL statements. This enables it to edit code as you work within SQL Editor and highlight errors and typographical errors in "real time." As you work, SQL Editor examines each clause in a statement and provides error reporting and other features as required.

SQL Editor identifies the following clauses and elements:

- **SELECT.** Specifies the field, constants, and expressions to display in the query results.
- **FROM.** Specifies one or more tables containing the data that the query retrieves from.
- **WHERE.** Specifies join and filter conditions that determine the rows that query returns. Join operations in a WHERE clause function in the same manner as JOIN operations in a FROM clause.
- **GROUP BY.** Specifies one or more columns used to group rows returned by the query.

Columns referenced in the SQL SELECT statement list, except for aggregate expressions, must be included in the GROUP BY clause. You cannot group by Memo, General or Blob fields.

- **HAVING.** Specifies conditions that determine the groups included in the query. If the SQL statement does not contain aggregate functions, you can use the SQL SELECT statement containing a HAVING clause without the GROUP BY clause.
- **ORDER BY.** Specifies one or more items used to sort the final query result set and the order for sorting the results.

As you develop code in SQL Editor, it automatically detects semantic errors on a line-by-line basis. Whenever an error is detected, the line is flagged by an icon located in the left-hand column of the editor.

```

CREATE TABLE dbo.benson
(
    job CHAR (8) NOT NULL,
    sal NUMERIC (38, 0) NOT NULL,
    loc NUMERIC (38, 0) NOT NULL,
    CONSTRAINT pjob PRIMARY KEY CLUSTERED (job)
)

go
IF OBJECT_ID('dbo.benson') IS NOT NULL
    PRINT '<<< CREATED TABLE dbo.benson >>>'
ELSE
    PRoNT '<<< FAILED CREATING TABLE dbo.benson >>>'

go
SELECT *
FROM tbo.benson;

```

Additionally, all semantic errors detected in SQL Editor are displayed in the Problems view.

| Description | Resource | Path | Location |
|---|---------------|------|----------|
| An unexpected token ""<<< FAILED CR Benson.sql | SQL Project 1 | | line 14 |
| An unexpected token "" was found. Exr File.sql | SQL Project 1 | | line 6 |
| Table benson cannot be resolved on 'da Benson.sql | SQL Project 1 | | line 19 |

Right-click the error and select **Go To** in order to find the error. DB Optimizer opens and navigates to the specific line of code containing the specified error.

i Automatic error detection functions, such as syntax checking and semantic validation are suspended if #define or #include directives are detected in an editor window. DB Optimizer does not perform #define/#include substitutions on execution.

Understanding Code Assist

When SQL Editor has finished analyzing a partial piece of code, it displays a list of data source objects for you to select from.

SQL Editor takes the following into consideration when analyzing code for a list of possible data source objects for insertion:

- Text to be inserted
- Original text to be replaced
- Content assist request location in original text
- The database object represented by the insertion text

Generally, insertion suggestions use the following format:

<insert

ion_text

t > -

<qualif

ication

_inform

ation >

Code assist is available for SELECT, UPDATE, INSERT, and DELETE statements (and MERGE statements on SQL Server 2008 and above), as well as stored procedures, and functions (built-in and user defined.)

Additionally, code suggestions can be made for DML statements nested within DDL statements. This functions in the same manner as code assist for statements that are not nested, and applies to CREATE PROCEDURE, FUNCTION, TRIGGER, TABLE, and VIEW statements.

When the code assist window is open, you can filter out singular object suggestions by pressing (Ctrl + Spacebar). This removes all objects from the assist window while retaining procedures and functions. To display objects again, press (Ctrl + Spacebar) again.

The following table displays a list of all possible object suggestions, and the format in which SQL Editor inserts the suggestions into a statement:

Object and Stored Procedure Suggestions

| Object Suggestion | Syntax/Example |
|-------------------|--|
| Table | (TABLE) [catalog].[schema] EMPLOYEE - (TABLE)HR |

| Object Suggestion | Syntax/Example |
|-------------------|--|
| Alias Table | (TABLE ALIAS) [catalog].[schema]tableName EMPLOYEE - (TABLE ALIAS)HRJOBS |
| Column | datatype - (Column) [catalog].[schema].tableName JOB_TITLE:varchar(20) - (Column)HRJOBS |
| Alias Column | datatype - (COLUMN ALIAS) [catalog].[schema].tableName.columnName JOB_TITLE:int-(COLUMN ALIAS)HR.JOBS.JOB_ID |
| Schema | (SCHEMA) [catalog] dbo-(SCHEMA)NorthWind |
| Catalog | (CATALOG) |
| Call | Call HR.ADD_JOB_HISTORY |

Function Suggestions

| Function Suggestion | Syntax/Example |
|---------------------|---|
| Built-in | SELECT A FROM HR.DEPARTMENTS WHERE HR.DEPARTMENTS AVG |
| User-Defined | SELECT + FROM HR.CLIENTS WHERE HR.F_PERSONAL |

 Function suggestions are only available for Oracle and DB2 platforms.

SQL Editor detects incomplete or erroneous code, processes the code fragments, and then attempts to apply the appropriate logic to populate the code.

As code is typed into SQL Editor, the application ‘reads’ the language and returns suggestions based on full or partial syntax input.

Depending on the exact nature of the code, the automatic object suggestion feature behaves differently; this enables SQL Editor to provide reasonable and ‘intelligent’ suggestions on coding.

Additionally, semantic validations can be made for DML statements nestled within DDL statements. This functions in the same manner as validation for top-level statements, and applies to CREATE PROCEDURE, FUNCTION, TRIGGER, TABLE, and VIEW statements.

The following chart displays the possible statement fragments that SQL Editor will attempt to suggest/populate with objects:

| Statement Fragment Elements | Object Suggestion Behavior |
|-----------------------------|---|
| SELECT | A list of tables, when selected automatically, prompts the user to select a column. |
| UPDATE and DELETE | A list of tables appears in the FROM and/or WHERE clause. |
| INSERT | A list of tables and views appears in the INSERT INTO and OPEN BRACKET clause prior to values. A list of columns based on the table or view name appears in the OPEN BRACKET or VALUES clause. |

In addition to DML statements, SQL Editor also suggests objects based on specific fragmented syntax per line of code:

| Statement Syntax | Object Suggestion Behavior |
|---|--|
| A partial DML statement (for example SEL ... indicates a fragment of the SELECT clause) | The keyword is completed automatically, assuming SQL Editor can match it. Otherwise, a list of suggested keywords is displayed. If the preceding character is a period, and the word prior is a table or view, a list of columns appears. If the word being typed is a part of a table name (denoted by a schema in front of it) the table name is autocompleted. If the word being typed has a part of a column name (denoted by a table in front of it) the column name is autocompleted. |
| Without typing anything. | A list of keywords appears. |
| A period is typed. | If the word prior to the period is a name of a table or view, a list of columns is displayed. If the word prior to the period is a schema name, a list of table names is displayed. If the word prior to the period is either a table name or a schema name, then both a list of columns and a list of table names is displayed |

To activate code suggestions

By default, code suggestions are automatically offered if you stop typing in SQL Editor for one second. You can turn off the automated suggestion feature on the Code Assist preferences page.

If automated code suggestion is disabled, you can still access the suggestion window using the following method:

1. Click the line that you want SQL Editor to suggest an object for.

2. Press (**CTRL + Spacebar**) on your keyboard. SQL Editor 'reads' the line and presents a list of tables, views or columns as appropriate based on statement elements.

i On a per platform basis, auto-suggestion behavior may vary. (For example, the WITH statement on DB2 platforms.)

To modify object suggestion parameters, including setting it from automatic to manual, see [Specify Code Assist Preferences](#).

You can speed up the performance of the code assist functionality by enabling data source indexing either when you connect to the data source, see [Working with data sources](#) or on the Preferences page, see [Specify Data Source Indexing Preferences](#).

Understanding Hyperlinks

SQL Editor supports hyperlinks that are activated when a user hovers their mouse over a word and presses the CTRL key. If a hyperlink can be created, it becomes underlined and changes color. When the hyperlink is selected, the creation script for the hyperlink object is opened in a new editor.

Hyperlinks can be used to link to tables, columns, packages, and other reference objects in development code. Additionally, hovering over a hyperlink on a procedure or function of a call statement will open it. You can also use the hyperlink feature on function calls in DML statements.

Clicking a hyperlink performs an action. The text editor provides a default hyperlink capability. It allows a user to click on a URL (for example, <http://www.idera.com>) and database object links.

Hyperlink options (look and feel) can be modified via the Hyperlinking subnode in the **Editors > Text Editors** node of the Preferences panel.

i Hyperlink functionality relies on certain objects being captured in the Object Index. If the index is turned off, or has been restricted in what information it captures, users will be unable to link them (as they are non-existent within the Index.) To specify object index types, see [Specify Data Source Indexing Preferences](#).

Understanding Code Formatting

Code formatting provides automatic code formatting in SQL Editor while you are developing code.

To access the code formatter, select the open editor you want to format and select Ctrl+Shift+F. The code is formatted automatically based on formatting parameters specified in the Code Formatter subnode of the SQL Editor node in the Preferences panel.

You can also format an entire group of files from Project Explorer. To do so, select the directory or file and execute the Format command via the shortcut menu. The files will be formatted automatically based on your formatting preferences. See [Specify Code Formatter Preferences](#) for more information.

The following examples display a list of code formatting parameters and the resultant output in SQL Editor, based on the same set of SQL statements.

Custom Code Formatting Example 1

The following chart indicates a list of custom code formatting parameters and their corresponding values. The chart is followed by the actual syntax as it would appear in SQL Editor, based on the formatting parameter values.

Compare the parameters and formatted code in Example 2 with this example for a concept of how custom formatting works.

| Custom Code Formatting Parameter | Value (if applicable) |
|---|------------------------------|
| Stack commas separated by lists? | Yes |
| Stack Lists with ___ or more items. | 3 |
| Indent Size? | 2 |
| Preceding commas? | Yes |
| Spaces after comma? | 1 |
| Trailing commas? | - |
| Spaces before comma? | - |
| Right align FROM and WHERE clauses with SELECT statement? | Yes |
| Align initial values for FROM and WHERE clauses with SELECT list? | Yes |
| Place SQL keywords on their own line? | No |
| Indent size? | - |
| Indent batch blocks? | Yes |
| Number of new lines to insert | 1 |
| Indent Size | 5 |
| Right Margin? | 80 |
| Stacked parentheses when they contain multiple items? | No |
| Stacked parentheses when the list contains ___ or more items. | - |
| Indent size? | 5 |
| New line after first parentheses? | No |

| Custom Code Formatting Parameter | Value (if applicable) |
|---|-----------------------|
| Indent content of conditional and looping constructs? | Yes |
| Number of new lines to insert? | 1 |
| Indent size? | 5 |

```

Begin
  If x=5
    SELECT apple
           \ pear
           \ orange 'Big Orange'
           \ strawberry
           \ orchard_name
           \ owner
    FROM fruit F, orchard O
  WHERE fruit_region in ('latin america'
           \ 'france'
           \ 'russia'
           \ 'canada'
           \ 'hawaii')
  and orchard not in (select region
                     from bad_growers bg, (select orchard
                                           from hybrid_growers
                                           where us_approved in

```

Custom Code Formatting Example 2

The following chart indicates a list of custom code formatting parameters and corresponding values. The chart is followed by the actual syntax as it would appear in SQL Editor based on the formatting parameter values. Compare the parameters and formatted code in Example 1 with this example for a concept of how custom formatting works.

| Custom Code Formatting Parameter | Value (if applicable) |
|-------------------------------------|-----------------------|
| Stack commas separated by lists? | Yes |
| Stack Lists with ___ or more items. | 2 |
| Indent Size? | 0 |
| Preceding commas? | -- |

| Custom Code Formatting Parameter | Value (if applicable) |
|---|------------------------------|
| Spaces after comma? | Yes |
| Trailing commas? | Yes |
| Spaces before comma? | 2 |
| Right align FROM and WHERE clauses with SELECT statement? | No |
| Align initial values for FROM and WHERE clauses with SELECT list? | -- |
| Place SQL keywords on their own line? | Yes |
| Indent size? | 4 |
| Indent batch blocks? | No |
| Number of new lines to insert | 1 |
| Indent Size | 5 |
| Right Margin? | 80 |
| Stacked parentheses when they contain multiple items? | Yes |
| Stacked parentheses when the list contains ___ or more items. | 2 |
| Indent size? | 2 |
| New line after first parentheses? | Yes |
| Indent content of conditional and looping constructs? | -- |
| Number of new lines to insert? | 1 |
| Indent size? | 5 |

B' 1=i le

.sql

Begin

If $x = S$

SELECT

apple

,

pear ,

orange

Big

Orange

! ,

strawb

erry ,

orchar

d

name ,

owner

FROM

fruit

F,

orchard

d 0

WHERE

fruit_

region

in

'lati

n

americ

a' ,

' franc

e' ,

' russi

a',



'canad



a',

Understanding Code Folding

SQL Editor features code folding that automatically sorts code into an outline-like structure within the editor window for easy navigation and clarity while developing code.

```

CREATE DEFAULT ET_FALSE AS 0
go
+ IF OBJECT_ID('PERFCNTR_KS.ET_FALSE') IS NOT NULL
go
- CREATE TABLE PERFCNTR_KS.ADDRESS_ROLE
(
    ADDRESS_ROLE_ID    numeric(10,0) IDENTITY(20000,1),
    NAME                varchar(128) NOT NULL,
    DESCRIPTION         varchar(255) NULL,
    LIST_ORDER          int NOT NULL,
    IS_DEFAULT          bit NOT NULL,
    IS_SYSTEM_REQUIRED bit NOT NULL,
    ROWTIMESTAMP        datetime CONSTRAINT DF__ADDRESS_R__ROWTI__1BF
    CONSTRAINT ADDRESSROLEPK
    PRIMARY KEY CLUSTERED (ADDRESS_ROLE_ID)
)
go
EXEC sp_bindefault 'PERFCNTR_KS.ET_FALSE', 'ADDRESS_ROLE.IS_DEFAULT'
go
EX sp_bindrule 'PERFCNTR_KS.TRUEORFALSE', 'ADDRESS_ROLE.IS_DEFAULT'

```

The editor window automatically inserts collapsible nodes in the appropriate lines of code for organizational purposes. This enables you to expand and collapse statements, as needed, while developing code in particularly large or complicated files.

Understanding Code Quality Checks


Code quality markers provide annotations that prevent and fix common mistakes in the code.

These notes appear in a window on any line of code where the editor detects an error, and are activated by clicking the light bulb icon in the margin or by pressing Ctrl + I.

For example, if a statement reads `select * from SCOTT.EMP, SCOTT.DEPT`, when you click the light bulb icon or press Ctrl + I, a window appears beneath the line of code that suggests Add join criteria.

When you click on a proposed fix, the statement is automatically updated to reflect your change.

The following common errors are detected by the code quality check function in the editor:

| Code Quality Check Type | Definition |
|--|---|
| Statement is missing valid JOIN criteria | <p>If a SELECT statement contains missing join criteria, when it is executed, it can produce a Cartesian product between the rows in the referenced tables. This can be problematic because the statement will return a large number of rows without returning the proper results.</p> <p>The code quality check detects missing join criteria between tables in a statement and suggests join conditions based on existing foreign keys, indexes, and column name/type compatibility.</p> <p>Example:</p> <p>The following statement is missing a valid JOIN criteria:</p> <pre>SELECT*FROM employeee,customerc,sales_orders WHERE e.employee_id = c.salesperson_id</pre> <p>The code quality check fixes the above statement by adding an AND clause:</p> <pre>SELECT*FROM employeee,customerc,sales_orders WHERE e.employee_id = c.salesperson_id AND s.customer_id = c.customer_id</pre> <div data-bbox="395 954 1423 1032" style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  This code quality check is valid for Oracle, DB2, and Sybase-specific join conditions. </div> |
| Invalid or missing outer join operator | <p>When an invalid outer join operator exists in a SELECT statement, (or the outer join operator is missing altogether), the statement can return incorrect results.</p> <p>The code quality check detects invalid or missing join operators in the code and suggests fixes with regards to using the proper join operators.</p> <p>Example:</p> <p>The following statement is missing an outer join operator:</p> <pre>SELECT * FROM employee e, customer c WHERE e.employee_id=c.salesperson_id(+)ANDc.state='CA'</pre> <p>The code quality check fixes the above statement by providing the missing outer join operator to the statement:</p> <pre>SELECT * FROM employee e,customer c WHERE e.employee_id = c.salesperson_id(+) AND c.state(+) = 'CA'</pre> |

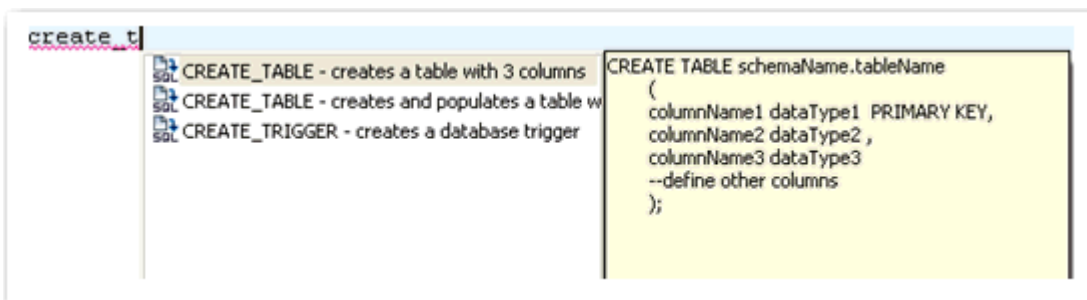
| Code Quality Check Type | Definition |
|------------------------------|---|
| Transitivity issues | <p>The performance of statements can sometimes be improved by adding join criteria, even if a join is fully defined. If this alternate join criterion is missing in a statement, it can restrict the selection of an index in Oracle's optimizer and cause performance problems.</p> <p>The code quality check detects possible join conditions by analyzing the existing conditions in a statement and calculating the missing, alternative join criteria.</p> <p>Example:</p> <p>The following statement contains a transitivity issue with an index problem:</p> <pre>SELECT * FROM item i, product p, price pr WHERE i.product_id = p.product_id AND p.product_id = pr.product_id</pre> <p>The code quality check fixes the above statement with a transitivity issue by adding the missing join condition:</p> <pre>SELECT * FROM item i, product p, price pr WHERE i.product_id = p.product_id AND p.product_id = pr.product_id AND i.product_id = pr.product_id\</pre> |
| Nested query in WHERE clause | <p>It is considered bad format to place sub-queries in the WHERE clause of a statement, and such clauses can typically be corrected by moving the sub- query to the FROM clause instead, which preserves the meaning of the statement while providing more efficient code.</p> <p>The code quality check fixes the placement of sub-queries in a statement, which can affect performance. It detects the possibility of moving sub- queries from the FROM clause of the statement.</p> <p>Example:</p> <p>The following statement contains a sub-query that contains an incorrect placement of a WHERE statement:</p> <pre>SELECT*FROM employee WHERE employee_id=(SELECT MAX(salary) FROM employee)</pre> <p>The code quality check fixes the above statement by correcting the sub- query issue:</p> <pre>SELECT employee.* FROM employee (SELECT DISTINCT MAX(salary) col1 FROM employee) t1 WHERE employee_id = t1.col1</pre> |

| Code Quality Check Type | Definition |
|---|--|
| <p>Wrong place for conditions in a HAVING clause</p> | <p>When utilizing the HAVING clause in a statement:</p> <p>It is recommended to include as few conditions as possible while utilizing the HAVING clause in a statement. DB Optimizer detects all conditions in a given HAVING statement and suggests equivalent expressions that can benefit from existing indexes.</p> <p>Example:</p> <p>The following statement contains a HAVING clause that is in the wrong place:</p> <pre>SELECT col_a, SUM(col_b) FROM table_a GROUP BY col_a HAVING col_a > 100</pre> <p>The code check fixes the above statement by replacing the HAVING clause with equivalent expressions:</p> <pre>SELECT col_a, SUM(col_b) FROM table_a WHERE col_a > 100 GROUP BY col_a</pre> |
| <p>Index suppressed by a function or an arithmetic operator</p> | <p>In a SELECT statement, if an arithmetic operator is used on an indexed column in the WHERE clause, the operator can suppress the index and result in a FULL TABLE SCAN that can hinder performance.</p> <p>The code quality check detects these conditions and suggests equivalent expressions that benefit from existing indexes.</p> <p>Example:</p> <p>The following statement includes an indexed column as part of an arithmetic operator:</p> <pre>SELECT * FROM employee WHERE 1 = employee_id - 5</pre> <p>The code quality check fixes the above statement by reconstructing the WHERE clause:</p> <pre>SELECT * FROM employee WHERE 6 = employee_id</pre> |
| <p>Mismatched or incompatible column types</p> | <p>When the data types of join or parameter declaration columns are mismatched, the optimizer is limited in its ability to consider all indexes. This can cause a query to be less efficient as the system might select the wrong index or perform a table scan, which affects performance.</p> <p>The code quality check flags mismatched or incompatible column types and warns that it is not valid code.</p> <p>Example:</p> <p>Consider the following statement if Table A contains the column col int and Table B contains the column col 2 varchar(3):</p> <pre>SELECT * FROM a, b WHERE a.col = b.col;</pre> <p>In the above scenario, the code quality check flags the 'a.col = b.col' part of the statement and warns that it is not valid code.</p> |

| Code Quality Check Type | Definition |
|-------------------------|---|
| Null column comparison | <p>When comparing a column with NULL, the !=NULL condition may return a result that is different from the intended command, because col=NULL will always return a result of false. Instead, the NULL/IS NOT NULL operators should be used in its place.</p> <p>The code quality check flags occurrences of the !=NULL condition and replaces them with the IS NULL operator.</p> <p>Example:</p> <p>The following statement includes an incorrect col = NULL expression:</p> <pre>SELECT * FROM employee WHERE manager_id = NULL</pre> <p>The code quality check replaces the incorrect expression with an IS NULL clause:</p> <pre>SELECT * FROM employee WHERE manager_id IS NULL</pre> |

Understanding SQL Templates

DB Optimizer provides code templates for DML and DDL statements that can be applied to the Editor via the (Ctrl + Spacebar) command. When you choose a template from the menu that appears, SQL Editor automatically inserts a block of code with placeholder symbols that you can modify to customize the code for your own purposes.

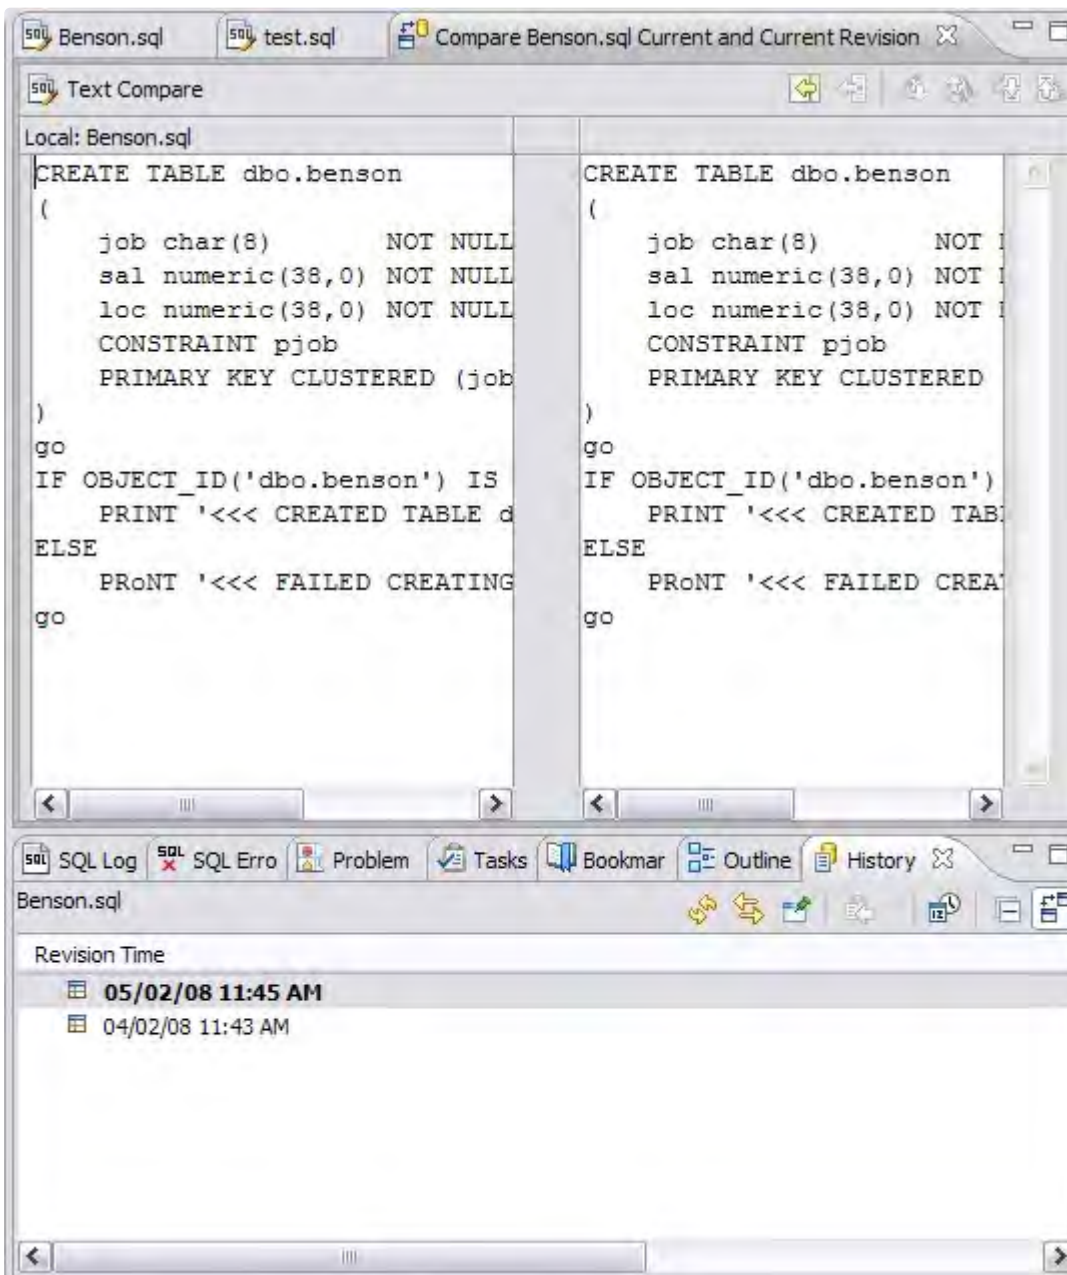


Code templates are available for DML, ALTER, DROP, CREATE, and platform specific commands.

A comprehensive set of DDL/DML templates are available, with statement alternatives varying by DBMS and specific DBMS versions. You can modify and create new templates via the SQL Templates panel on the Preferences dialog. See for more information on how to create and alter SQL code templates.

View change history

Each time an SQL file is saved, the local history of that file is recorded (changes made). Using the Local History command, you can view all changes made to the file. Local History is accessed via the shortcut menu of SQL Editor and selecting **Compare With > Local History**.



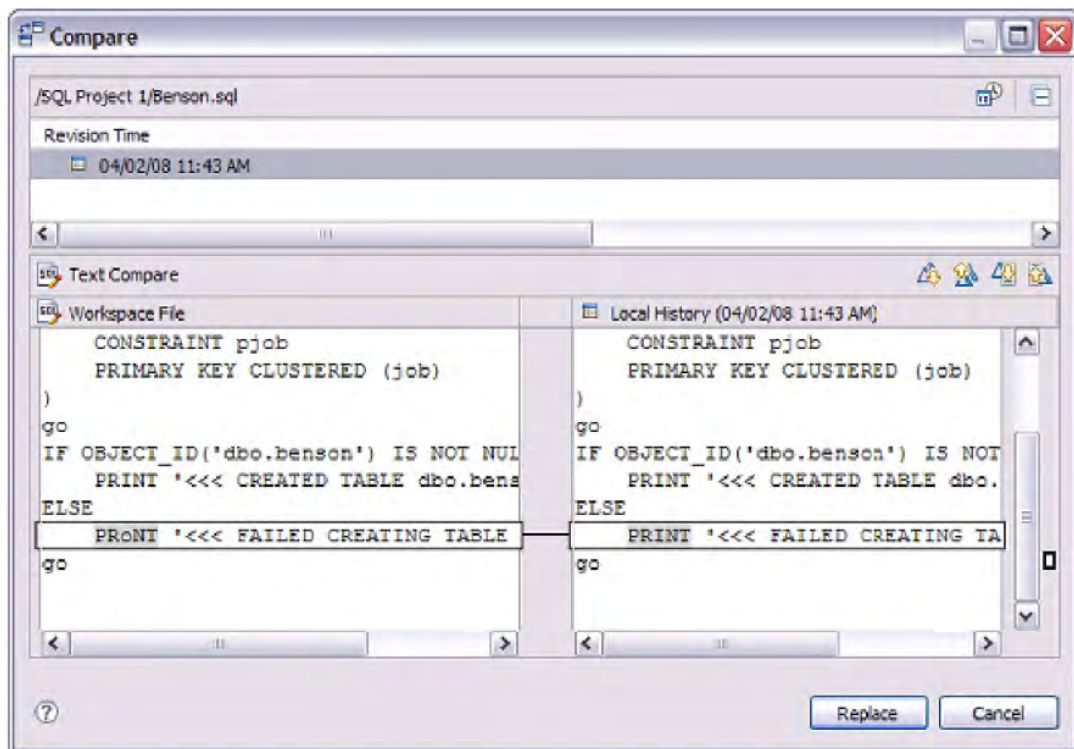
- The History view displays all recorded times the file was changed since its inception/introduction into the workspace.
- Double-click a time in the **History** view to access the **Text Compare** panel. It displays the text of the file after the change occurred at the time indicated in the **History** view.

Revert to an old version of a file

The **Replace With > Local History** command provides you with the ability to revert a SQL file back to a previously recorded local history.

To replace the contents of a file with the contents of a previously saved version via local history

1. Right-click the SQL Editor and select **Replace With > Local History** from the shortcut menu. The **Replace from Local History** dialog appears.



2. In the **Local History of ...** panel, select a previously recorded version of the file by clicking the appropriate timestamp.
3. Click **Replace**.

The contents of the currently-opened file revert to the contents of the file at the history point you selected in the dialog.

Alternatively, from the shortcut menu, select **Replace With > Previous from Local History** to replace the contents of the file with DB Optimizer's last recorded history point.

Delete an SQL file

To delete an SQL file

1. Right-click its icon in the SQL Project Explorer.
2. Select **Delete**. This action removes the file from both the SQL project and the file system.

Executing SQL files

DB Optimizer can execute SQL code directly on registered data sources. Files are executed via the Execute SQL command in the Run menu, or by clicking the green arrow button on the toolbar.

When an SQL file is open in the Workspace, select it and choose a database and an associated catalog on which you want to execute the file via the lists in the Toolbar. You can click the execute icon to execute code on the specified database and catalog, start a transaction or commit a transaction, or modify SQL session options prior to execution.

To execute a file

Open the SQL file you want to run, ensure it is associated with the correct database, and click **Execute**. DB Optimizer executes the code on the data source you specified. Results are displayed in the Results view and can be exported into a file via the Data Export wizard, or displayed in multiple file formats (HTML, XML, and TXT formats).

To execute a transaction

1. Ensure that the auto commit feature is turned off. See [Specify SQL Execution preferences](#) for more information on how to turn off auto commit.
2. Open the transaction file you want to run.
3. Ensure that file is associated with the correct database, and then click **Start Transaction**. DB Optimizer executes the transaction on the data source you specified. Once the transaction runs, you can execute the file as normal.

i Click **Commit** or **Rollback** to finish or cancel a transaction.

To commit a transaction

1. Open the transaction file you want to commit.
2. Ensure that file is associated with the correct database, and then click **Commit Transaction**. DB Optimizer commits the transaction on the data source you specified.

i You can set transactions to auto-commit prior to execution on the SQL Execution node of the **Preferences** panel.


See Also:

- [Associate an SQL file with a data source](#)
- [Configure an SQL session](#)
- [Execute SQL code](#)
- [View and save results](#)

Associate an SQL file with a data source

When working with files, SQL Editor enables developers to view and change the data source to which they are connected.

The bread crumb line in SQL Editor is used to display and specify a data source in relation to the specified SQL Editor file. The menu contains a list of all registered data sources. Additionally, on platforms that support catalogs, these are displayed as well.



Changing a catalog via the drop down lists is the equivalent of issuing a USE DATABASE command on SQL Server, Sybase, and MySQL platforms. Any change will not affect the current connection, and the list automatically updates to display the name of the newly selected data source.

If no registered database is associated with a SQL file (as would be the case if a user started a new, unsaved file), the list is empty. This indicates that the file is not connected to a registered data source.

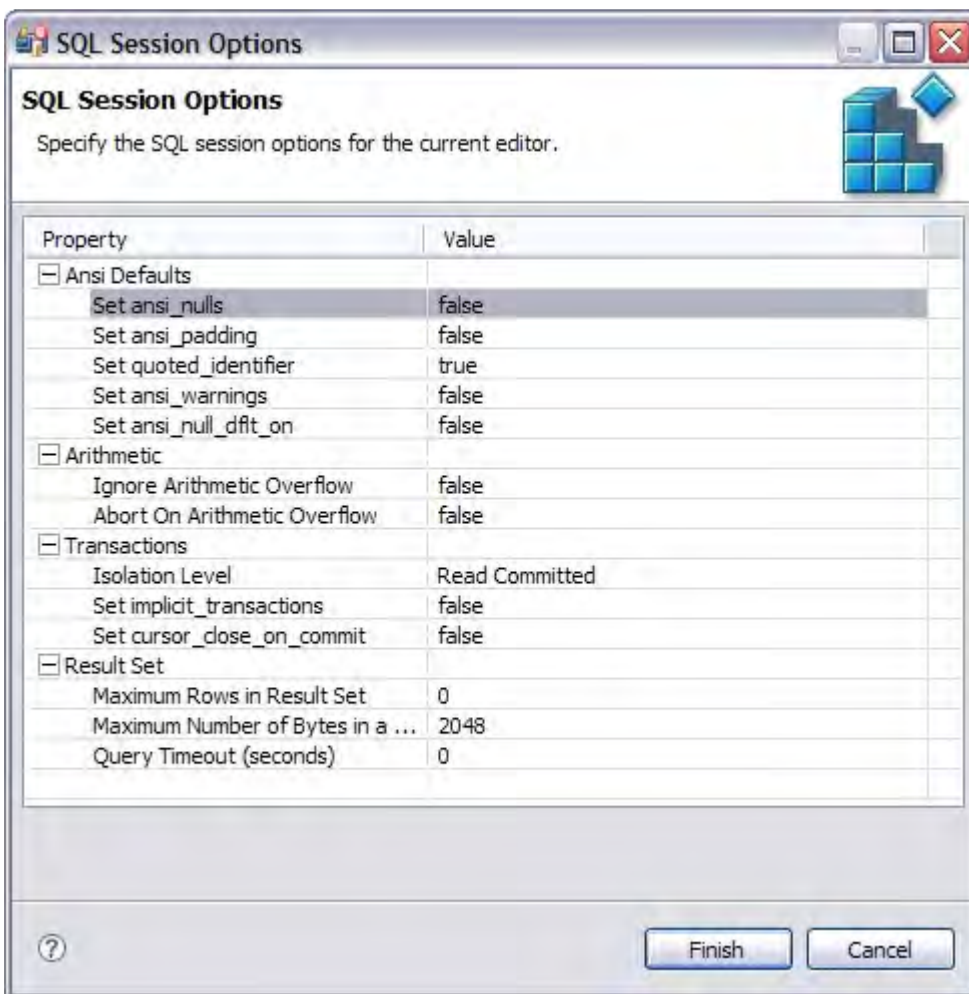
To change or associate a registered data source with a SQL file

1. Click the database list.
2. Select the name of a registered database from the list provided. Depending on the state of the code in SQL Editor, DB Optimizer's behavior differs when the connection is made.

i If you are receiving multiple syntax errors, always check that the file is associated with the correct data source and corresponding database/catalog before troubleshooting further.

Configure an SQL session

The SQL Session Options dialog provides configuration parameters that indicate to DB Optimizer how to execute code in the development environment.



To modify SQL session options

1. Click the SQL Session Options icon in the Toolbar. The **SQL Session Options** dialog appears.
2. Click on individual parameters in the Value column to change the configuration of each property, as specified.
3. Click **Finish**. The session options will be changed and DB Optimizer will execute the code as specified when you execute it.

Session options only apply to the corresponding editor and are not retained when executing multiple SQL files.

Execute SQL code

Files can be launched from within the DB Optimizer development environment for execution on a registered data source. Files are executed via the commands in the Run menu.

When a SQL file is open in the Workspace, select it and choose a database and an associated catalog on which you want to execute file using the drop down menus in the Toolbar. You can click the execute icon to execute the code on the specified database and catalog, start a transaction or commit a transaction, or modify the SQL session options prior to execution.

To execute code

1. Open the SQL file you want to run.
2. Ensure it is associated with the correct database.
3. Click the Execute icon. DB Optimizer executes the code on the data source you specified. Results are displayed in the same tab or in a new tab.

To execute a transaction

1. Open the transaction file you want to run.
2. Ensure it is associated with the correct database.
3. Click the Start Transaction icon. DB Optimizer executes the transaction on the data source you specified.

To commit a transaction

1. Open the transaction file you want to commit.
2. Ensure it is associated with the correct database.
3. Click the Commit Transaction icon. DB Optimizer commits the transaction on the data source you specified.



You can set transactions to auto-commit prior to execution on the SQL Execution node of the Preferences panel in DB Optimizer.

View and save results

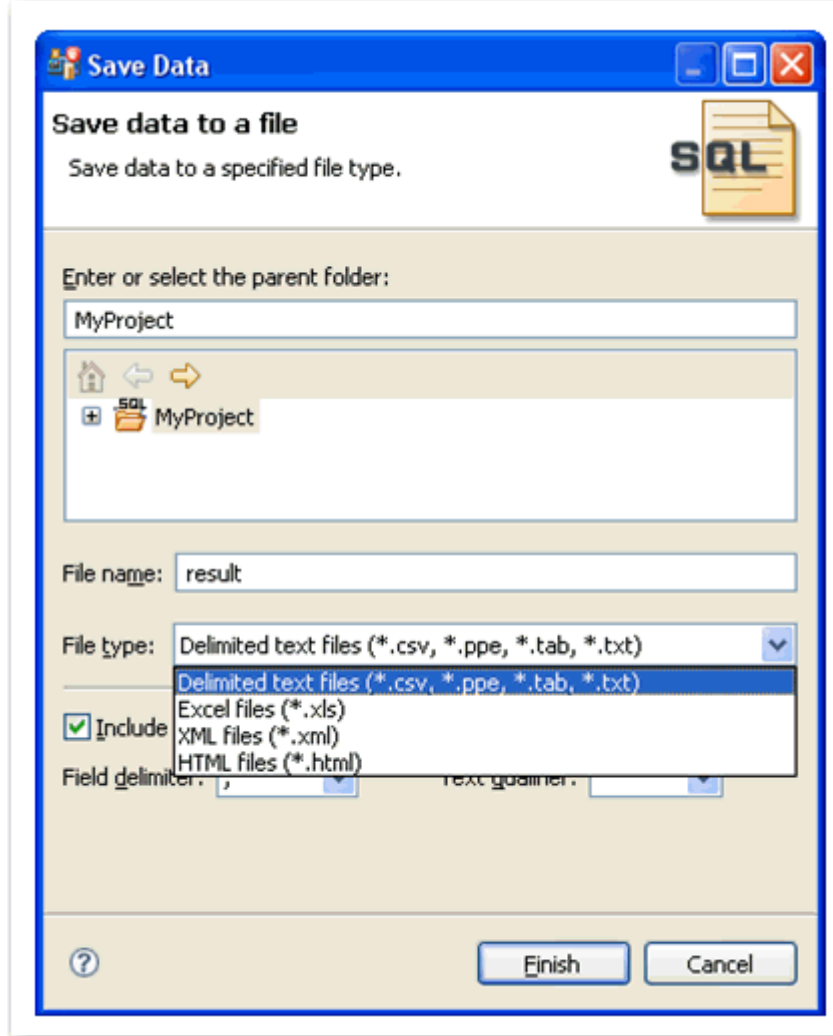
Once you execute a file, the results are displayed in the Results view. Here, you can examine the outcome of the execution process, as well as save the results in other file formats, as needed.

You can view results in the following formats:

- HTML
- XML
- TXT

To save results

1. Right-click on the **Results** view and select **Save Data**. The **Save Data** dialog appears.



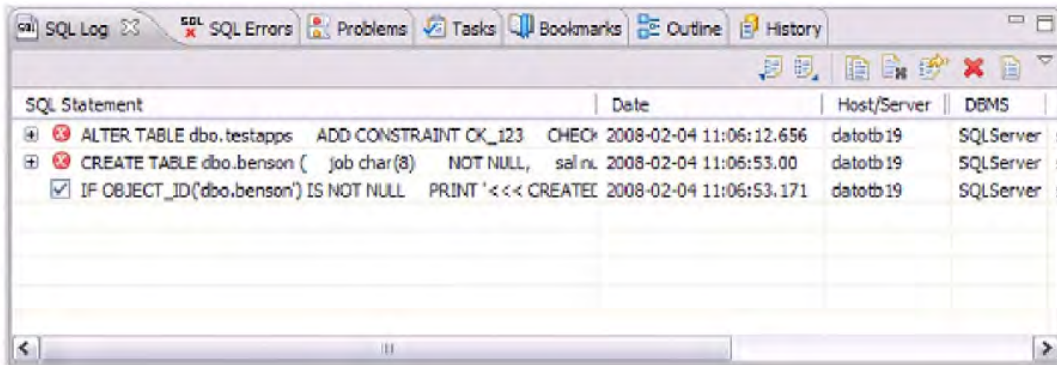
2. Select the project name to which you want to save the results, enter a file name, choose the file parameters, and then choose a file format from the drop down menu. You can select delimited text file, Excel, XML, or HTML file formats.
3. Click **Finish**. The results are saved in the directory location and format that you specified.

Troubleshooting

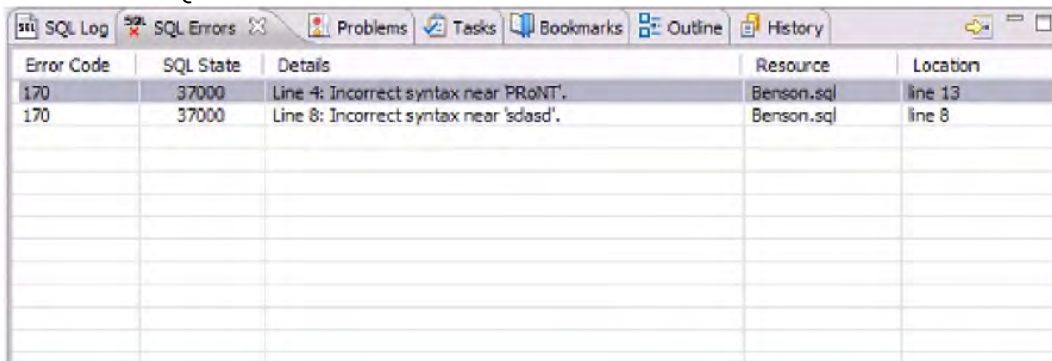
DB Optimizer contains a number of views used exclusively to log and monitor the SQL development process.

- The **SQL Log** captures all SQL commands executed by SQL Editor and the system. SQL Log entries are listed by SQL Statement name, Date issued, Host/Server, Service, User, Source, and the Time (in milliseconds) it

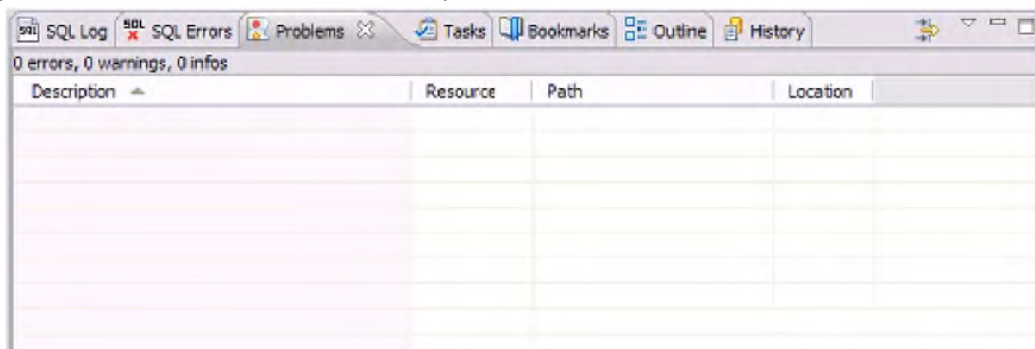
took to execute the command.



- The **SQL Errors** log automatically logs all SQL errors encountered when SQL commands are executed through DB Optimizer. Errors are listed by Error Code, SQL State, error Details, Resource, and the Location of the error in the SQL file.



- The **Problems** view captures syntactic and semantic errors and warnings in the files of the workspace. These entries typically take the form of error messages or warnings issued by the system over the course of a procedure execution. Problems are organized by Description (which indicates the type of problem logged), Resource, file Path, and Location. Using the Problems view, you can apply quick fixes to issues that DB Optimizer detects, as well as locate other problems that have similar attributes.



See Also


- [View log details](#)
- [Maintain logs](#)
- [Filter logs](#)
- [Import and export error logs](#)
- [Find and fix SQL code errors](#)
- [Find and fix other problems](#)

View log details

The SQL Error Log and Problems views contain functionality that enable you to view details regarding individual log entries, and in some cases, locate or fix those issues automatically.

To view details about SQL Errors entries

Right-click the error whose details you want to view, and then select **SQL Error Details**.



The SQL Error Details dialog provides information about the specified SQL error. Additionally, you can double-click the error to view the problem code in SQL Editor.

To view details about Problems

Right-click the entry whose details you want to view, and then select **Properties**. The **Properties** dialog appears, summarizing the issue.

Maintain logs

The SQL Log and SQL Errors views both contain commands that enable you to save, restore, or otherwise move log entries into files outside of DB Optimizer. Additionally, both views also contain commands that enable the clearing of the view.

The current editor option displays user statements only as generated by the active editor.

To maintain log entries

All entries automatically captured by the Error Log are written to a file (.log suffix) that resides in the Workspace .metadata folder.

To remove all messages

Right-click in the SQL Log view and select **Clear Log Viewer**.

To delete the .log file

In the shortcut menu, select **Delete Log**. If entries are created after the Delete Log command is issued, DB Optimizer automatically generates a new .log file in the .metadata subfolder.

i You cannot recover old error log entries once you delete the .log file. To prevent data loss, archive the .log file via the Export command prior to deletion.

To clear the Error Log view without deleting the .log file

Select **Clear Log Viewer** from the shortcut menu. The View is cleared of entries, but these entries still exist in the .log file.

To restore the Error Log view based on the entries in the .log file

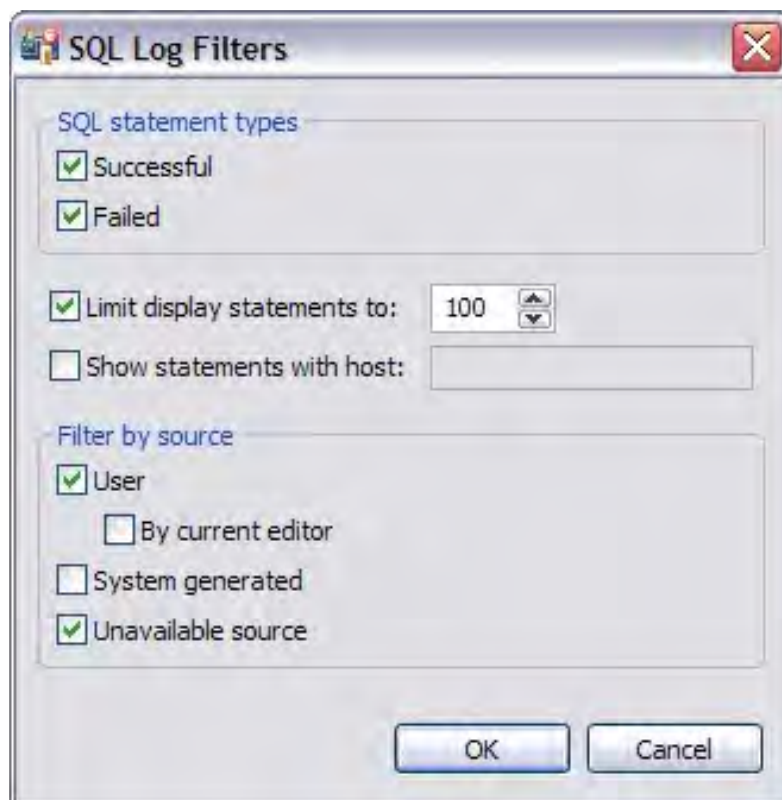
Select **Restore Log** from the shortcut menu. The View is restored based on the entries in the .log file.

Filter logs

You can apply filters to Problems, SQL Log, and the SQL Errors views to limit searches when troubleshooting and pinpointing specific processes within the system.

To filter the SQL Log


1. On the SQL Log view, select the Toolbar Menu icon  (the downward-pointing arrow in the right-hand corner of the view), and then choose Filters. The SQL Log Filters dialog appears.

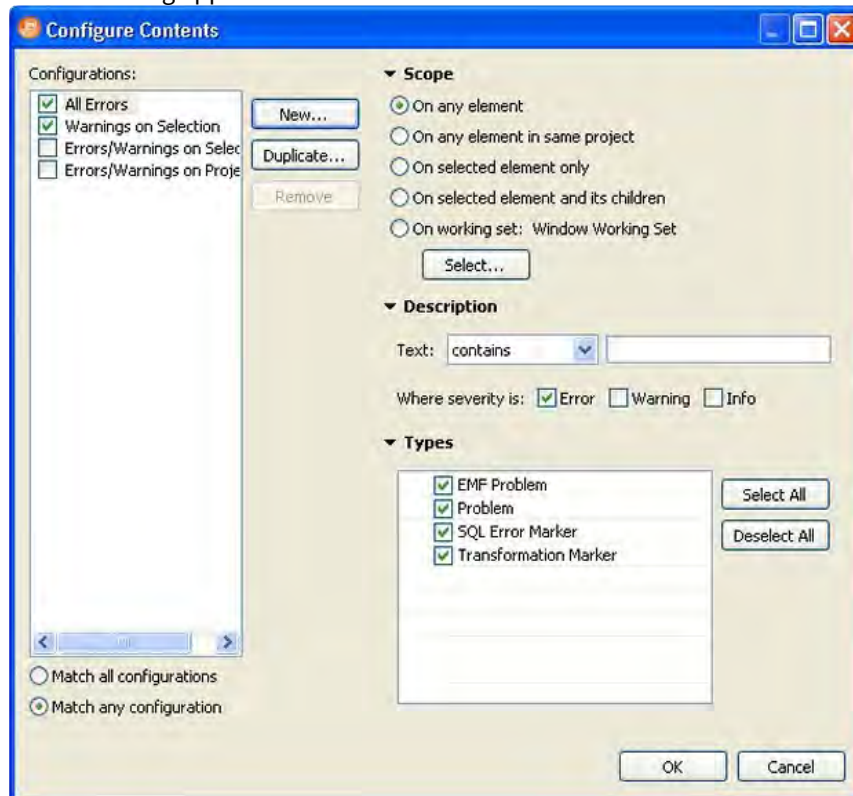



2. In the SQL Statement Types frame, select **Successful** or **Failed** to filter by the type of Error Log entries.
3. Select **Limit display statements** to indicate a maximum limit of the number of entries displayed in the Error Log, and enter the maximum entry value in the corresponding field.
4. Select **Show statements with host** to indicate that only entries from a specific data source are to be displayed, then type the name of the data source (as it appears in the Database Explorer) in the corresponding field.

5. In the **Filter by Source** pane, specify **User**, **System Generated**, or **Unavailable Source** to filter statements by the type of source from where they originated.

To view and filter the Problems log

1. On the Main Toolbar, click **Window > Show View > Other > General > Problems**.
2. On the Problems view, select the options icon  , and then choose **Configure Contents**. The Configure Contents dialog appears.



The Configure Contents dialog enables you to create multiple filter profiles that you can apply to the log by clicking the options icon , clicking Show, and then choosing the filter to apply. The Configurations panel on the left-hand side of the dialog displays all existing filter profiles stored in the Workspace. Selecting a configuration displays its filter parameters, and selecting the check box associated with its name enables the filter in the Problems view (only problems that match the criteria defined in the Filters dialog will appear in the view).

The ability to define different profiles enables the selection of multiple filter profiles. For each profile selected, the profile criterion is applied to the View, concurrently. You can filter problems by:

- Working Set
- Character String
- Problem Severity
- Problem Type
- A combination of the above four filter options

| Profile Criteria | Description |
|------------------|---|
| Working Set | <p>The options located in the Scope area of the dialog enable you to filter problems based on defined Working Sets. A Working Set is a collection of user-defined Project files that you can organize, as needed, in DB Optimizer. Select an option, and then click Select to define a Working Set to which the parameters apply. If no Working Sets exist, you need to define one or more via the New button on the Select Working Set dialog.</p> <p>Select one or more Working Sets to which you want the criteria to apply. If no Working Sets exist, or none suitably match the current filter criteria, click New or Edit to define a new Working Set, or edit an exist Working Set, respectively.</p> |
| Character String | <p>Use the Description list to select contains or doesn't contain, as needed, and type the character string in the field below the list. The Problems view is filtered to only contain, or omit, problem descriptions that fully or partially match the string value.</p> |
| Problem Severity | <p>In the Where severity is area, choose Error, Warning, Info or some combination of the three check boxes. Only entries whose severity matches the check boxes you have selected remain visible in the Problems view.</p> |
| Problem Type | <p>The options in the Types list on the right-hand side of the dialog enable you to filter problems by type. For example, deselect Problem to remove any system entries from the view, or deselect SQL Error Marker to remove any SQL code entries from the view.</p> |

Once you have defined and/or selected the appropriate filter profiles, the profiles will appear in the **Show** submenu in the Toolbar Menu of the Problems view. Select or deselect the profiles from the submenu, as needed.

Import and export error logs

Error messages are written to a file named .log located in the Workspace directory .metadata folder. This file can (and should) be cleared periodically via the Delete Log command to prevent performance issues with regards to system memory and file size. However, the Export command enables you to archive log files prior to deletion. The files created by the Export command can then be imported back into the Error Log as needed at a later point in time.

To export the SQL Log

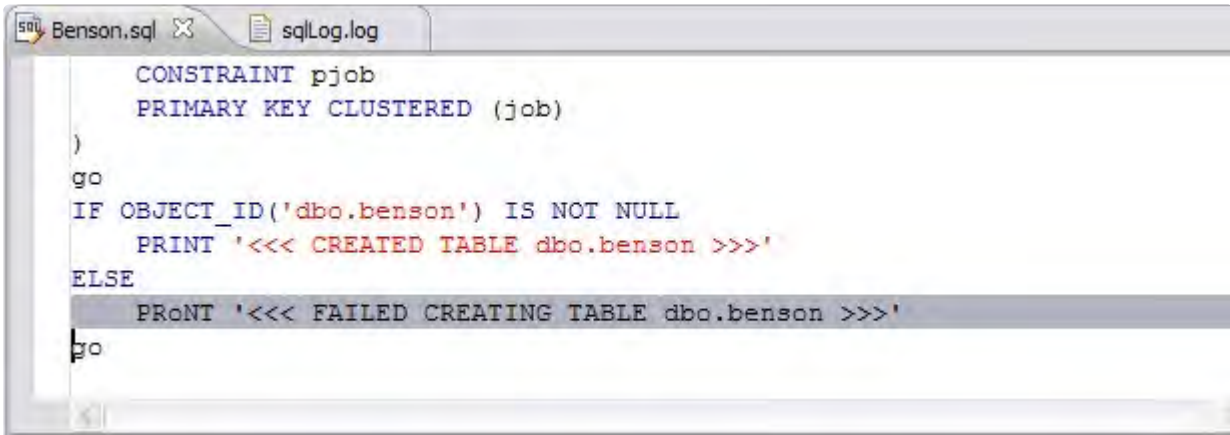
Right-click the SQL Log view, and then choose **Export Log**. The log is saved in the specified directory path with a .log extension.

To import the Error Log

Right-click the SQL Log view, and then choose **Import Log**. Select the previously exported .log file. The Error Log view is restored with the entries from the specified export file.

Find and fix SQL code errors

The SQL Errors view contains an option that enables you to navigate directly to the resource associated with an error entry.



```

Benson.sql  sqlLog.log
CONSTRAINT pjob
PRIMARY KEY CLUSTERED (job)
)
go
IF OBJECT_ID('dbo.benson') IS NOT NULL
PRINT '<<< CREATED TABLE dbo.benson >>>'
ELSE
PRINT '<<< FAILED CREATING TABLE dbo.benson >>>'
go
  
```

To navigate to the source of a SQL error entry

Right-click the entry to which you want to navigate, and then select **Go To**. The file to which the error applies automatically opens in a new instance of SQL Editor, and the line is highlighted in the window.

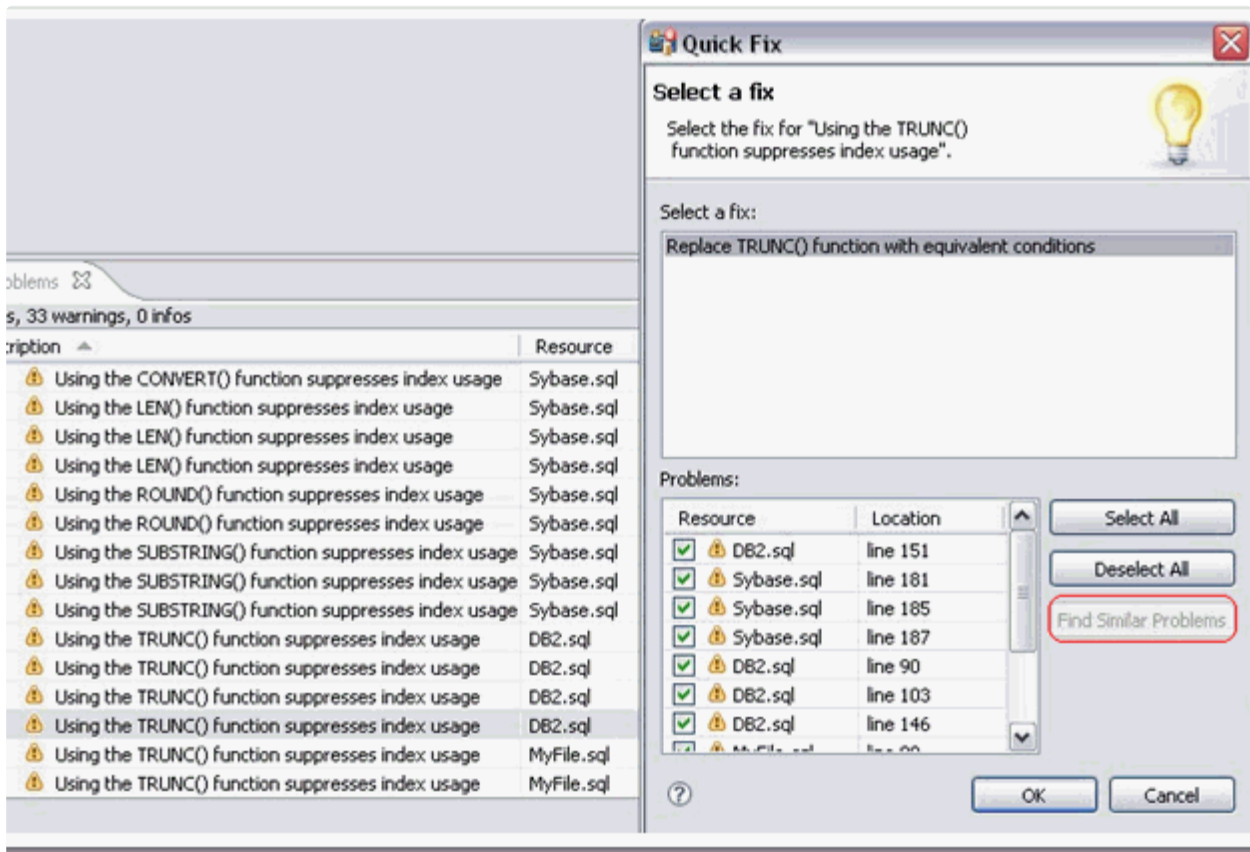
Find and fix other problems

By default, the Problems view organizes problems by severity. You can also group problems by type, or leave them ungrouped.

The first column of the Problems view displays an icon that denotes the type of line item, the category, and the description. Click the problem and DB Optimizer will open the SQL file and automatically highlight the line that triggered the issue.

You can filter Problems to view only warnings and errors associated with a particular resource or group of resources. You can add multiple filters to the view, as well as enable/disable them as required. Filters are additive, so any problem that satisfies at least one of the filters will appear.

Problems can sometimes be fixed via the Quick Fix command in the shortcut menu. The **Quick Fix** dialog enables you to apply a fix to a problem detected by the view. The dialog also provides a list of similar problems to the one you selected, and enables you to apply a fix to multiple problems at the same time.



To apply a quick fix to an issue in the Problem view

1. Right-click on a problem in the list, and then select **Quick Fix** from the menu. The **Quick Fix** dialog appears.
2. Select a fix from the list provided, and then click **OK**. DB Optimizer attempts to resolve the issue.

To find similar issues

1. In the **Quick Fix** dialog, click **Find Similar Problems**. The Problems list populates with all of the issues that are similar to your initial selection.
2. Use the check boxes beside the problems to select them, and then choose a fix and click **OK**. DB Optimizer attempts to resolve all of the specified issues.


Using profiling

For details on working with profiling, see the following topics:

- [Understanding profiling](#)
- [Understanding the interface](#)
- [Running a profiling session](#)
- [Configuring profiling](#)

Understanding profiling

Profiling continuously samples the data source to build a statistical model of the load on the database. Profiling can be used to locate and diagnose problematic SQL code and event-based bottlenecks. Additionally, profiling enables you to investigate execution and wait time event details for individual stored routines. Results are presented in the SQL Profiling Editor, which enables users to identify problem areas and subsequently drill down to individual, problematic SQL statements.

 DB Optimizer supports the Oracle Real Application Cluster (RAC). When you want to profile a RAC you can profile the entire cluster in one profiling session. For more information, see [Building profiling configurations](#).

Profiling filters out well performing, light weight SQL and collects information on heavy weight SQL. SQL that is heavy weight is either long running queries or queries that are short but run so often that they put load on the database

Profiler takes snapshots of user and session activity once a second and builds up a statistical model of the load on the database. The sampled data is displayed in three ways:

- Load on the database measured in average number of sessions active
- Top Activity - Top SQL, Event, and Sessions, Object I/O, and Procedures
- Details - Detail on a SQL statement, Session, Event, Object I/O or Procedure

The graph on the top of the screen shows the load on the database and can quickly indicate visually how the database is functioning. The database could be

- Idle
- Lightly loaded
- Heavily loaded
- Bottlenecked

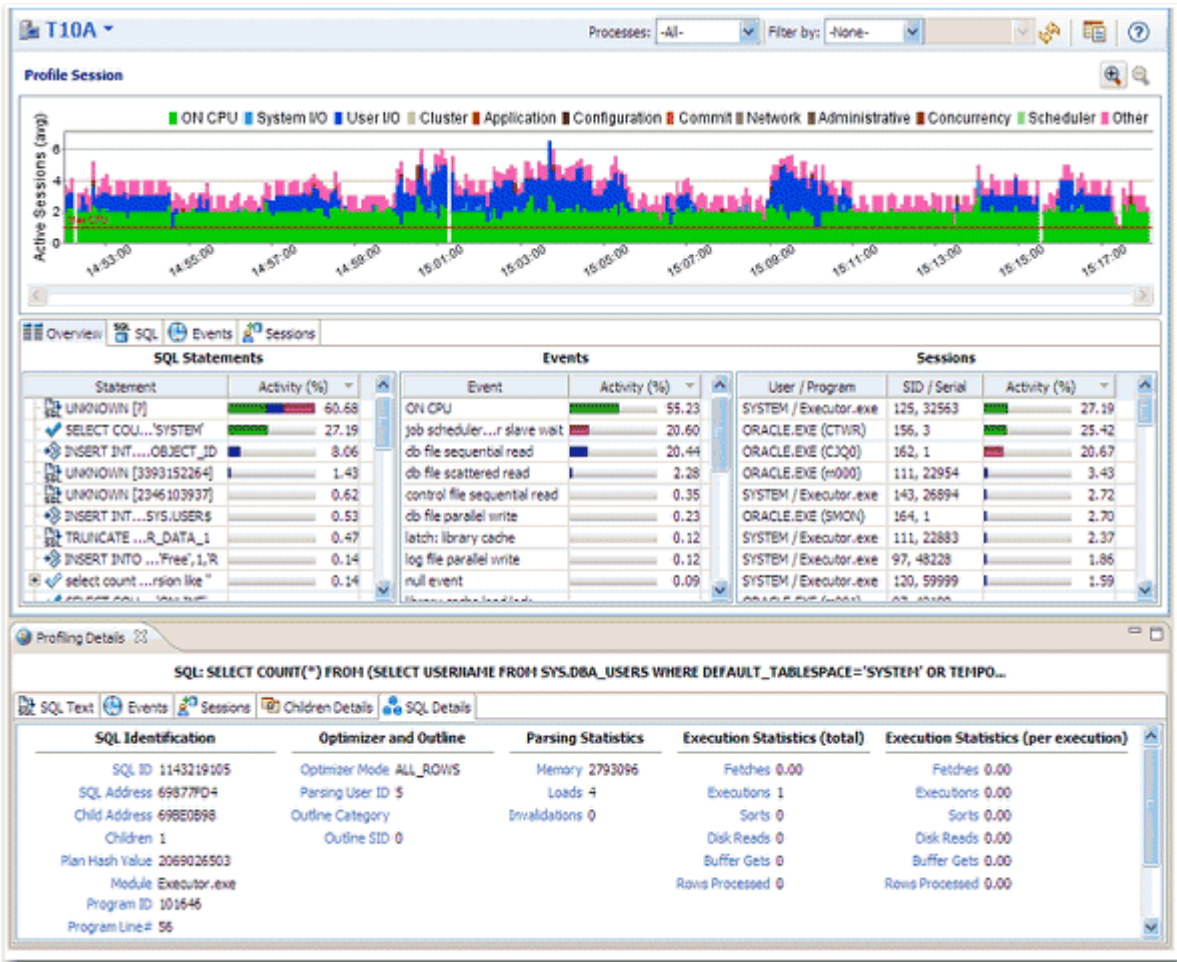
Problems can come from any one or more of four areas

- Machine CPU or Engine, slow disks (network)
- Application locks, invalid SQL
- Database cache sizes, log files, etc
- Inefficient SQL

Details of profiling sessions can be saved to an .oar file, which you can access through the SQL Project Explorer, or if you are profiling to an Oracle 9i, 10g, or 11g data source, you can see the profiling sessions in the Profiling Repository node of the data source in the Data Source Explorer, or in the Profiling Repositories node. Profiling direct to a repository means the session details are sent directly to the hard disk of the data source. If you have enough free disk space, you can collect data for a longer time than if you were profiling to a file which accumulates in memory until you save the session to disk.

Understanding the interface

The profiling interface is divided into three major parts: Load graph, Top Activity, and Profiling Details view.



- The **Load Graph** is located on the top section of the editor and provides a display of the overall load on the system. The bars represent individual aspects of the enterprise, and the view can be used to find bottlenecks.
- **Top Activity** is located on the middle section of the editor and displays where the load originates. Specifically, the top SQL statements, top events that the database spends time in, as well as the top activity sessions.
- The **Profiling Details View** is located on the bottom section of the editor and displays detailed information on any item selected in the middle section. For example, an SQL statement, an Event, or a Session.

The graphical portion of the profiling editor presents the distribution of sessions executed over the length of the profiling process, and those that were waiting in DBMS-specific events. It provides a first and most important step in identifying problem areas. Results can be viewed in real-time.

The Load Graph and Top Activity Section compose one view in the editor, while the Profiling Details view is a separate interface component that only activates when an item in the Top Activity Section is specified.

i Use a 1280 x 1024 monitor resolution when viewing profiling information. Smaller resolution sizes can obscure details in the view.

Running a profiling session

Profiling provides the continuous monitoring of a data source and builds a statistical model based of database load based on the user's state every second. The created profile can then be saved to file, and the data can be saved, analyzed, and optimized by importing statements into the tuning component and running a tuning job.

The following list provides the general workflow and overhead tasks, when attempting to monitor data sources and store query information.

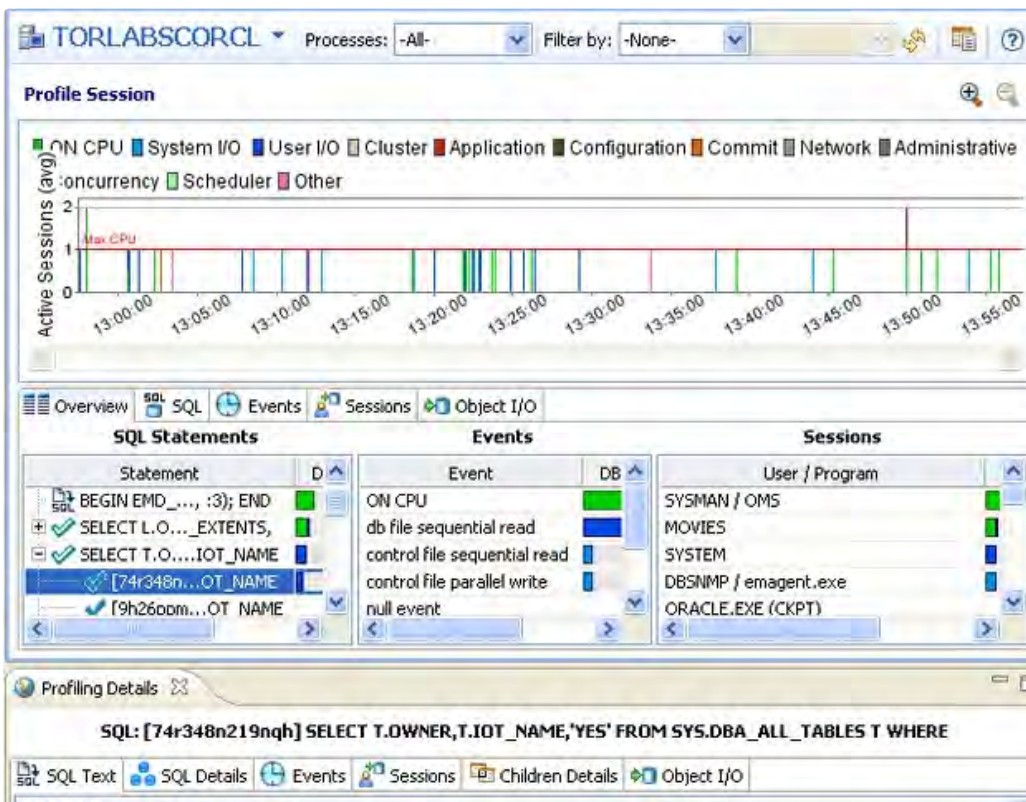
1. [Execute a profiling session](#)
2. [Work with Session Results](#)
3. [Creating profiling reports](#)
4. [Saving profiling sessions](#)
5. [Import statements to tuning](#)

In addition to the workflow tasks outlined above, the profiling interface also enables a number of important functions to help in statement analysis and diagnosis. This additional, or extra, functionality can be found in [Other profiling commands](#).

Furthermore, in some cases you will need to configure system variables and parameters in order to get the results you need from the application. See [Configuring DBMS properties and permissions](#) for more information on how to configure profiling and your registered data sources prior to running a session.

Execute a profiling session

Profiling is monitored and managed via profiling's three major interface components: the Load Chart, Top Activity Section, and Profiling Details view.




To execute a profiling session

1. In Data Source Explorer, right click on the data source you want to profile and select **Profile As** from the menu, and then choose **Data Source 1**.
2. In the **Profile Configurations** dialog, select the configuration to use for this profiling session. If you haven't already created a profile, see [Building profiling configurations](#) for a description of the profiling configuration options you can choose.


The profiling session begins. Alternatively, clicking the Profiling icon on the Toolbar automatically runs a profiling session for the last data source you selected.

Once a profiling session launches, it runs until you stop it. When a session has run for a length of time, you can then interpret and analyze the results. See [Working with session results](#).

To stop a profiling session

You can stop a profiling session at any time by clicking the **Stop** button [] in the upper left-hand side of the Profile Session screen or by clicking the Stop button in the Progress Window.

Executing a session from the Command Line

 This is not supported when using DB Optimizer InstantOn.

You can launch a profiling session from the command line using the following syntax:

```
dboptim
```

```
izer.exe
```

e

profile

ds:ROM*

L*ABORC

L10G_1

duratio

n:20

tofile:

C:

\ testpr

ofile.o

ar

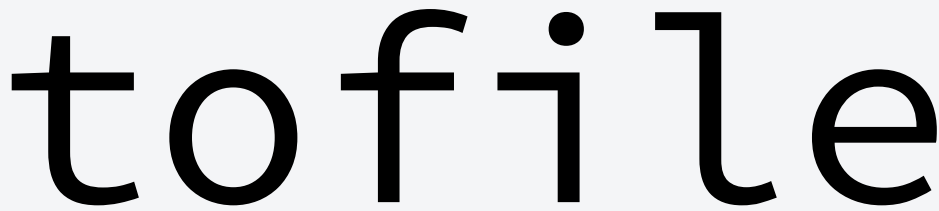
In the above command, the user has specified

```
ROM * L * A
```

```
BORCL10G
```

```
_1
```

as the data source, and indicates a profiling session of 20 minutes. The



tofile

variable

specifies the directory and name of the file to which the profiling session will be saved.

Working with session results

Results are displayed in the Profiling Session editor whenever a profiling session is executed. Results can appear in real time (if real time profiling is enabled) or once a session has finished its execution.

Results are displayed in the three major interface components of the editor, which you can use to analyze the overall efficiency and capacity of queries running on the data source, to various levels of detail:

The Profiling UI has three correlated sections:

- Selection in Chart fills the top activity section data, distributed in Overview/SQL/Events/ Sessions/Object I/O.
- Selection in any tab of Top Activity fills the Profiling Details with top selection type related data.

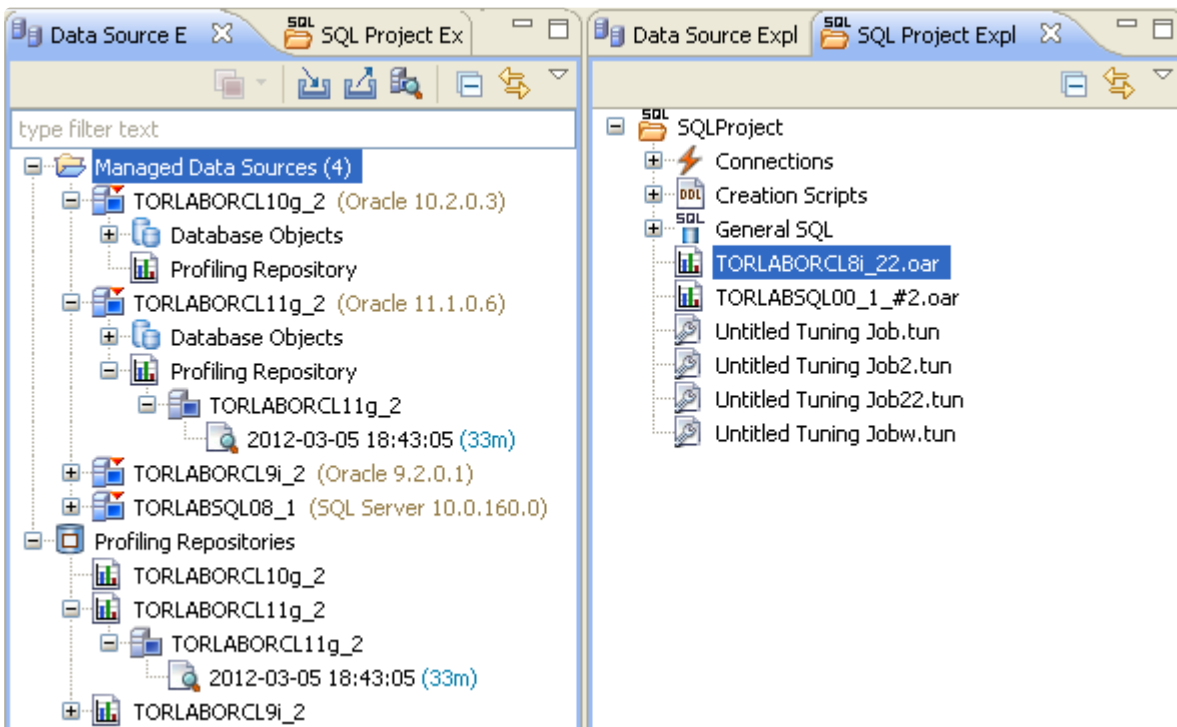
For more information, see:

- [Opening an existing profiling session](#)
- [Filtering results](#)
- [Analyze the load chart](#)
- [Analyze the Top Activity section](#)
- [Analyze profiling details](#)

Opening an existing profiling session

Saved profiling session data is stored in either an SQL Project or in a Profiling Repository on an Oracle data source. You can find profiling sessions saved as .oar files in the SQL Project Explorer. You can find profiling sessions saved to a profiling repository in the Data Source Explorer, either in the Profiling Repository node of the data source or in the Profiling Repositories.

To view a saved profiling session, locate it in either the **SQL Project Explorer** or in the **Data Source Explorer** and double-click the icon to open it in the Profiling Session window.




Filtering results


You can display filtered subsets of the original profiling results set for each section of profiling based on DBMS platform type:

- **IBM DB/2 for Windows, UNIX, and Linux:** Application, Creator ID, Cursor Name, Package Name, Statement Type, and User Name.
- **Microsoft SQL Server:** Application, Command, Database, Hostname, NT domain, Net Address, and User Name.
- **Oracle:** Processes (Background or User), Action, Application, Hostname, Module, Schema Name, and User Name. When profiling a RAC, there is also an instance filter that appears to let you limit the profiling results shown to a specific instance.
- **Sybase:** Application, Database, Host, IP Address, KPSID, SPID, and User Name.

You filter results using the filter controls in the upper, right-hand part of the profiling editor.




To filter profile editor results

1. Use the **Filter by** menu to select a filter type. The second menu becomes active based on your selection in the first menu.
2. Use the second menu to specify a value.
3. Click **Refresh** () to update the profiling details.
The profiling editor is updated to show only results associated with your choice.

Select **-None-** from the **Filter by** list to restore the unfiltered results.

Analyze the load chart

The Load Chart is located on the top section of the Profile Session editor and provides a display of the overall load on the system. The bars represent individual aspects of the enterprise, and the view is used to discover bottlenecks.



The most important part of the previous screenshot is the Average Active Sessions (AAS) graph. AAS shows the performance of the database measured in the single powerful unified metric AAS. AAS easily and quickly shows any performance bottlenecks on the database when compared to the Max Engines (for Sybase) or Max CPU (for Oracle) line. The Max Engines line is a yardstick for performance on the database. When AAS is larger than the Max CPU line, there is a bottleneck on the database. Bottleneck identification is that easy.

AAS or the average number or sessions active, shows how many sessions are active on average (over a 5 second range in DB Optimizer) and what the breakdown of their activity was. If all the users were running on CPU then the AAS bar is all green. If some users were running on CPU and some were doing I/O, represented by blue, then the AAS bar will be partly green and partly blue.

The line Max Engines or Max CPU represents the number of CPU processors on the machine. If we have one CPU/Engine then only one user can be running on the CPU/Engine at a time. If we have two CPUs/Engines then only two users can be on CPU at any instant in time. Of course users can go on and off the CPU/Engine extremely rapidly. When we talk about sessions on the Engines we are talking about the average number of sessions on CPU/Engine. A load of one session on the Engine thus would be an average which could represent one user who is consistently on the CPU/Engine or many users who are on the CPU for short time slices. When the number of Engines becomes a resource bottleneck on the database we will the average active sessions in CPU/Engine state go over the Max Engine/Max CPU line. The number of sessions above the Max Engine line is the average number of sessions waiting for CPU/Engine.

The Max CPU is a yardstick for performance on the database. The number of CPUs or Engine on the data source is information DB Optimizer obtains during the profiling process. However, sometimes the number of CPUs or Engines is not reported. In these cases, it might be desirable to change the default number of CPUs/Engines from one to a number more closely matching the actual system running the data source. You might also want to change the Max CPU/Engine line to reflect the performance impact of adding or removing a CPU or Engine from the system.

To change the Max CPU or Max Engine count in the Load Graph

1. From the Profile Session window, right-click anywhere on the AAS graph and select **Edit Engine Count** or **Edit CPU Count**.
2. In the **EngineCount** dialog that appears select **Useacustomvalue**, enter a new value, and then click **OK**. The AAS or Load Chart **MaxCPU** or **MaxEngine** line is updated immediately to reflect the change.

The Load Chart is designed as a high level entry point to profile session results. Subsequently, you can use the Top Activity and Profiling Details views to examine more detailed information on waiting and executing sessions over the length of the session. Alternatively, you can select one or more bars on the graph to populate the Top Activity section (and subsequently, the Details View) with information on a specific subset of the graph.

The Load Chart displays the distribution of waiting and executing sessions over the length of a profiling session.

- Time is displayed on the X axis. You can zoom in and zoom out on the graph via the icons in the upper right hand corner of the graph, once a profiling session is stopped.
- The Y axis shows the average number of sessions waiting or executing. Each supported platform has a specific set of wait event times.

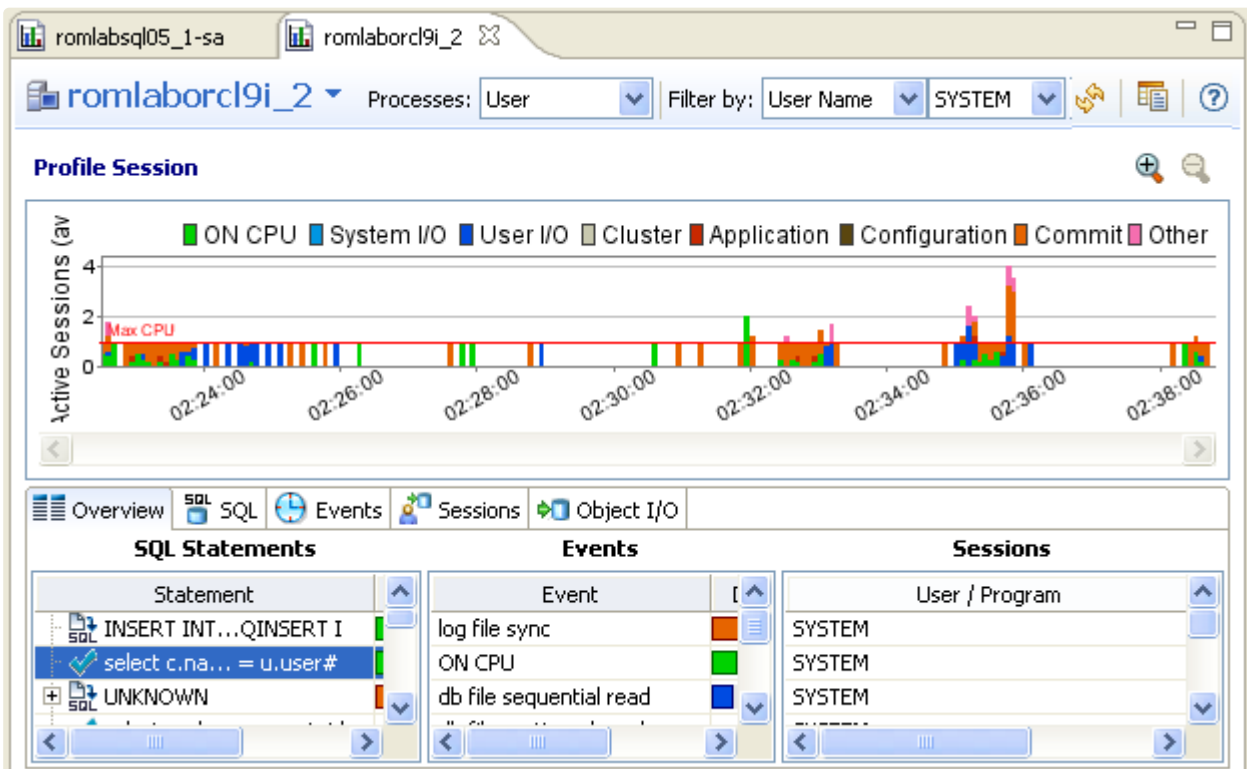
| DBMS | Wait Event Category |
|------------|--|
| IBM DB2 | Fetch, Cursor, Execution, Operation, Transaction, Connectivity, Lock, Other |
| Oracle | On CPU, System I/O, User I/O, Cluster, Application, Configuration, Commit, Other |
| SQL Server | CPU, Lock, Memory, Buffer, I/O, Other |
| Sybase | CPU, Lock, Memory, I/O, Network, Other |

* A chart legend displays a color and code scheme for executing and waiting session categories, in the upper right-hand corner of the view.

Analyze the Top Activity section

The Top Activity section is located in the middle section of the editor and displays where the load originates. Specifically, the top SQL statements, top events that the database spends time in, as well as the top activity sessions.

The Top Activity section is composed of a series of tabs that provide detailed statistics on individual SQL statements and sessions that were waiting or executing over the length of a profiling session.



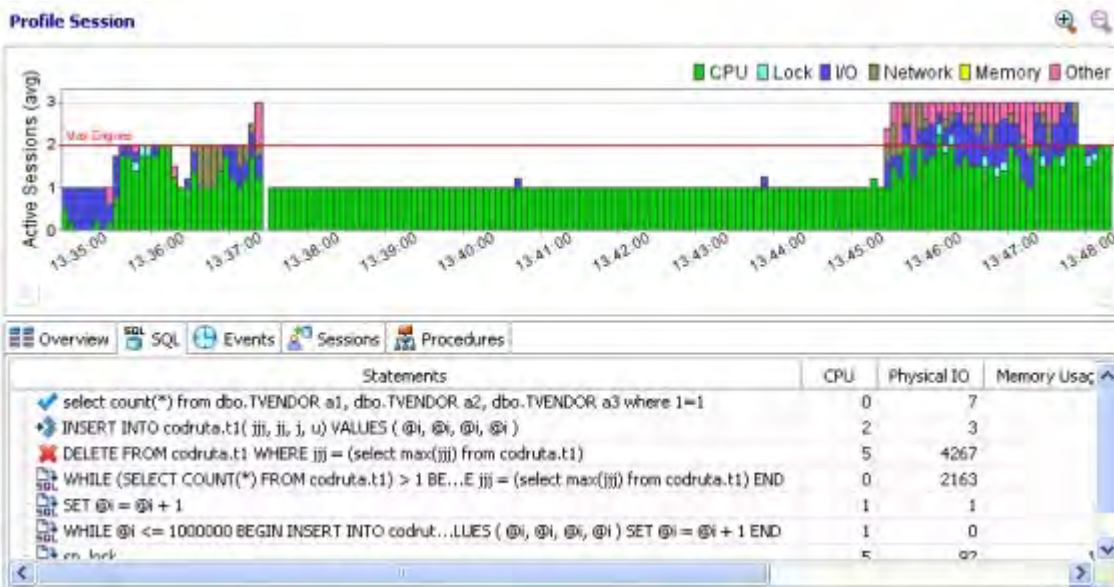
- The top **SQL** tab provides more detailed information than provided on the Overview tab, in terms of executing SQL statements and procedures. For more information, see [Top SQL Tab](#).
- The top **ExecutionActivity** (DB 2 Specific) tab provides details about the statements and procedures that ran. This is DB 2 specific. For more information, see [Top Execution Activity Tab \(DB2 Specific\)](#).
- The top **Events** tab displays information about wait events profiled by the execution process. For more information, see [Top Events Tab](#).
- The top **Sessions** tab displays information about sessions profiled by the execution process. For more information, see [Top Sessions Tab](#).
- The top **Blockers** tab (Oracle) displays information about blocking sessions. For more information, see [Top Blockers Tab \(Oracle Specific\)](#).
- The top **Object/I/O** tab (Oracle-Specific) tab does not appear in the Top Activity Section unless the data source being profiled is an Oracle platform. This tab displays information about the I/O profiled by the execution process. For more information, see [Top Object I/O Tab \(Oracle-Specific\)](#).
- The top **Procedures** tab (Oracle, SQL Server, and Sybase-specific) displays information about procedures observed during profiling. For more information, see [Top Procedures Tab \(Oracle, SQL Server, and Sybase Specific\)](#).

When you select any item from the Top Activity tabs, details are displayed in the Profiling Details view. The tabs that appear in **Profiling Details** will be different depending on the database platform and whether you selected a statement, session, or an event. This is to accommodate the parameter specifics of the item you selected.

Top SQL tab

The Profile editor's SQL tab shows a representation of all SQL statements that are executing or waiting to execute over the length of the profiling session or within the currently selected graph bars.

i The image below depicts results achieved for a Sybase database. The columns displayed on this tab differ depending on the database platform.



Statements

Statements can be grouped by type by right-clicking the view and selecting **Organize > By Type**.

i Statements are grouped when they differ only by their clause values. This enables the roll-up of SQL statements that only differ by a variable value. For example: select * from emp where empno=1; and select * from emp where empno=2. A '+' symbol appears beside rollup statements. You can click the symbol to expand and view the different statement predicates.

Additionally, the SQL tab displays two other groupings of statements:


| Group | Description |
|---------|--|
| OTHER | Includes all recognized statements other than INSERT, SELECT, UPDATE, DELETE, and MERGE statements. |
| UNKNOWN | Statements that are not recognized by the application. SQL Query Tuner has been improved to query the caching more often and more intelligently so that UNKNOWN appears less frequently in the Top SQL tab. The system queries the data source for SQL text in 15 second intervals. Unknown may still appear infrequently as the SQL text may have been removed by the database. |

All statements are displayed in a tree structure with the following statement components:

| Statement Component | Description |
|---------------------|---|
| Subject | The DML statement type (and FROM clause, as appropriate). |
| Predicate | The WHERE clause. |


| Statement Component | Description |
|---------------------|---|
| Remainder | Any statement component following the WHERE clause. |

For example, all statements with common subjects are shown as a single entry with multiple children; one child for each unique predicate. Predicates are similarly broken down by remainders.

 Right-clicking the **SQL** tab and selecting **Organize By** lets you choose between **Statement Type** grouping and **None**. The **None** option disables grouping by statement.

Statistics

Statistics are provided for statements and statement components. The statistics let you evaluate costs and spot wait event problems not just at the level of entire SQL statements, but also at the level of statement components. For each subject, predicate or remainder entry, the following statistics are provided:

 Columns displayed on the top SQL tab differ depending on the data source platform.

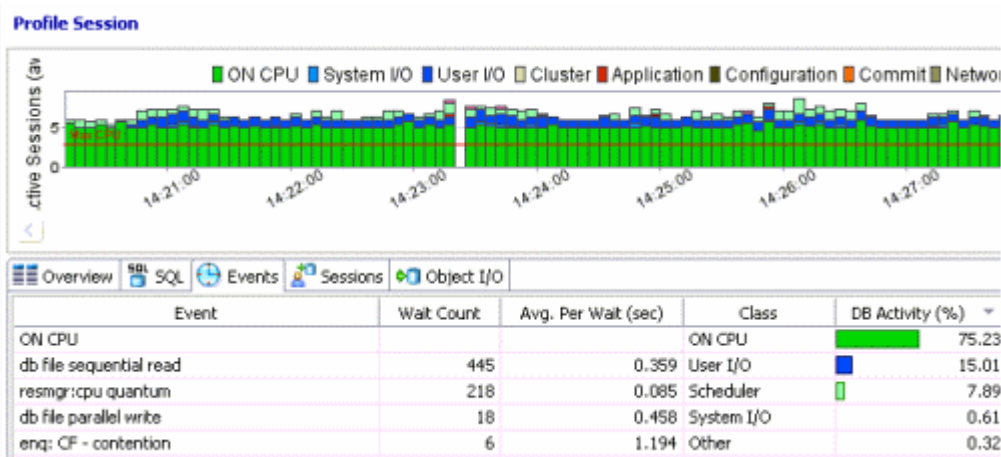
| Statistic | Shown for Platform | Notes |
|--------------------|---------------------------------|--|
| Executions | SQL Server, Oracle, Sybase, DB2 | The number of active executions for the statement or statement component over the length of the profiling session or the selected graph bars. |
| Avg. Elapsed (sec) | Oracle, DB2 | The average amount of time that elapsed while executing the statement during the profiling period. This column appears for only SQL Server, DB2 and Oracle data sources. |
| DB Activity (%) | SQL Server, Oracle, Sybase, DB2 | A graphical representation of the distribution of execution and wait time for the statement or statement component. |
| SQL ID | Oracle | The ID value of the SQL statement. This statistic only appears on Oracle data sources. |
| Child Number | Oracle | The child number in the database. This statistic only appears on Oracle data sources. |
| Parsing User ID | Oracle | The ID of the user who parsed the statement. This statistic only appears on Oracle data sources. |
| Plan Hash Value | Oracle | The execution value of the statement. This statistic only appears on Oracle data sources. |

Top Execution Activity tab (DB2 specific)

In addition to the statistics displayed on the Top SQL tab, DB2 platforms have an additional tab in the Profile Session editor named Execution Activity, which contains the following statistical rows: Rows Read, Rows Written, Fetch Count, Statement Sorts, Sort Time, and Sort Overflows.

Top Events tab

The Top Events tab displays information about wait events on the resources involved in the profiling process. This display is used to tune at the application or database configuration level. For example, if the top events are locks, then application logic needs to be examined. If top events are related to database configuration, then database setup should be investigated.



i The columns that display are data source-dependent. For example, the Wait Count and Avg. Per Wait (sec) columns display only for an Oracle data source.

Top Sessions tab

The Sessions tab provides information about individual sessions. This tab provides information about sessions that are very active or bottlenecked.

The screenshot shows the 'Sessions' tab in the Profile Session tool. It displays a table with the following columns: User Name, Program, SID, Serial#, Activity (%), Machine, and Session Type. The data is as follows:

| User Name | Program | SID | Serial# | Activity (%) | Machine | Session Type |
|-----------|-------------------|-----|---------|--------------|---------------------------|--------------|
| SYSTEM | Executor.exe | 125 | 32563 | 27.19 | EMBARCADERO\ROWEBITA01 | USER |
| | ORACLE.EXE (CTWR) | 156 | 3 | 25.42 | TORLABORCL10G_1 | BACKGROUND |
| | ORACLE.EXE (CJQ0) | 162 | 1 | 20.67 | TORLABORCL10G_1 | BACKGROUND |
| | ORACLE.EXE (m000) | 111 | 22954 | 3.43 | TORLABORCL10G_1 | BACKGROUND |
| SYSTEM | Executor.exe | 143 | 26894 | 2.72 | EMBARCADERO\ROWEANITUCA01 | USER |
| | ORACLE.EXE (SMON) | 164 | 1 | 2.70 | TORLABORCL10G_1 | BACKGROUND |
| SYSTEM | Executor.exe | 111 | 22883 | 2.37 | EMBARCADERO\ROW/SNOVAC01 | USER |
| SYSTEM | Executor.exe | 97 | 48228 | 1.86 | EMBARCADERO\ROW/SNOVAC01 | USER |
| SYSTEM | Executor.exe | 120 | 59999 | 1.59 | EMBARCADERO\ROW/SNOVAC01 | USER |
| | ORACLE.EXE (m001) | 97 | 48199 | 1.54 | TORLABORCL10G_1 | BACKGROUND |
| | ORACLE.EXE (SMON) | 164 | 1 | | TORLABORCL10G_1 | BACKGROUND |

Top Blockers tab (Oracle Specific)

The Blockers tab provides details on sessions holding blocking locks.

| User Name | Program | SID | Serial# | Blocking (%) | Machine | Session Type | Client Info |
|-----------|----------------|-----|---------|--------------|--------------|--------------|-------------|
| SYSTEM | JDBC T...lient | 91 | 113 | 79.10 | rowcb...iu02 | USER | |
| | ORA...WR) | 166 | 1 | 15.67 | ROML...RS03 | BACK...UND | |
| SYSTEM | JDBC T...lient | 124 | 26549 | 4.48 | rowcb...iu02 | USER | |
| | ORA...ON) | 161 | 1 | 0.75 | ROML...RS03 | BACK...UND | |

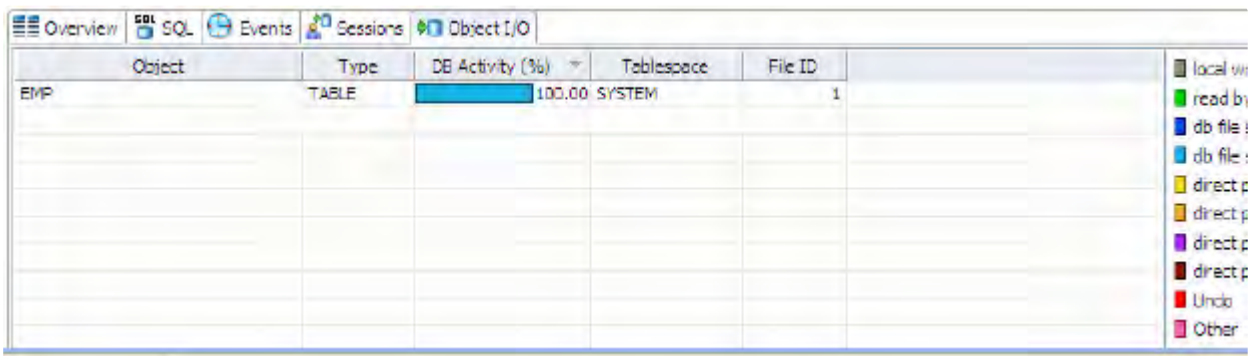
The following parameters are displayed on the Blockers tab:

| Value | Description |
|--------------|---|
| User Name | The user name under which the session was run. |
| Program | The name of the executable under which the session was run. |
| SID | The SID value of the session. |
| Serial | The serial number of the machine from which the session executed. |
| Blocking (%) | A graphical representation of the percentage of total blocked sessions being blocked by a blocking session. |
| Machine | The machine name and network location of the machine from which the session executed. |
| Session Type | The type of session. |
| Client Info | The name/type of the client from which the session initiated. |

 For more detailed information, see [Viewing Details on the Blockers Tab \(Oracle\)](#).

Top Object I/O tab (Oracle specific)

The **Object I/O** Tab is specific to the Oracle data source platform, and displays information about Oracle I/O loads on the profiled data source.

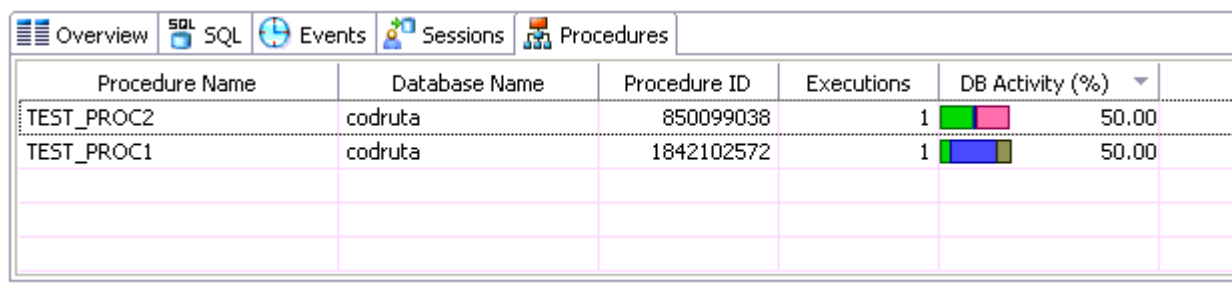


The following parameters are displayed on the I/O tab:

| Value | Description |
|-----------------|---|
| Object | The name of the data source object affecting the Oracle I/O. |
| Type | The object type. For example, table, partition, or index. |
| DB Activity (%) | Use the color chart on the right-hand side of the I/O tab to view the I/O load on the data source during the profiling session. |
| Tablespace | The name of the tablespace where the object resides. |
| File ID | The unique ID value of the file from where specified object resides. |

Top Procedures tab (Oracle, SQL Server, and Sybase specific)

The **Procedures** tab is specific to Oracle, SQL Server, and Sybase data source platforms. It displays information about Procedure loads on the profiled data source.



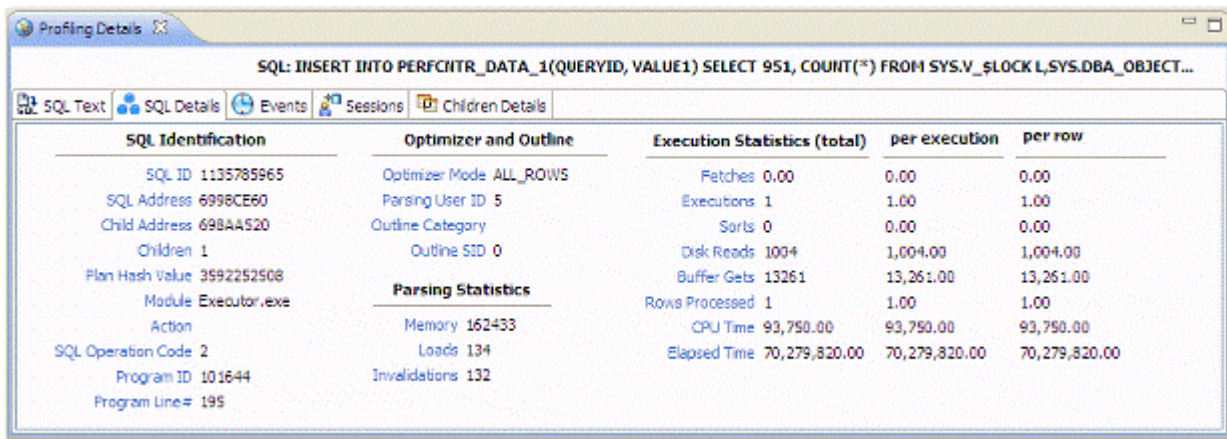
The following parameters are displayed on the Procedures tab:

| Value | Description |
|----------------|---|
| Procedure Name | The name of the procedure affecting the database performance. |

| Value | Description |
|---|--|
| Database Name (SQL Server, Sybase only) | The name of the database where the procedure resides. |
| Oracle (Oracle only) | The owner of the schema in which the procedure resides. |
| Procedure ID | A unique ID created when the procedure is created. |
| Executions (SQL Server, Sybase only) | The number of times the procedure was executed during the session. |

Analyze profiling details

The Profiling Details view displays detailed information on any item selected in the Top Section View, such as an SQL statement, an Event, a Session or a Procedure.



Depending on the data source platform you have specified, the tabs that appear in the view will be different, in order to accommodate the parameter specifics of the statement you have selected.

Depending on the top activity selected and the profiled platform types, some tabs may not be available.

Statement selected

When a **Statement** is selected, the following Profile Detail tabs are available.

| Tab Name | Description | Supported Platform | | | |
|----------|---|--------------------|--------|------|------------|
| | | Oracle | Sybase | DB 2 | SQL Server |
| SQL Text | Displays the full code of the selected SQL statement. | Yes | Yes | Yes | Yes |

| Tab Name | Description | Supported Platform | | | |
|-----------------------|--|--------------------|--------|------|------------|
| | | Oracle | Sybase | DB 2 | SQL Server |
| SQL Details | Provides details on statement, like execution statistics. | Yes | No | Yes | No |
| Events | Provides database activity details about events the statement is associated with. | Yes | Yes | Yes | Yes |
| Sessions | Shows which sessions executed this statement. | Yes | Yes | Yes | Yes |
| Blockers | Shows which sessions held blocking locks against the session associated with this statement. | Yes | No | No | No |
| Children Details | Lists all copies of the cursor or SQL query, if Oracle has cached multiple copies of the same statement. | Yes | No | No | No |
| Object I/O | If the SQL query has done physical I/O, then these are the objects, such as tables, and indexes that were read to satisfy the query. Temporary objects with not have values in Object and Type columns. | Yes | No | No | No |
| Procedures | Shows which procedures contain the selected statement. | No | Yes | No | Yes |
| Bind Variable Details | Shows bind variable information for SQL captured during the Profiling session. | Yes | No | No | No |

Event selected

When an **Event** is selected, the following Profile Detail tabs are available.

| Tab Name | Description | Supported Platform | | | |
|----------|--|--------------------|--------|------|------------|
| | | Oracle | Sybase | DB 2 | SQL Server |
| SQL | Shows which SQL statements waited on this event. | Yes | Yes | Yes | Yes |


| Tab Name | Description | Supported Platform | | | |
|------------|---|--------------------|--------|------|------------|
| | | Oracle | Sybase | DB 2 | SQL Server |
| Sessions | Provides information about the sessions associated with the event. | Yes | Yes | Yes | Yes |
| Blockers | Shows which sessions held blocking locks against the session associated with this event. | Yes | No | No | No |
| Procedures | Shows which procedures contain the selected event. | No | Yes | No | Yes |
| Raw Data | Raw data that was sampled from the database, specifically the following: <ul style="list-style-type: none"> • Sample time • SID • Serial # • User name • Program • Sql ID • P1 • P2 • P3 | Yes | No | No | No |
| Analysis | Displays for "buffer busy waits" and "cache buffer chains latch" waits. The analysis shows data and documentation to assist in solving these bottlenecks. | Yes | No | No | No |

Session selected

When a **Session** is selected, the following Profile Detail tabs are available.

| Tab Name | Description | Supported Platforms | | | |
|----------|--|---------------------|--------|------|------------|
| | | Oracle | Sybase | DB 2 | SQL Server |
| Sessions | Provides parameters regarding the session. For example, database server connection information, and data regarding the client tool and application. | Yes | Yes | No | Yes |

| Tab Name | Description | Supported Platforms | | | |
|------------|---|---------------------|--------|------|------------|
| | | Oracle | Sybase | DB 2 | SQL Server |
| Blockers | Shows which sessions held blocking locks while this session was active. | Yes | No | No | No |
| SQL | Shows which SQL statements this session ran. | Yes | Yes | Yes | Yes |
| Events | Shows which events this session waited on. | Yes | Yes | Yes | Yes |
| Procedures | Shows which procedures ran the selected session. | No | Yes | No | Yes |

 When right-clicking on a SQL statement in the Top Activity Section in Profiling, if the SQL statement is run by a different user than the user who is running DBO, than the User Mismatch dialog appears, with an example of the following message: "This query was executed by [SOE] and you are currently connected as [system]. We recommend you reconnect as [SOE] to tune the SQL. Would you like to continue anyway?" This message indicates that the statement is being tuned by a user other than the user who originally ran the query, and tables may be missing based on the different schemas. Click **OK** to run the query, or click **Cancel** and run tuning under the original user.

Blocking Session selected

When a **Blocking Session** is selected, the following Profiling Detail tabs are available.

| Tab Name | Description | Supported Platform | | | |
|------------------|--|--------------------|--------|------|------------|
| | | Oracle | Sybase | DB 2 | SQL Server |
| Blocked Sessions | Provides identifier and V\$SESSION session information on the sessions being locked by the blocking session. | Yes | No | No | No |
| Session Details | Provides parameters regarding the session. For example, database server connection information, and data regarding the client tool and application. | Yes | No | No | No |
| SQL | Shows the SQL statements associated with the lock. | Yes | No | No | No |
| Events | Shows which events the blocking session waited on. | Yes | No | No | No |

Procedure selected

When a **Procedure** is selected, the following Profile Detail tabs are available.

| Tab Name | Description | Supported Platform | | | |
|----------|--|--------------------|--------|-----|------------|
| | | Oracle | Sybase | DB2 | SQL Server |
| SQL Text | Shows the SQL text of the selected procedure. | No | Yes | No | Yes |
| SQL | Shows which SQL statements this procedure ran. | No | Yes | No | Yes |
| Events | Shows which events the selected procedure waited on. | No | Yes | No | Yes |
| Sessions | Provides parameters regarding the session. For example, database server connection information, and data regarding the client tool and application. | No | Yes | No | Yes |

This section also addresses the following topics:

- [Viewing Details on the SQL Tab](#)
- [Viewing Details on the Sessions Tab](#)
- [Viewing Details on the Events Tab](#)
- [Viewing Details on the Procedures Tab](#)

Viewing details on the SQL tab

In the **Top Activity Session**, selecting a statement entry on the **SQL** tab displays information in the **Profiling Details** view. The graph portion and details on the event category tabs on the new editor pertain only to the selected statement. Additionally, new tabs become available:

- **SQL Text tab:** Shows the full code of the SQL statement. For more information, see [SQL Text](#).
- **SQL Details tab:** Displays execution details. This tab is only displayed for Oracle data sources. For more information, see [SQL Details](#).
- **Events tab:** Displays information about the events the selected statement is associated with.

For more information, see [Events](#).

- **Sessions tab:** Displays information about the sessions that the selected statement is associated with. This tab is displayed only for Oracle data sources. For more information, see [Sessions](#).
- **Procedures tab:** Displays information about the procedures that contain the selected statement. This tab is displayed only for SQL Server and Sybase data sources. For more information, see [Procedures](#).

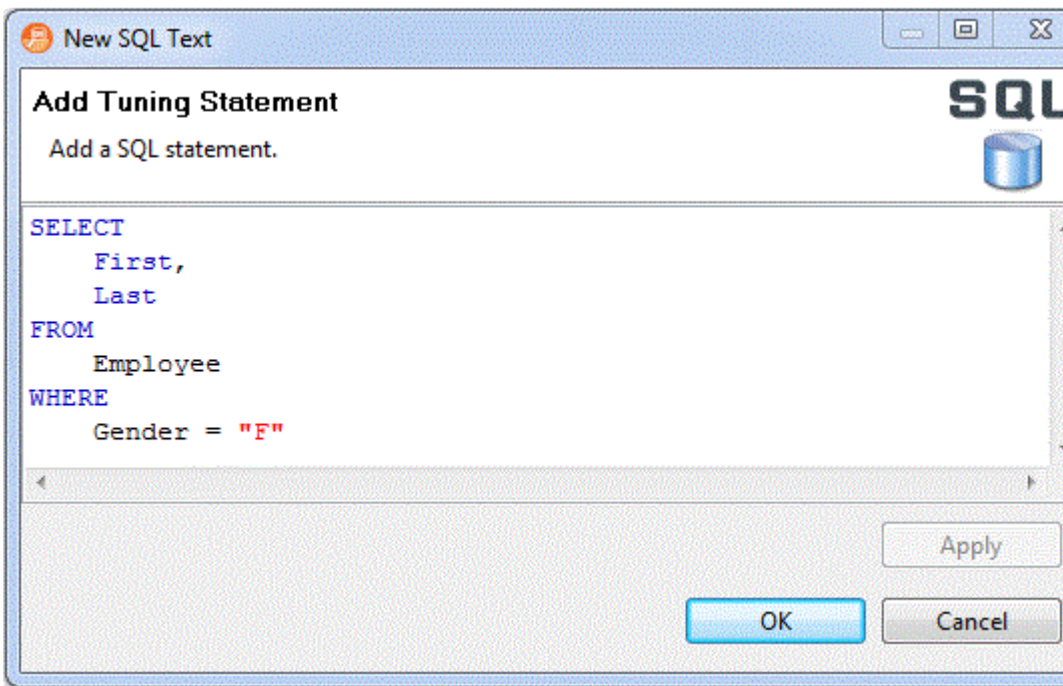
To select a SQL tab statement entry

- On the **SQL** tab, click on a statement with no child nodes or on a leaf node in the statement structure.

The new profiling editor page opens, as reflected by the bread crumb trail at the top left of the editor. You can continue to drill down into the statement, as needed.

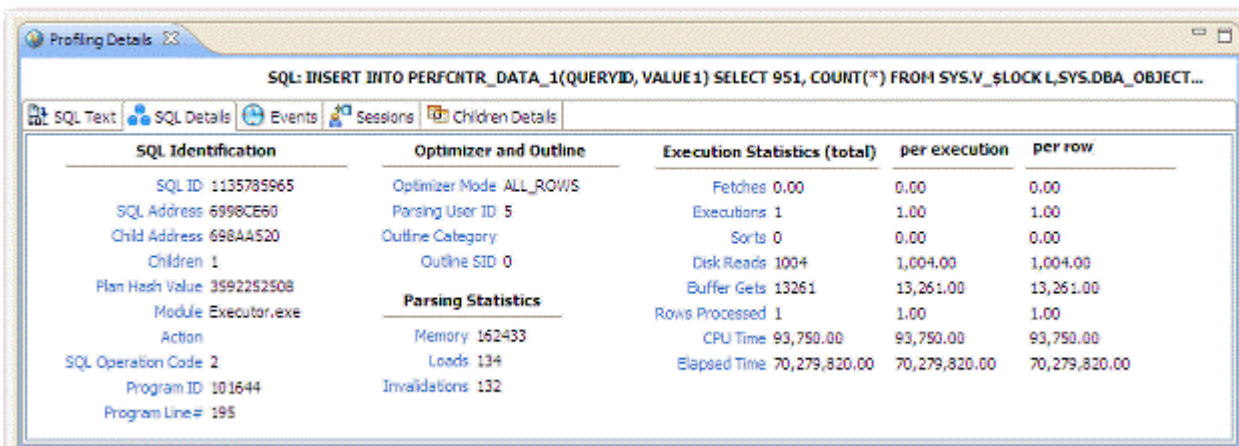
SQL Text

The SQL Text tab displays the full code of the SQL statement.



SQL Details

The SQL Details tab provides information and the execution of the statement and other information related to how it is running. It is only applicable to Oracle data sources:



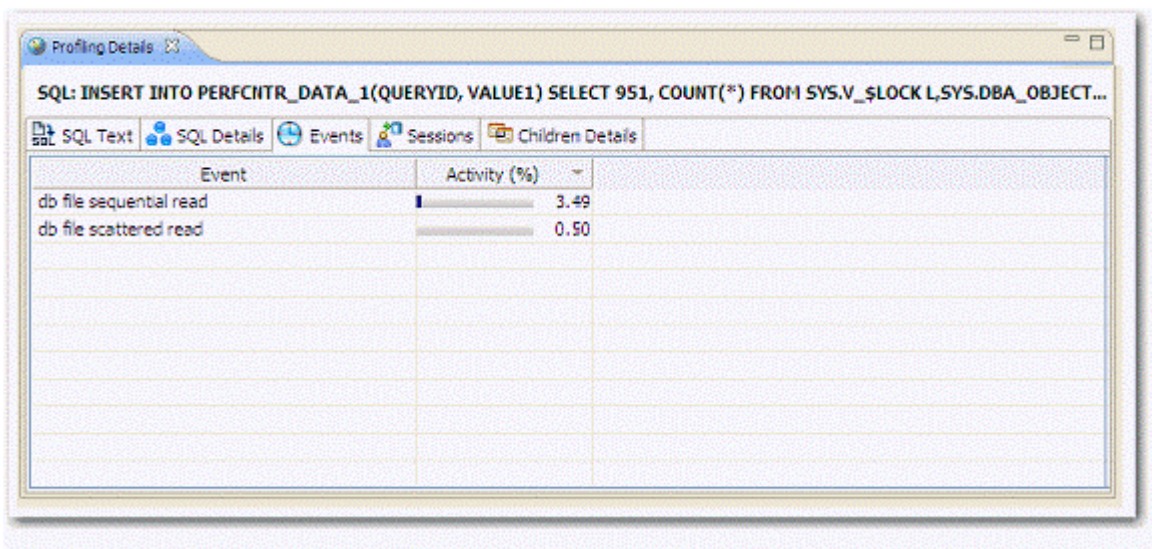
SQL Details include:

| Parameters | Description |
|------------------------------|--|
| SQL Identification Values | The SQL ID value of the statement. |
| Optimizer and Outline Values | Optimizer-specific values pertaining to the parsing user ID value and outline SID. |

| Parameters | Description |
|----------------------|---|
| Parsing Statistics | Information regarding memory, loads, and invalidation values. |
| Execution Statistics | The execution statistics of the statement. This category includes disk reads, buffer gets, rows, and values that represent CPU and elapsed time. |

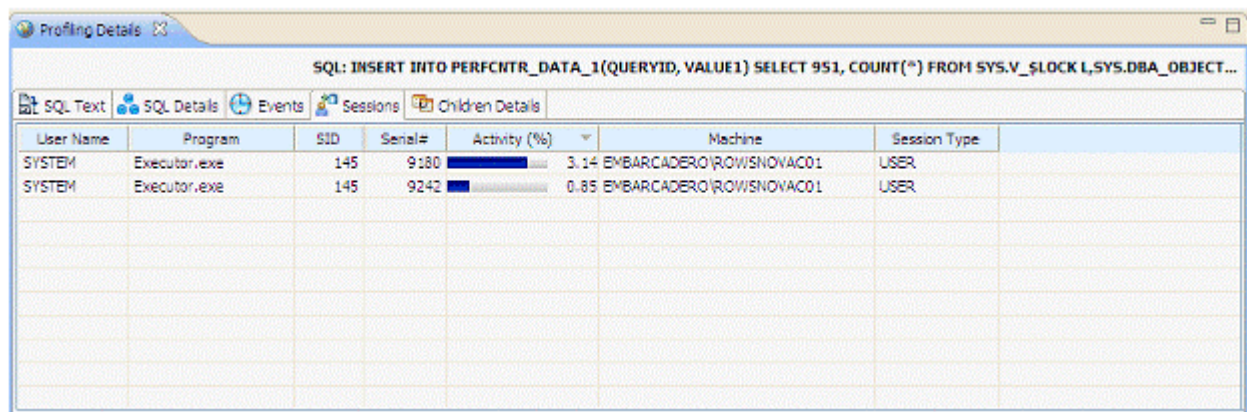
Events

The Events tab provides details about the events that the statement is associated with.



Sessions

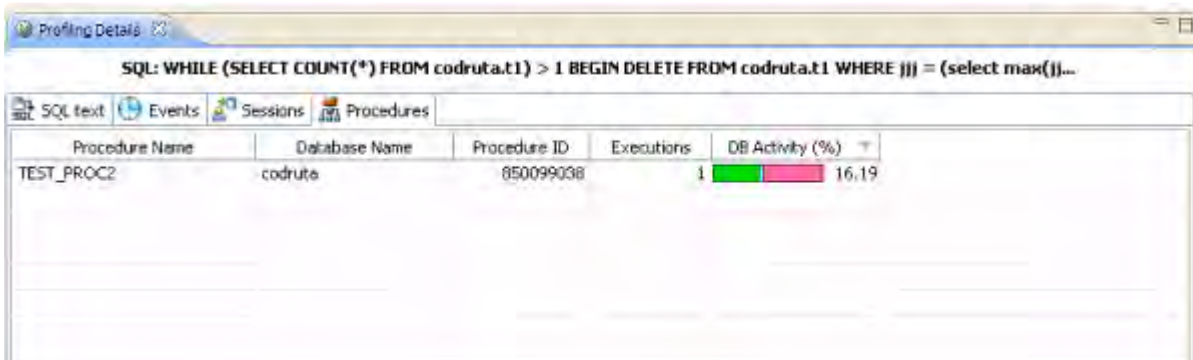
The Sessions tab provides information about any sessions the statement is associated with:



Session details include information on different parameters, depending on the platform. For example, on Oracle platforms, the following parameters are displayed: User Name, Program, SID, Serial #, Activity (%), Network Machine Name, and Session Type.

Procedures

The Procedures tab provides information about any procedures containing the selected statement.



The following parameters are displayed on the Procedures tab:

| Value | Description |
|-----------------|---|
| Procedure Name | The name of the procedure that contains the selected statement. |
| Database Name | The name of the database where the procedure resides. |
| Procedure ID | The unique ID value of the file where the specified procedure resides. |
| Executions | The number of times the procedure was executed. |
| DB Activity (%) | Use the color chart on the right-hand side of the Procedures tab to view the procedures load on the data source during the profiling session. |

Viewing details on the Sessions tab

In the **Top Activities Section**, selecting a statement entry on the **Sessions** tab displays information in the **Profiling Details** view. The graph portion and details on the event category tabs on the new editor pertain only to the selected statement. Additionally, new tabs become available.

Selecting an event type entry on an event category tab opens a new profiling editor page. The graph portion and details on the **Sessions** tab and event category tabs on the new editor page pertain only to the selected wait event and to SQL statements that waited in that event.

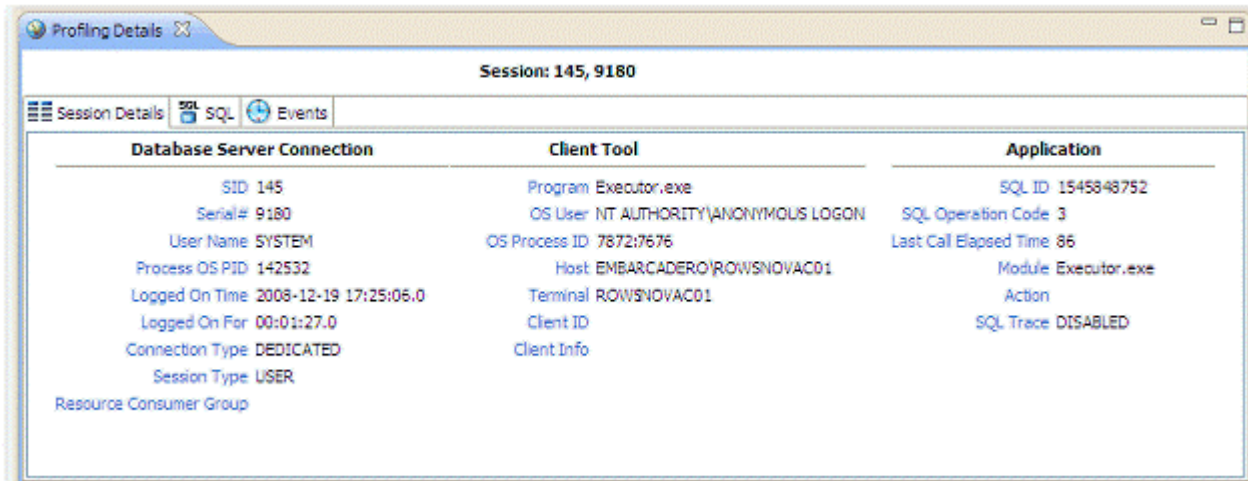
- **Session Details tab:** Shows system details about the selected session. For more information, see [Session Details](#).
- **SQL tab:** Displays information about the SQL files that the selected session is associated with. This tab only appears on Oracle platforms. For more information, see [SQL](#).
- **Events tab:** Displays the time and parameter information about the selected session. For more information, see [Events](#).
- **Procedures tab:** Displays the details of any procedures run in the selected session. For SQL Server and Sybase data sources only. For more information, see [Procedures](#).

Session Details

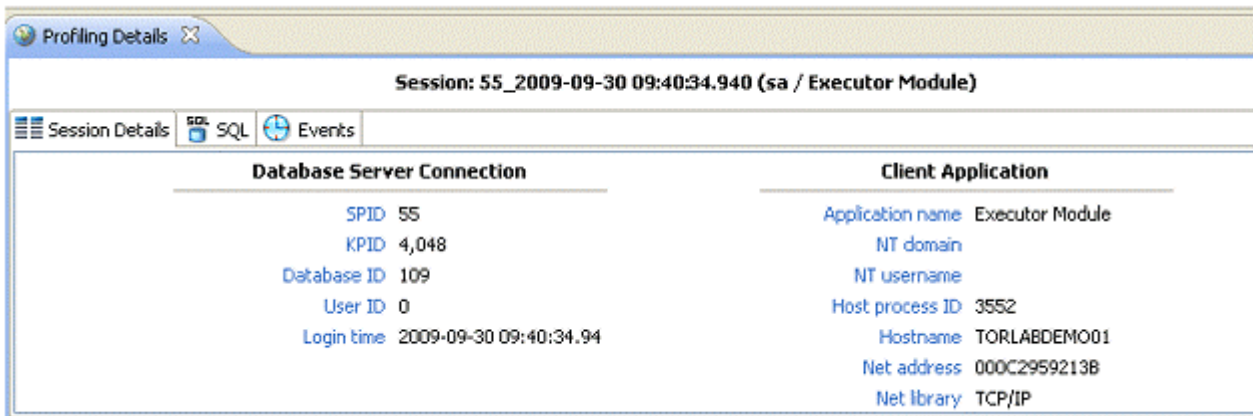
The Session tab provides further information about the selected session. The following are examples of the session details provided for different platforms.

 The fields that display vary depending on the database platform.

Oracle profiling details

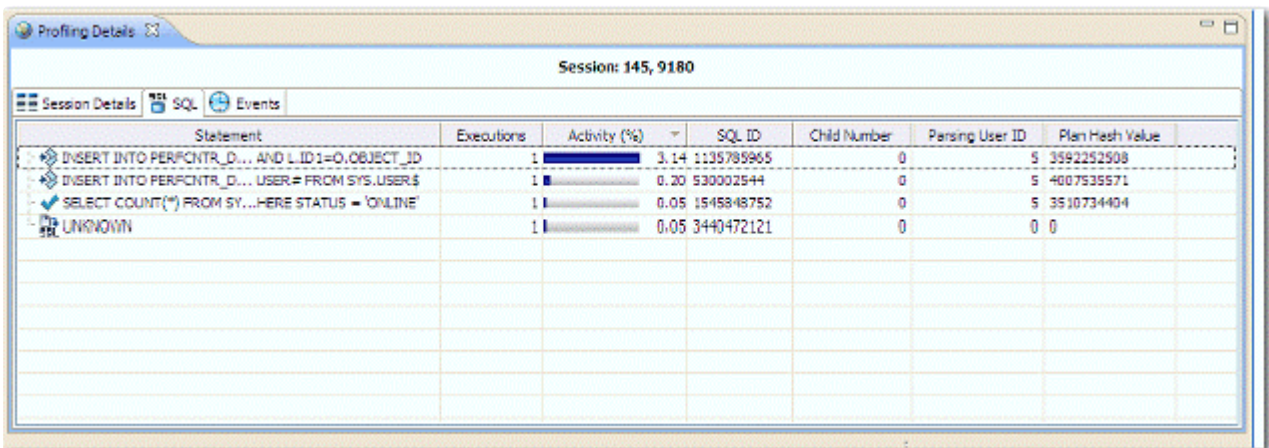


Microsoft SQL Server



SQL

The SQL tab displays information about the statements associated with the session.

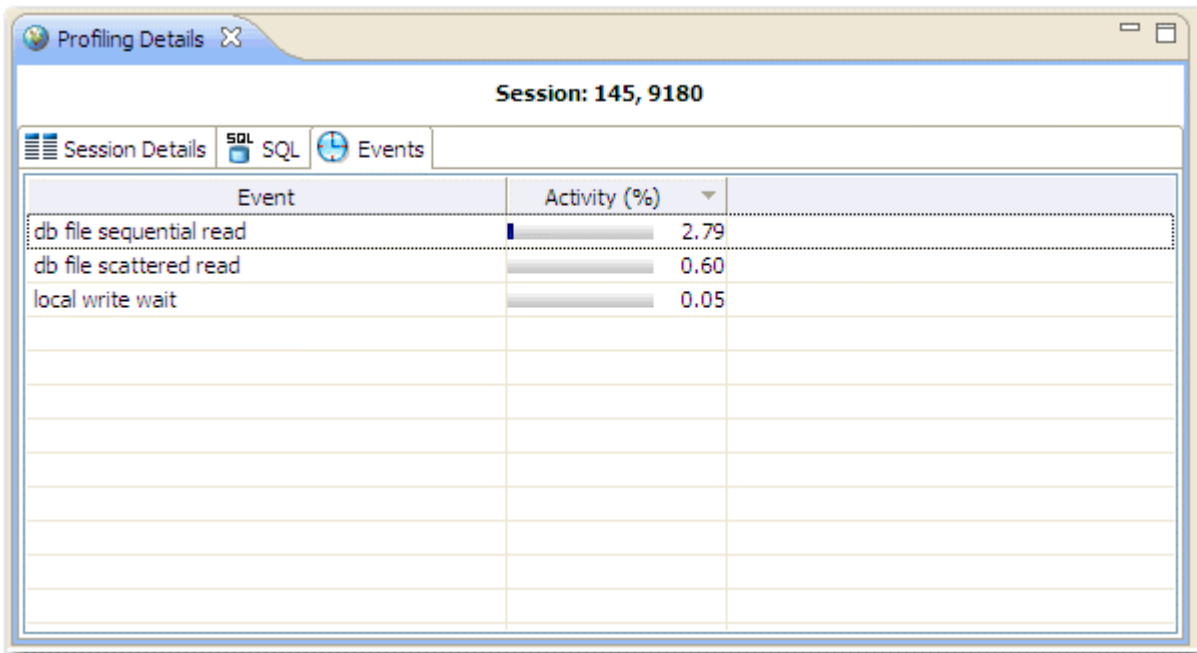


SQL statements are listed by the following parameters:

| Value | Notes |
|-----------------|---|
| Statement | The name of the statement. |
| Executions | The number of times the statement was executed during the session. |
| Activity (%) | A graphical representation of the distribution of execution and wait time for the statement or statement component. |
| SQL ID | The SQL ID value of the statement. |
| Child Number | The child number in the database. |
| Parsing User ID | The ID of the user who parsed the statement. Plan Hash ValueThe execution value of the statement. |
| Plan Hash Value | The execution value of the statement. |

Events

The Events tab provides details about the events that the session is associated with.

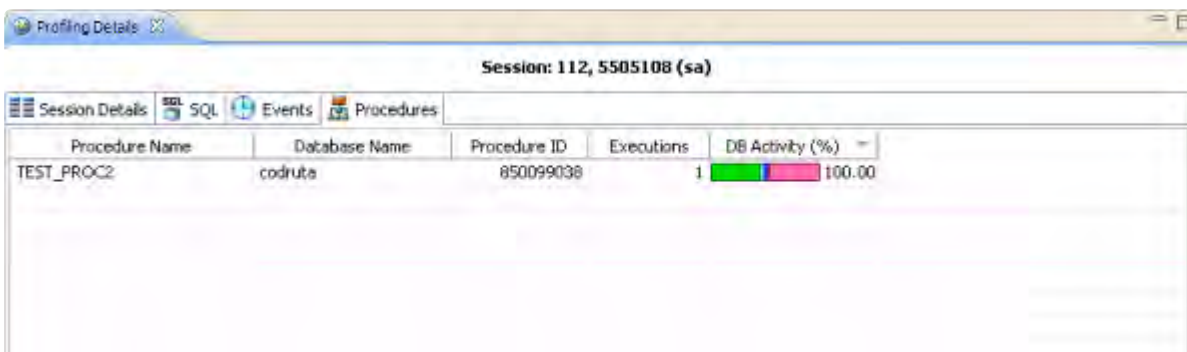


Events are listed by the following values:

| Value | Notes |
|--------------|---|
| Event Name | The name of the event. |
| Activity (%) | A graphical representation of the distribution of execution and wait time for the statement or statement component. |

Procedures

For SQL Server and Sybase data sources only, the Procedures tab provides details about the procedures that the session is associated with




The following parameters are displayed on the Procedures tab:

| Value | Description |
|-----------------|---|
| Procedure Name | The name of the procedure that ran during the selected session. |
| Database Name | The name of the database where the procedure resides. |
| Procedure ID | The unique ID value of the file where the specified procedure resides. |
| Executions | The number of times the procedure was executed during the session. |
| DB Activity (%) | Use the color chart on the right-hand side of the Procedures tab to view the procedures load on the data source during the profiling session. |

Bind Variable details

For Oracle data sources, profiling captures the bind variables and their attributes. Select an SQL statement in the Profiling Session and the details of the captured bind variables for that statement are displayed here.



The following parameters are displayed on the Bind Variable Details tab:

| Value | Description |
|---------------|--|
| SQL ID | SQL identifier used by the data source. |
| Child Number | A new child number is generated for the SQL ID of the query whenever the plan changes, for example the value of a bind variable is changed, and the query is executed again. |
| Position | The position of the variable within the SQL text. For example, given the query, <code>select * from T1 where C1 = :a and C2 = :b and C3 = :c and C4 = :d</code> , the position of a is 1, b is 2, c is 3 and d is 4. |
| Variable Name | The name of the variable. |
| Variable Type | The data type of the variable. |

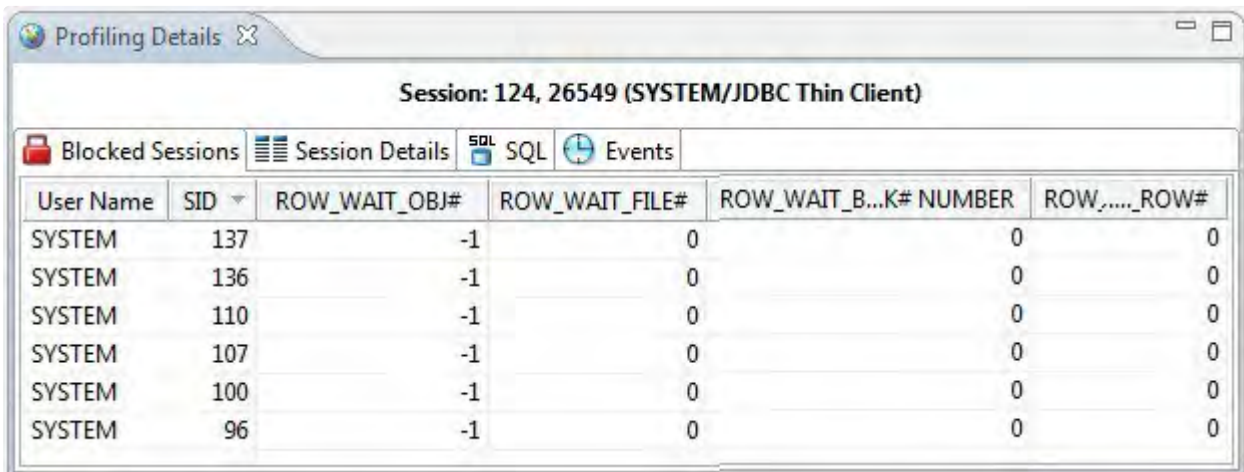
| Value | Description |
|----------------|----------------------------|
| Variable Value | The value of the variable. |

Viewing details on the Blockers tab (Oracle)

In the **Top Activities Section**, selecting an entry on the **Blocked Sessions** tab displays information on sessions holding blocking locks in the **Profiling Details** view.

Blocked Sessions

The Blocked Sessions tab provides general information on blocked sessions and the details identifying the specific row locked.

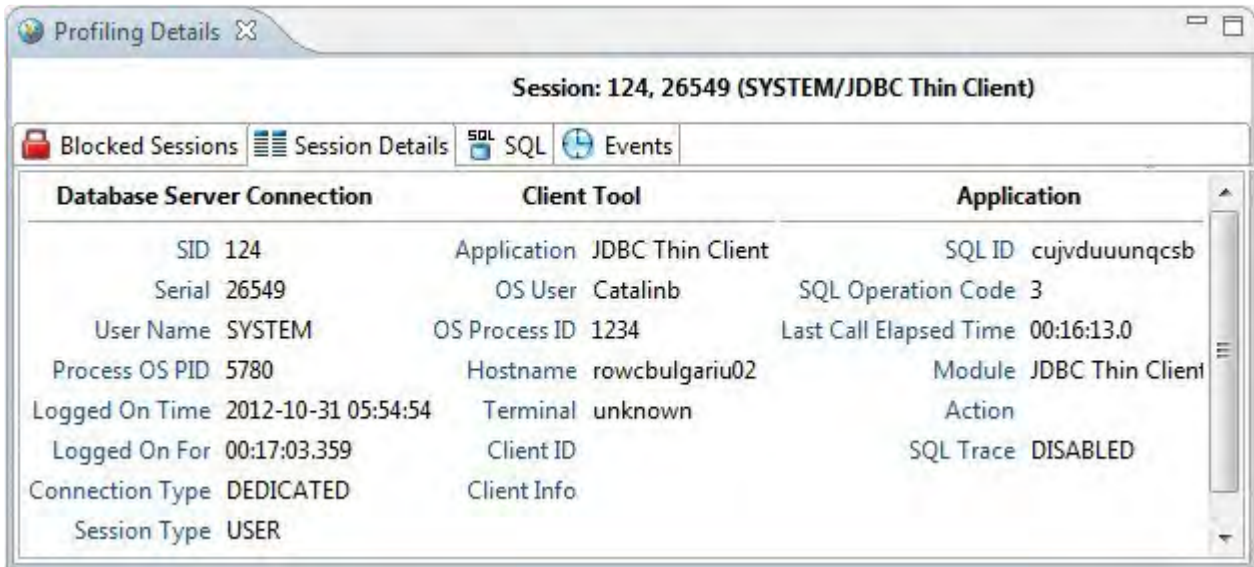


This tab provides the following columns for each blocked session:

| Value | Notes |
|-----------------|---|
| User Name | The user name under which the blocking session was run. |
| SID | The SID value of the blocking session. |
| ROW_WAIT_OBJ# | Object ID of the table containing the row specified in ROW_WAIT_ROW#. |
| ROW_WAIT_FILE# | Identifier of the datafile containing the row specified in ROW_WAIT_ROW#. |
| ROW_WAIT_BLOCK# | Identifier of the block containing the row specified in ROW_WAIT_ROW#. |
| ROW_WAIT_ROW# | The current row being locked. |

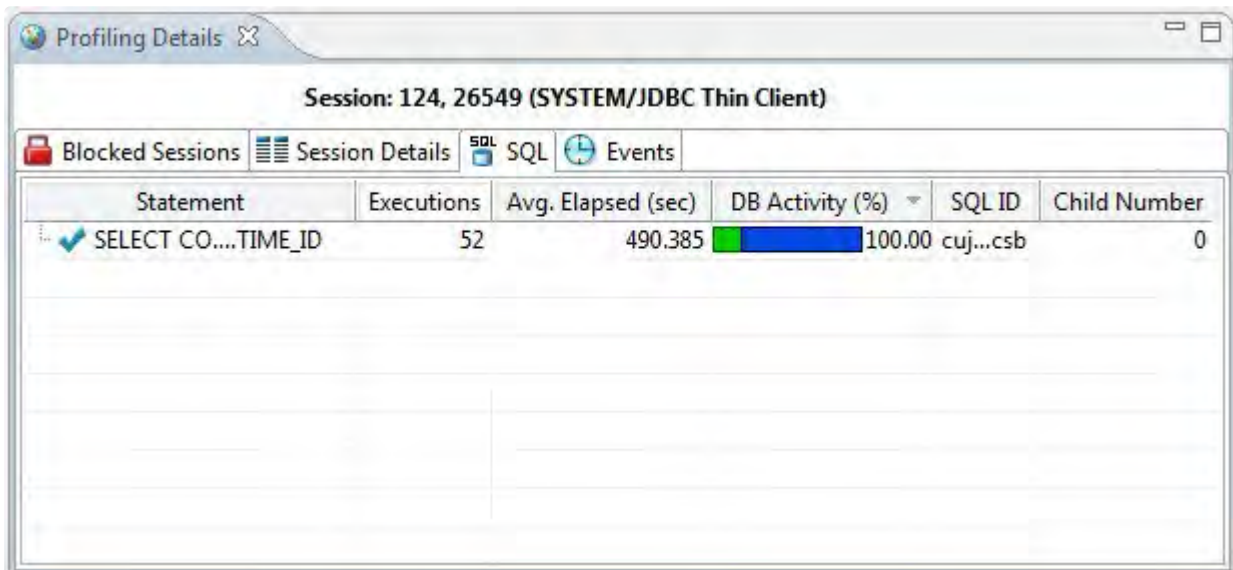
Session Details

The Session Details tab provides information on the server connection, client, and application associated with the blocking session.



SQL

The SQL tab displays information about the statements associated with the blocking session.



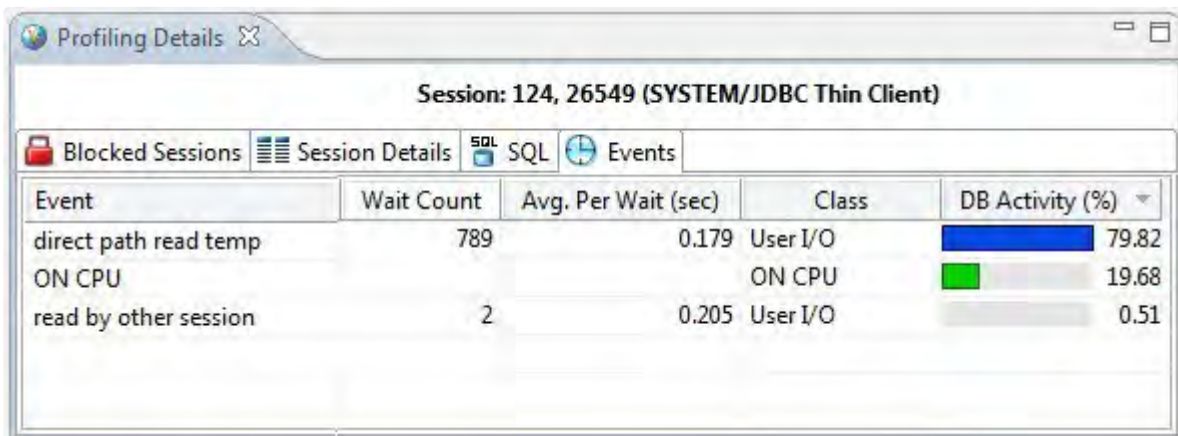
SQL statements are listed by the following parameters:

| Value | Notes |
|-----------|----------------------------|
| Statement | The name of the statement. |

| Value | Notes |
|--------------|---|
| Executions | The number of times the statement was executed during the session. |
| Activity (%) | A graphical representation of the distribution of execution and wait time for the statement or statement component. |
| SQL ID | The SQL ID value of the statement. |
| Child Number | The child number in the database. |

Events

The Events tab provides details about the events that the blocking session is associated with.



Events are listed by the following values:

| Value | Notes |
|-----------------|---|
| Event | The name of the event. |
| DB Activity (%) | A graphical representation of the distribution of execution and wait time for the statement or statement component. |

Viewing details on the Events tab


In the **Top Activities Section**, selecting a statement entry on the **Event** tab displays information in the **Profiling Details** view. The graph portion and details on the event category tabs on the new editor pertain only to the selected statement. Additionally, new tabs become available.

Selecting an event type entry on an event category tab opens a new profiling editor page. The graph portion and details on the **Events** tab and event category tabs on the new editor page pertain only to the selected wait event and to SQL statements that waited in that event.

- **SQL** tab: Shows the statements involved in the selected event. For more information, see [SQL](#).
- **Sessions** tab: Displays information about the sessions that the selected event is associated with. For more information, see [Sessions](#).
- **Procedures** tab: Displays information about the procedures that ran during the selected event. For more information, see [Procedures](#).

SQL

The SQL tab displays information about the SQL statements involved in the selected event.



| Value | Notes |
|-----------------|---|
| Statement | The name of the statement. |
| SQL ID | The ID value of the SQL statement. |
| Child Number | The child number in the database. |
| Parsing User ID | The ID of the user who parsed the statement. |
| Plan Hash Value | The execution value of the statement. |
| CPU | Cumulative CPU time for the process. (measured in "ticks", an arbitrary unit of time) |
| Physical IO | Cumulative disk reads and writes for the process. (total count) |
| Memory Usage | Number of pages in the procedure cache that are currently allocated to this process. A negative number indicates that the process is freeing memory allocated by another process. |
| Executions | The number of times the statement was executed. |
| Activity (%) | A graphical representation of the distribution of execution and wait time for the statement or statement component. |

Sessions

The Sessions tab displays the sessions and related information regarding those that were associated with the selected event.

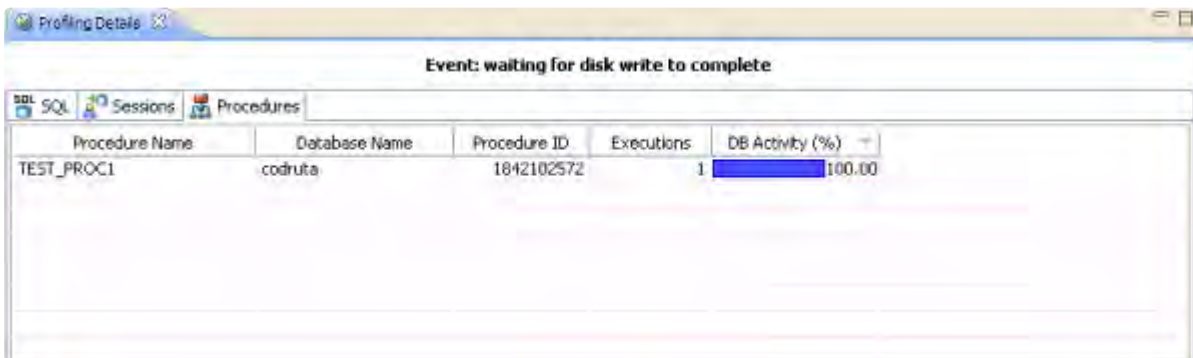
| User Name | Program | SID | Serial# | Activity (%) | Machine | Session Type |
|-----------|-------------------|-----|---------|--------------|-------------------------|--------------|
| SYSTEM | Executor.exe | 125 | 32563 | 31.42 | EMBARCADERO\ROWEBITA01 | USER |
| | ORACLE.EXE (CTWR) | 156 | 3 | 29.48 | TORLABORCL10G_1 | BACKGROUND |
| | ORACLE.EXE (DBW0) | 167 | 1 | 0.75 | TORLABORCL10G_1 | BACKGROUND |
| | ORACLE.EXE (MMON) | 161 | 1 | 0.45 | TORLABORCL10G_1 | BACKGROUND |
| IGNITE | ORACLE.EXE (J002) | 149 | 26903 | 0.10 | TORLABORCL10G_1 | USER |
| | ORACLE.EXE (SMON) | 164 | 1 | 0.10 | TORLABORCL10G_1 | BACKGROUND |
| | ORACLE.EXE (PSP0) | 169 | 1 | 0.10 | TORLABORCL10G_1 | BACKGROUND |
| IGNITE | ORACLE.EXE (J015) | 118 | 1 | 0.05 | TORLABORCL10G_1 | USER |
| SYS | ORACLE.EXE (J007) | 124 | 3890 | 0.05 | TORLABORCL10G_1 | USER |
| SYSTEM | Executor.exe | 145 | 9242 | 0.05 | EMBARCADERO\ROWSNOVAC01 | USER |
| IGNITE | ORACLE.EXE (J000) | 159 | 5 | 0.05 | TORLABORCL10G_1 | USER |

The following parameters are displayed on the Sessions tab:

| Value | Notes |
|---------------|---|
| User Name | The user name under which the session was run. |
| Program | The name of the executable under which the session was run. |
| SID | The SID value of the session. |
| Serial Number | The serial number of the machine from which the session executed. |
| Activity (%) | A graphical representation of the distribution of execution and wait time for the statement or statement component. |
| Machine | The machine name and network location of the machine from which the session executed. |
| Session Type | The type of session. |

Procedures

The Procedures tab displays the procedures and related information regarding those that were associated with the selected event.



The following parameters are displayed on the Procedures tab:

| Value | Notes |
|-----------------|--|
| Procedure Name | The name of the procedure that ran during the event. |
| Database Name | The name of the database where the procedure resides. |
| Procedure ID | The unique ID of the procedure. |
| Executions | The number of times the procedure ran during the event. |
| DB Activity (%) | A graphical representation of the distribution of execution and wait time for the procedure. |

Viewing details on the Procedures tab

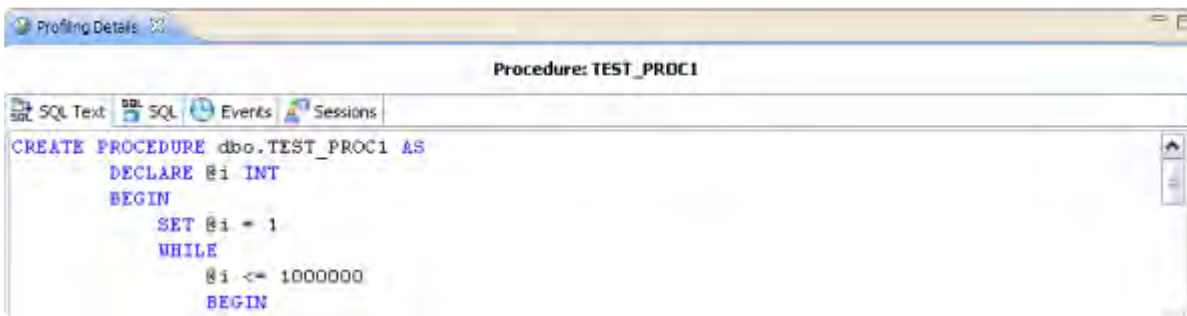
In the **Top Activities Section**, selecting a procedure entry on the **Procedure** tab displays information in the **Profiling Details** view. The graph portion and details on the procedure category tabs on the new editor pertain only to the selected procedure. Additionally, new tabs become available.

Selecting a procedure type entry on a procedure category tab opens a new profiling editor page. The graph portion and details on the **Procedure** tab and procedure category tabs on the new editor page pertain only to the selected procedure and to SQL statements that waited in that procedure.

- The **SQL Text** tab shows the SQL of the procedure. For more information, see [SQL Text](#).
- The **SQL** shows the statements involved in the procedure. For more information, see [SQL](#).
- The **Events** displays the time and parameter information about the selected procedure. For more information, see [Events](#).
- The **Sessions** displays information about the sessions that the selected procedure is associated with. For more information, see [Sessions](#).


SQL Text

The SQL Text tab displays the full code of the procedure.



SQL

The SQL tab displays information about the SQL statements involved in the selected procedure.



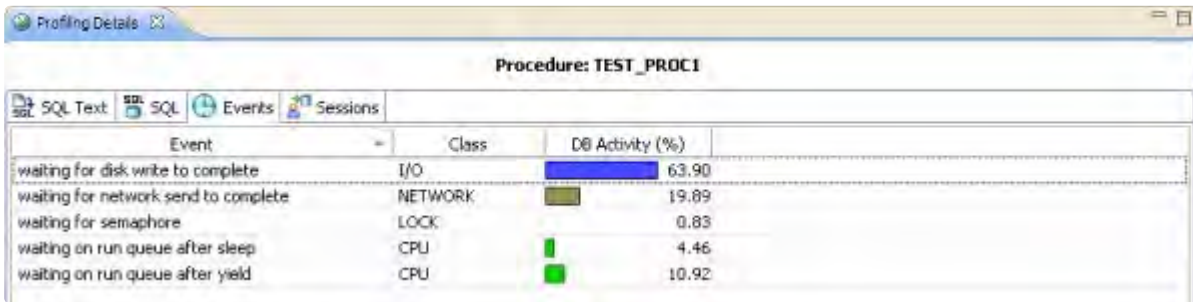
The SQL tab displays the following parameters about the statement:

| Value | Notes |
|-----------------|--|
| Statement | The name of the statement. |
| SQL ID | The ID value of the SQL statement. |
| Child Number | The child number in the database. |
| Parsing User ID | The ID of the user who parsed the statement. |
| Plan Hash Value | The execution value of the statement. |
| CPU | Cumulative CPU time for the process. (measured in "ticks", an arbitrary unit of time) |
| Physical IO | Cumulative disk reads and writes for the process. (total count) |
| Memory Usage | Number of pages in the procedure cache that are currently allocated to this process. A negative number indicates that the process is freeing memory allocated by another process. |
| Executions | The number of times the statement was executed. |

| Value | Notes |
|--------------|---|
| Activity (%) | A graphical representation of the distribution of execution and wait time for the statement or statement component. |

Events

The Events tab provides details about the events that the session is associated with.




Events are listed by the following values:

| Value | Notes |
|--------------|--|
| Event Name | The name of the event. |
| Class | The wait group the event in the selected procedure belongs to. |
| Activity (%) | A graphical representation of the distribution of execution and wait time for the event. |

Sessions

The Sessions tab displays the sessions and related information regarding those that were associated with the selected procedure.



The following parameters are displayed on the Sessions tab:

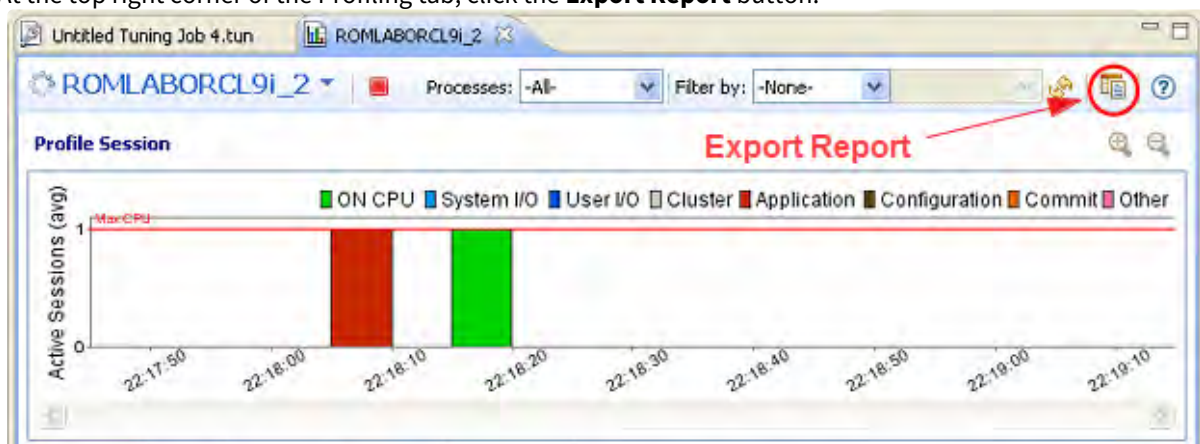
| Value | Notes |
|-----------|--|
| User Name | The user name under which the session was run. |

| Value | Notes |
|---------------|---|
| Program | The name of the executable under which the session was run. |
| SID | The SID value of the session. |
| Serial Number | The serial number of the machine from which the session executed. |
| Activity (%) | A graphical representation of the distribution of execution and wait time for the statement or statement component. |
| Machine | The machine name and network location of the machine from which the session executed. |
| Session Type | The type of session. |

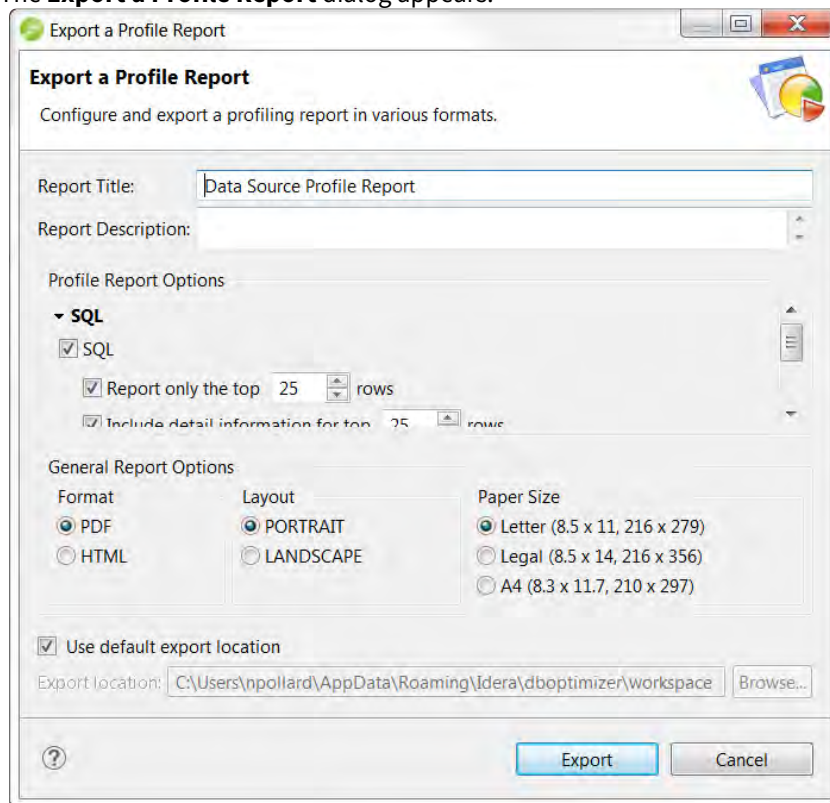
Creating profiling reports

After profiling a data source, you can create an HTML or PDF Report of the profiling session. You can choose the details to include in the report.

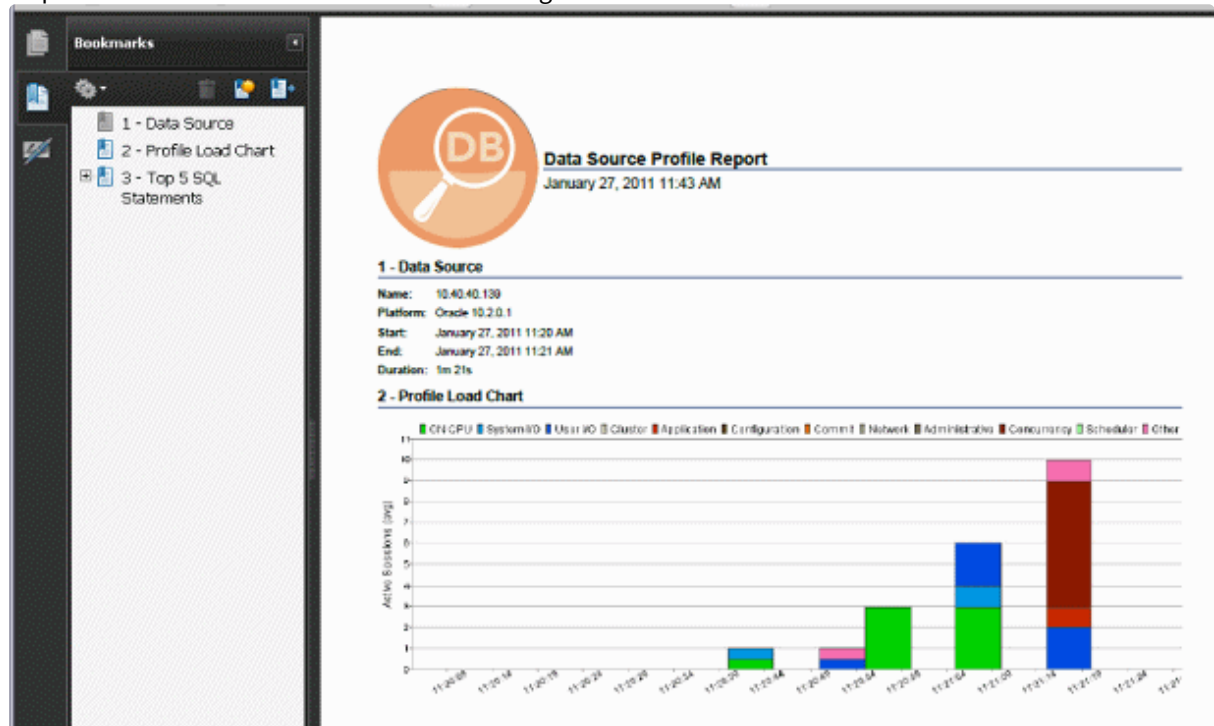
1. At the top right corner of the Profiling tab, click the **Export Report** button.



The **Export a Profile Report** dialog appears.



2. Enter a **Report Title** and **Description**.
 3. In the **Profile Report Options** area, click the triangles to expand the options.
 4. Select your options, enter the location for the report, and then click **Export**.
- A report in PDF format will resemble the following:



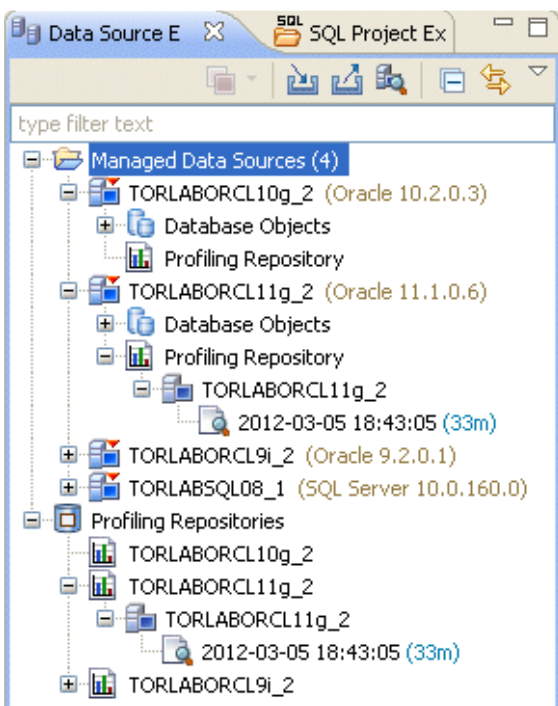
Saving profiling sessions

The profiling session is saved automatically or when you try to close it according to the choices made in the Profile Configuration dialog. For information on configuring profiling sessions, see [Building profiling configurations](#). Profiling sessions can be saved in the current workspace in an archive file with a **.oar** suffix or for Oracle users, into a Profiling Repository.

The .oar archive file is named with a default file name of:


- The name of the data source if the session was not initiated from a named launch configuration
- The name of the launch configuration if the session was initiated from a named launch configuration.

If you are using an Oracle data source and have configured DB Optimizer to automatically save profiling sessions in a Profiling Repository within an Oracle data source, then the profiling session is saved in the Profiling Repository under the name of the data source. Each profiling session for that data source is named using a date and time stamp. As you can see below, the duration of the profiling session is also saved with the session data.



For information on working with the Profiling Repository, see [Work with the Profiling Repository](#). The time period of the saved session is the amount of data on the chart. The maximum amount of data on the chart is determined when profiling is started (1-hour default). You can specify the amount of time to profile the data source in the Profile Configurations dialog and you can also stop the profiling at any time.

Saving the profile lets you open the archive at a later time for subsequent analysis by yourself or by other DB Optimizer users. Use standard DB Optimizer file techniques to save, open, or close SQL Profiling archives. If you open a profiling archive on a machine on which the associated data source is not registered, a **Data source not available** warning appears in the profiling editor header. Use the associated control to specify a data source already defined on the machine or to register a new data source.



Work with the Profiling Repository

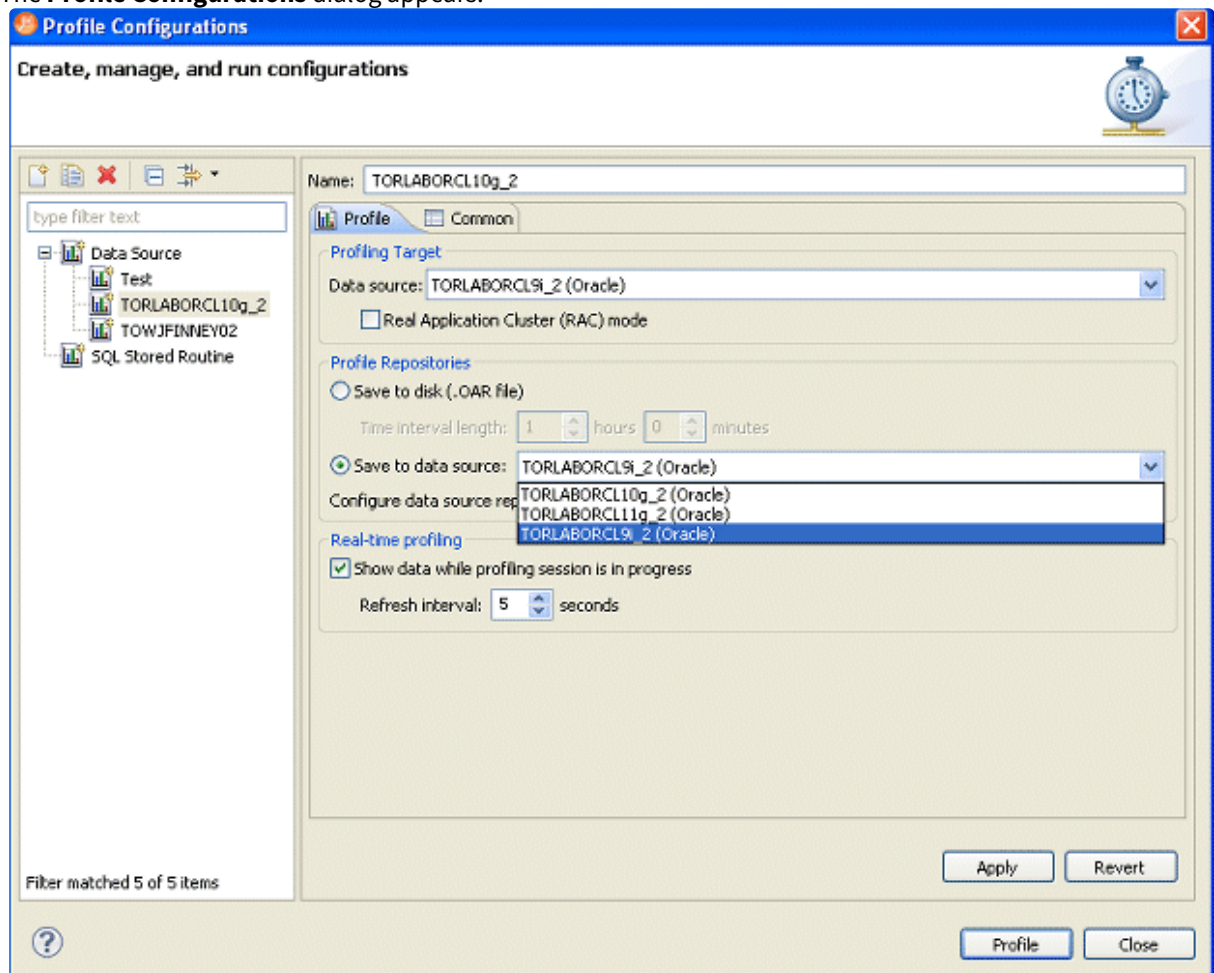
i Saving wait-time statistics to the Profiling Repository is not supported in Developer Editions of DB Optimizer.

The Profiling Repository is only available when profile session data is saved to an Oracle data source. You can specify a profiling repository on any supported version of Oracle data source. For information configuring profiling repositories, see [Specify Profile Repositories Preferences](#). You can specify which profiling to save a session to using the Profile Configurations dialog. For more information, see [Building profiling configurations](#). When the system is configured to automatically save profiling information to a Profiling Repository, DB Optimizer can profile 24 hours a day 7 days a week, thus providing much more statistical data for analysis. Also, since a Profiling Repository resides on a data source and not on the local disk, other DB Optimizer users can also view and analyze the profiles.

To start saving profile sessions to the Profiling Repository

1. From the Data Source Explorer right-click the data source you want to profile for and select **Profile As** from the menu, then choose **Profile Configurations**.

The **Profile Configurations** dialog appears.



2. In the **Name** field of the **Profile** tab, enter a name for the profile session for this data source.
3. In the Profiling Target Data source field, select a data source from the list.
4. In the **Profiling Repository** area of the **Profile** tab, click **Save to data source** and then from the list of registered profiling repositories, choose the Oracle data source to which you want to store the profiling session information.

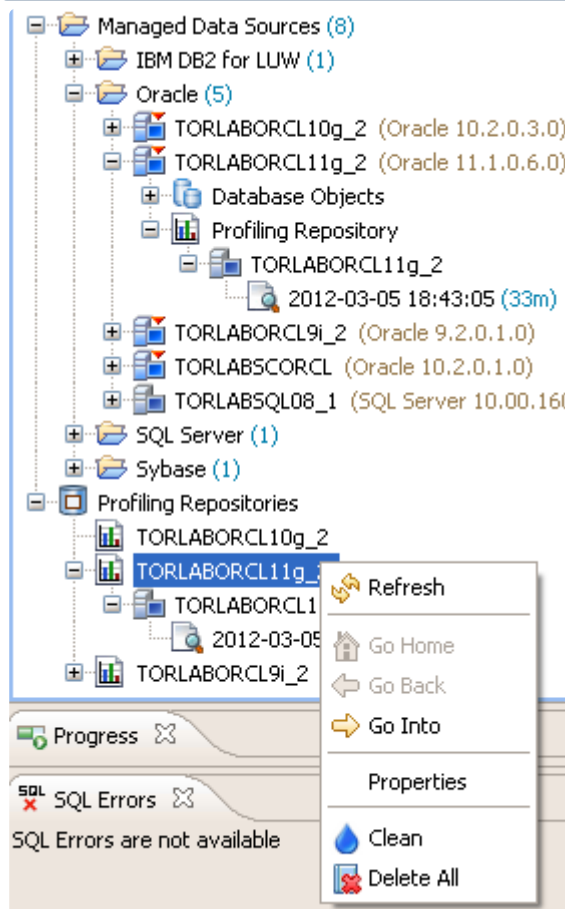
- Click **Apply** and then click **Profile** to start a profiling session immediately.

Any new profiling session that you start continues until you manually stop it. The profile session can be for as long as you like, days or weeks even. When the profiling session is stopped, a profile file is stored in the profile group for this data source. The name of the profile file is the date and time when the profile finished.

To delete profile sessions saved in the Profiling Repository

- In the **Data Source Explorer**, locate and then click the Profiling Repository.

i You access the profile repositories either in the Managed Data Sources node or in the Profiling Repositories.



If the data source is not already connected, DB Optimizer connects to the Profiling Repository data source.

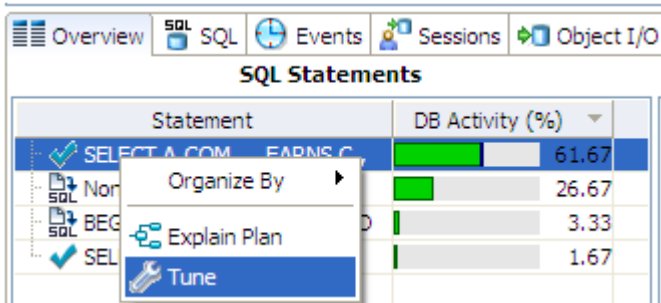
- To delete all sessions in the Profiling Repository, right-click **Profiling Repository** and then select **Delete All**. To delete a specific profiling session, expand the Profiling Repository and the data source containing the profiling session, and right-click the name of the profiling session, and then select **Delete**.

The profiling session data is deleted, however, some information about the data source is retained in order to expedite future profiling on this data source. If you are certain you will not want to retain this information, right-click the Profiling Repository and then select **Clean**.

i You can also delete the profiling sessions saved to a repository by unregistering the repository. For more information see [Specify Profile Repositories Preferences](#).

Import statements to tuning

The profiling feature lets you submit one or more SQL tab statements for tuning by the tuning feature. This lets you take advantage of tuning's hint-based and transformation-based suggestions, detailed execution statistics, and explain plan costing, in tuning a statement.



To open a tuning job on a statement appearing on the SQL tab of the profiling editor

Select one or more statements, right-click and select **Tune** from the context menu. Tuning opens on the selected statement.

For more information, see [Tuning SQL statements](#).

Other profiling commands

In addition to the default viewing options provided by the views, profiling also provides the following features and functionality:

- **Zooming In and Out.** For more information, see [Zooming In and Out](#).
- **Filtering Results.** For more information, see [Filtering Results](#).

Zooming in and out

To zoom in or out on the Load Graph


In the upper right-hand corner of the Load graph, click the Zoom In or Zoom Out icons, respectively.

i The Zoom In and Zoom Out commands are only available when a session has been stopped.

By default, the information contained on the Load Chart spans the entire length of the profiling session. You can select one or more bars of the graph to have the tabbed view populated with statistics for only the selected subset of the graph.

To display statistics for one or more bars on the graph, use the following method

Click-drag across one or more bars.



Configuring profiling

This section addresses the following topics:

- [Configuring DBMS properties and permissions](#)
- [Building profiling configurations](#)
- [Specify profile alerts preferences during configuration](#)
- [Specify profile repositories preferences during configuration](#)
- [Using SQL Load Editor/Tester](#)

Configuring DBMS properties and permissions


Profiling supports the following DBMS platforms:

- IBM DB/2 for Windows, UNIX, and Linux
- Microsoft SQL Server
- Oracle
- Sybase

The following describe how to set up a platform to utilize Optimizer on supported database platforms:

- [Configuring IBM DB/2 for Windows, UNIX, and Linux](#)
- [Configuring Microsoft SQL Server](#)
- [Configuring Oracle](#)
- [Configuring Sybase](#)

Configuring IBM DB/2 for Windows, Unix, and Linux

 The connected profiling user should be a member of the DB2 SYSMON group.

By default, DB2 Monitor flags are set to OFF. As a result, when attempting to launch a Profile job on a DB2 data source, users may encounter the following message: "One or more errors have occurred that prevent session profiling against this data source." Examine the details below and consult your data source administrator and/or the data source documentation to resolve the problem(s)."

You can resolve this error using one of two methods:

- Enabling DB2 Monitor Flags via IDERA DBArtisan
- Command Line Option

To resolve the error through DBArtisan

1. Ensure the following DB2 Monitor Flags are turned on in DB2:
 - dft_mon_uow
 - dft_mon_stmt
 - dft_mon_lock
 - dft_mon_bufpool
 - dft_mon_sort
 - dft_mon_table
 - dft_mon_timestamp

Yes



- Restart the DB2 data source to enable the changes, then launch SQL Query Tuner and begin profiling.

To resolve the error through the command line

This solution must be performed from DB2 CLP, on the DB2 server. If you attempt to perform these tasks through a client, an error message results.

1. Navigate to **Start > Programs > IBM DB2/COMMAND LINE TOOLS > COMMAND LINE PROCESSOR**.
2. Turn the monitor switches on using the following commands:

db2

update

dbm cfg

using

dft_mon

_lock

on

dft_mon

_bufpoo

l on

dft_mon

_sort

on

dft_mon

_stmt

on

dft_mon

_table

on

timesta

mp on

dft_mon

`_uow on`

`db2stop`

db2star

t

3. Ensure that the switches are turned on by connecting to the server with the following command:

Db2

connect

to

database

e

username

e

password

d

password



The following screenshot provides an example of the input and output from the server:

```

db2 => connect to gim user db2admin
Enter current password for db2admin
db2 => get monitor switches

Monitor Recording Switches

Switch list for db partition number 0
Buffer Pool Activity Information (BUFFERPOOL) = ON 03/05/2009 19:14:06.61
Lock Information (LOCK) = ON 03/05/2009 19:14:06.61
Sorting Information (SORT) = ON 03/05/2009 19:14:06.61
SQL Statement Information (STATEMENT) = ON 03/05/2009 19:14:06.61
Table Activity Information (TABLE) = ON 03/05/2009 19:14:06.61
Take Timestamp Information (TIMESTAMP) = ON 03/05/2009 18:50:44.00
Unit of Work Information (UOW) = ON 03/05/2009 19:14:06.61

```

You can set view and set Monitor Flags via DBArtisan. Ensure that the New Value field for each variable is set to , as shown below.

Configuring Microsoft SQL Server

Perform the following tasks to ensure that SQL Server is compatible with Optimizer:

- The current user must meet one of the following requirements:
 - Be a member of sysadmin, or have the VIEW SERVER STATE permission enabled.
 - Be a member of sysadmin, or have the SELECT permission enabled.

Configuring Oracle

Oracle users need access to V\$ views. In order to configure Oracle to provide users with these privileges:

- If you are setting up Oracle 10 or later, ensure you are logged in as sys or system with the sysdba role, or the SELECT_CATALOG_ROLE has been granted to *user_name*.
- If you are setting up an earlier version of Oracle, ensure you are logged in as sys or system with the sysdba role.

Configuring Sybase

Perform the following tasks to ensure that Sybase is compatible with DB Optimizer:


- Ensure the following system configuration properties are activated:
- Enabling Monitoring (sp_configure "enable monitoring", 1)
- Wait Event Timing (sp_configure "wait event timing", 1)
- SQL Batch Capture (sp_configure "SQL batch capture", 1)
- Max SQL Text Monitored (sp_configure "max SQL text monitored", 2048) 2048 is the suggested value, must be greater than 0. This defines the maximum size a SQL statement that can be captured.

The following options are specific to Sybase 15.0.2 and higher.

- SQL Text Pipe Active (sp_configure "sqltext pipe active", 1)
- SQL Text Pipe Max Messages (sp_configure "sql text pipe max messages", 512) 512 is the suggested value, must be greater than 0. This defines the maximum number of SQL statements that Sybase tracks at any given time. You may want to increase this number if you observe a lot of UNKNOWN statements in the Profile Editor.
- Statement Statistics Active (sp_configure "statement statistics active", 1)
- Per Object Statistics Active (sp_configure "per object statistics active", 1)

Additionally, perform the following tasks, as necessary:


- If a user does not have mon_role enabled, the user will not be able to access Adaptive Server's monitoring tables.
- If the monProcess table is missing, the user will not be able to view currently connected sessions.
- If the sysprocesses table is missing, the user will not be able to view information about Adaptive Server processes.
- If the monWaitEventInfo table is missing, the user will not be able to view information about wait events.
- If the monProcessSQLText table is missing, the user will not be able to view currently executing SQL statements.

 These packages should only be installed by the DBA.

Profiling enables you to create a set of launch configurations to store the basic properties for each profiling session that you run on a regular basis. A launch configuration enables you to start profiling sessions from a single menu command, rather than re-define configuration parameters each time you want to run one.

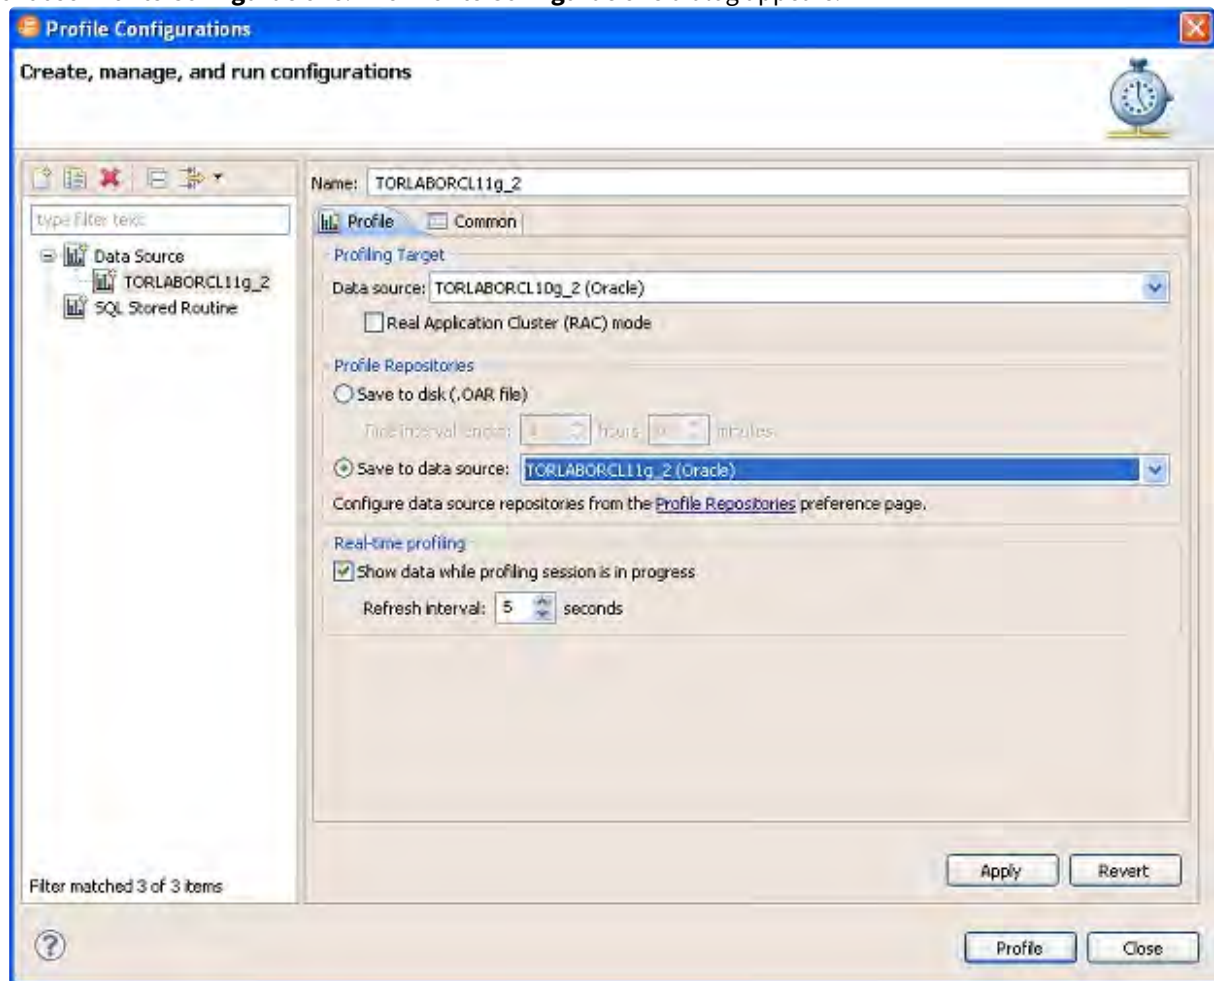
Building profiling configurations

Profiling enables you to store parameters related to specific profiling sessions, in a profile configuration for stored routines. Multiple configurations can be created for each data source in your enterprise and saved with unique names that identify them in the application.

 On all supported platforms, support for stored routines includes functions and procedures. On Oracle, stored routine support also includes package functions and package procedures.

To create a profile configuration

1. Right-click the data source you want to build a configuration for and select **Profile As** from the menu, then choose **Profile Configurations**. The **Profile Configurations** dialog appears.



2. Select the name of the data source and modify the parameters on the **Profile** tab, as needed.
3. In the **Name** field, provide a name for the launch configuration. Select a name that makes the launch configuration unique and easily identified once it is saved in the application.
4. In the **Profiling Target** area, click **Real Application Cluster (RAC) mode** if the target database is an Oracle RAC. This enables you to profile the entire cluster in one profiling session. (In general, profiling a RAC entails querying the GV\$ views.)

i When profiling a RAC, you can also filter the profiling details to show only the details for a selected instance. At the top of the Profiling view, click the Instances list and choose the instance you want to examine.

5. In the **Profiling Repository** area, choose to save the profile session to disk or if you are profiling an Oracle data source then you can choose to save the profile session to a data source registered as a Profiling Repository.

If you choose to save the session to disk, in the **Time interval length** area specify the length of the profiling session. When you try to close the Profile Session, you will be prompted to save the file and can then name the file as desired.

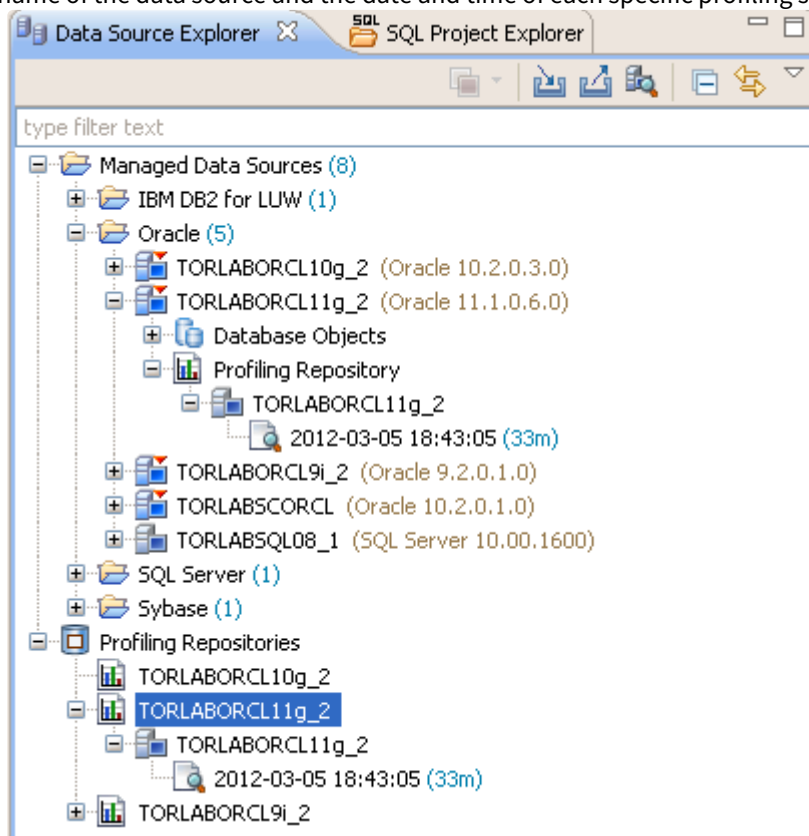
A .oar file saved to disk opens very quickly from **File > Open** dialog and has a limit of just under 1000 hours of profiling data. Profiling to a data source directly allows the system to capture more data for a longer period

- of time, until you decide to stop the profiling session. The profiling session is automatically saved to the Profiling Repository where other SQL Query Tuner users can also view the session for their own analysis.
- Click **Apply**. The launch configuration is stored in the application.
 - Once a launch configuration is defined, you can execute it in profiling. For more information, see [Running a profiling session](#).

i The parameters provided when you select the data source name in the left pane control session parameters for the specified data source. To set these controls, see [Configuring DBMS properties and permissions](#).

The following describes fields and options of the Profile tab that require further explanation.

- **Name** indicates the name of the profile configuration.
- **Data source** indicates the name of the data source to which the profile applies.
- **Save to disk/Save to data source** gives you the option to save your profiling session to a .oar file which you can access from within SQL Query Tuner or you can save the profiling session to a data source of your choice if you are using an Oracle data source. SQL Query Tuner will create a Profiling Repository, similar to the following on the selected Oracle data source. The structure of the Profiling Repository is created from the name of the data source and the date and time of each specific profiling session.



Saving your profiling sessions to a live data source enables you to better organize your profile session data for later review.

- **Time Interval Length** indicates how many hours of the session to save to disk. Since the profile session continues until you manually stop it the session length may exceed the time interval length. For example, the time interval length is set to four hours but the profiling session continues for 10 hours. In this case only the last four hours of data is retained. This parameter also indicates the total width of the time load graph.

The longer a profile is, the larger the saved file will be. For heavily loaded databases, the time interval length value should not exceed eight hours.

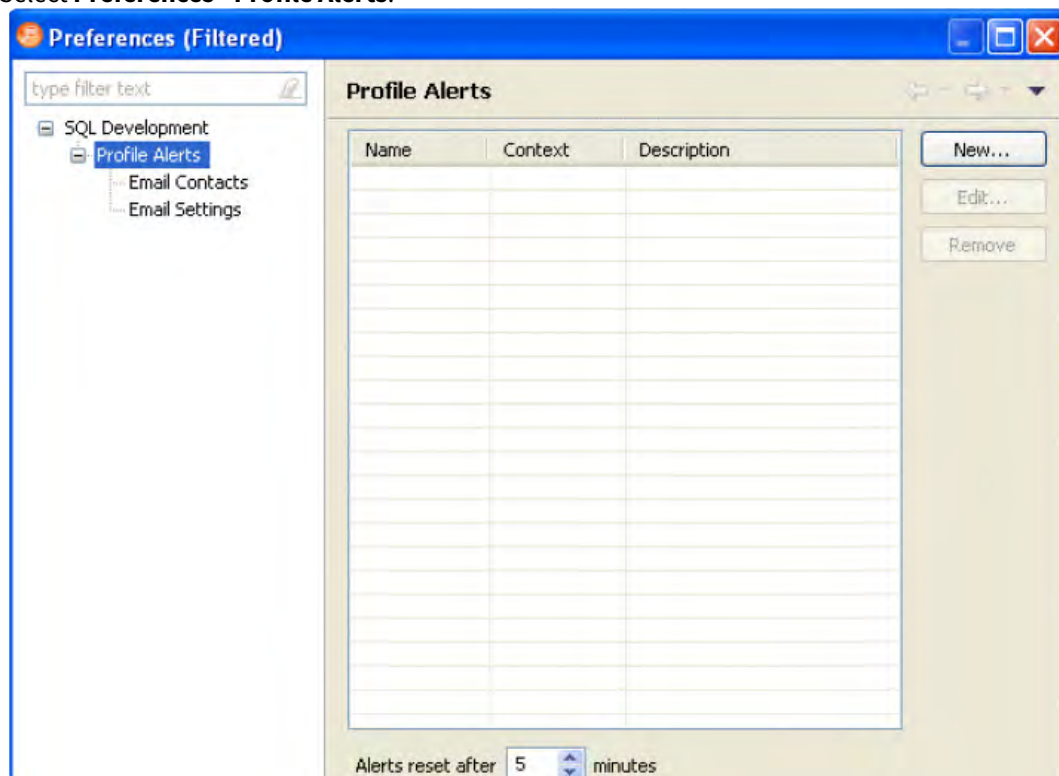
- The **Show Data While Profile Session is in Progress** check box enables "real time" profiling, which refreshes the data of the session as profiling runs. The **Refresh Interval** specifies how often in seconds profiling updates this data.

i Profiling can run sessions based on ad hoc parameters you designate before executing the profiling process. However, by building profile launch profiles, it is a much more efficient method of managing standard, frequent, or common profiling sessions.

Specify profile alerts preferences during configuration

You can configure DB Optimizer to send you an alert via email if during a profiling session it detects that database activity has met or exceeded the threshold that you specified.

1. Select **Preferences > Profile Alerts**.



2. To create a profile alert, click **New**.

Create Profiling Alert

Alert name must be specified.

Name:

Data source: TORLABSQL00_1 (SQL Server 8.0.2039.0)

Fire alert when

Active sessions

Sessions in wait:

Sessions in wait class:

All active sessions not on CPU

meet or exceed

1

1

Send email when alert fires

Display system notification when alert fires

OK Cancel

To edit an existing profile alert, select the alert you want to modify and then click **Edit**. The **Create Profiling Alert** or **Edit Profiling Alert** dialog appears.

3. Complete the create/edit profiling alert dialog as required and then click **OK**.

4. In the **Preferences** tree, click **Email Contacts**.

If you select **Display system notification when alert fires**, you will receive an alert notification in your Windows system tray when the alert fires.

If you select **Send email when alert fires**, you must specify email contacts and email server settings.

5. To create a new email contact, click **New**.

To edit an existing email contact, select the contact you want to modify and then click **Edit**.

Create Email Contact

Specify the details of the contact to add.

Name:

Email address:

Data Sources

- TORLABSQL00_1 (SQL Server 8.0.2039.0)
- TORLABSQL00_1_#2 (SQL Server 8.0.2039.0)
- sfvpclb01.embarcadero.com (Oracle 10.2.0.1)
- TORLABORCL8i_3 (Oracle)
- TORLABORCL8i_2 (Oracle 8.1.7.4)
- ROMLABORCL8i_1 (Oracle 8.1.7.4)
- ROMLABORCL9i_2 (Oracle 9.2.0.1)
- TORLABORCL9i_1 (Oracle)
- DSQUERY (SQL Server)
- HOBO (SQL Server)
- local (SQL Server)

Select All Deselect All

OK Cancel

6. From the list of **Data Sources**, select the data sources for which this contact should receive an email notification that an alert has fired, and then click **OK** to save this contact.

7. To configure your email settings, click **SQL Development > Profile Alerts > Email Settings**.

The screenshot shows the 'Email Settings' dialog box. On the left, a tree view shows the navigation path: SQL Development > Profile Alerts > Email Settings. The main panel is titled 'Email Settings' and contains an 'Email Configuration' section with the following fields:

- SMTP Host: [Empty text box]
- SMTP Port: [25]
- Sender Address: [Empty text box]
- Authorization Required
- Username: [Empty text box]
- Password: [Empty text box]

At the bottom right of the dialog are three buttons: 'Apply', 'OK', and 'Cancel'. A help icon (?) is located at the bottom left of the dialog.

8. Complete the Email Settings as required and click **Apply**.

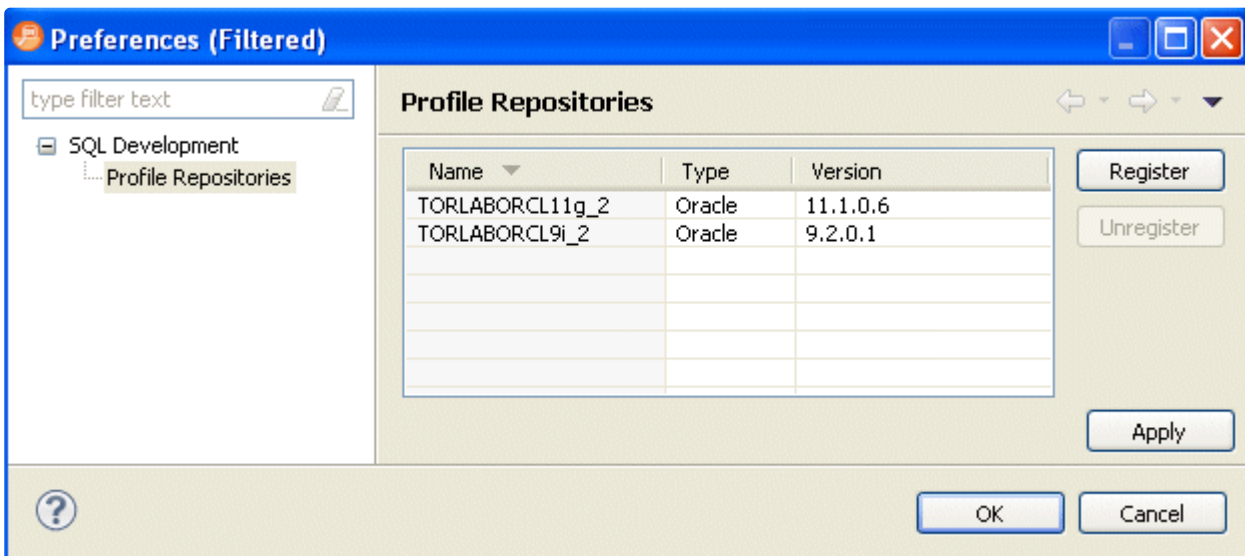
- **Sender Address:** This is an email user configured on your email server. SQL Query Tuner uses this address to send alerts to the email contacts defined.

Specify profile repositories preferences during configuration

Use the Profile Repositories preferences to register and unregister supported version of Oracle data sources as profiling repositories. A data source must be added to the list of DB Optimizer Managed Data Sources before it can be used as a repository.

To access the Profile Repositories Preferences

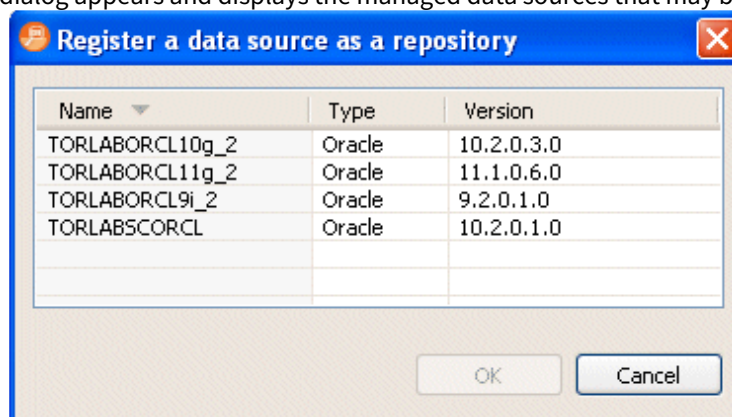
Select **Preferences > Profile Repositories**.



Register a profile repository

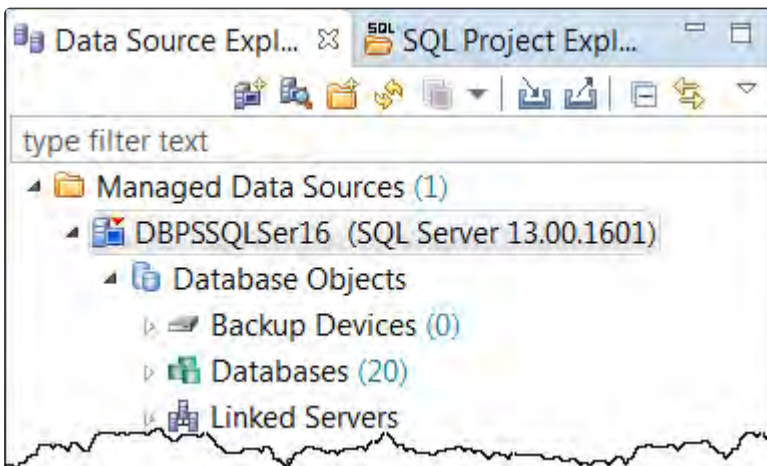
To register a new profile repository

1. In the Profile Repositories Preferences dialog, click **Register**. The **Register a data source as a repository** dialog appears and displays the managed data sources that may be used.

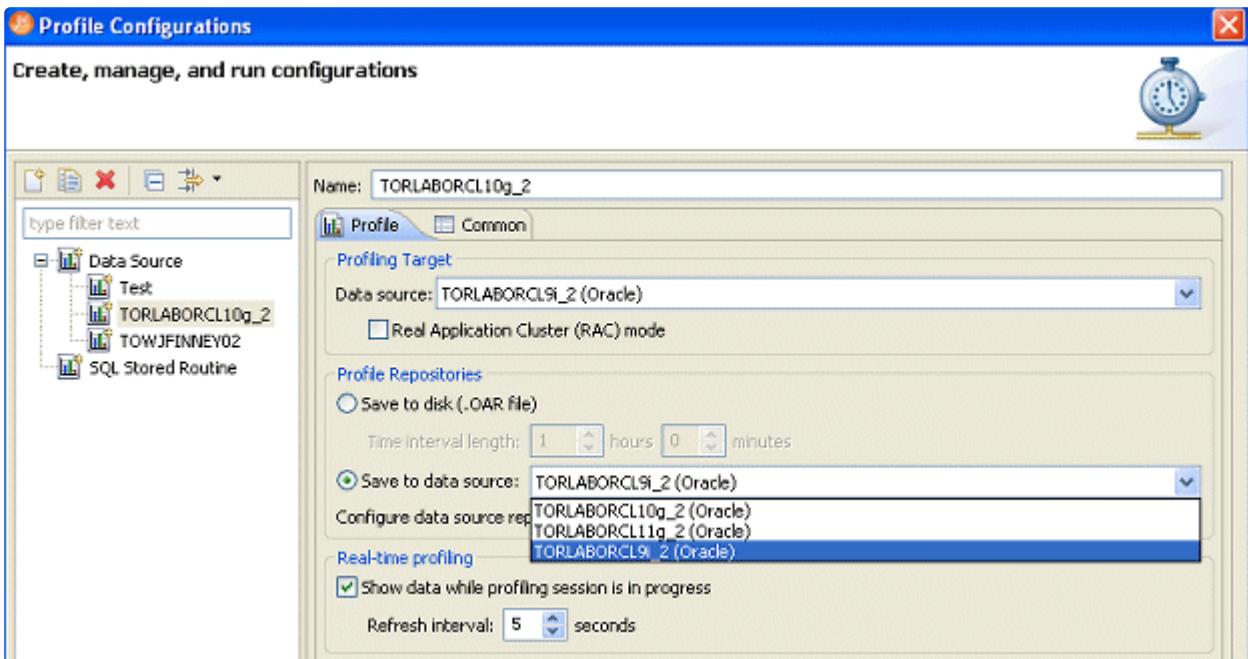


2. Select the data source you want to register as a profile repository, and then click **OK**.
3. Click **OK** again.

A data source registered as a profile repository appears in the **Profiling Repositories** node of the **Data Source Explorer** and also has a **Profile Repository** child node in under its entry in the **Managed Data Sources** folder. Profiling sessions are saved in the repository under a name that is comprised of the date, time, and duration of the profiling session.




When you next choose to profile a data source, using the **Profile As > Profiling Configurations** option, you can specify the profile repository in which you want to save the profiling session data. You can also access the Profile Repository Preferences from this dialog.



Unregister a profiling repository


To unregister a profiling repository

In the Profile Repositories Preferences dialog, select a profile repository from the list, and then click **Unregister**. Unregistering a data source prompts you to optionally delete all sessions on the repository.



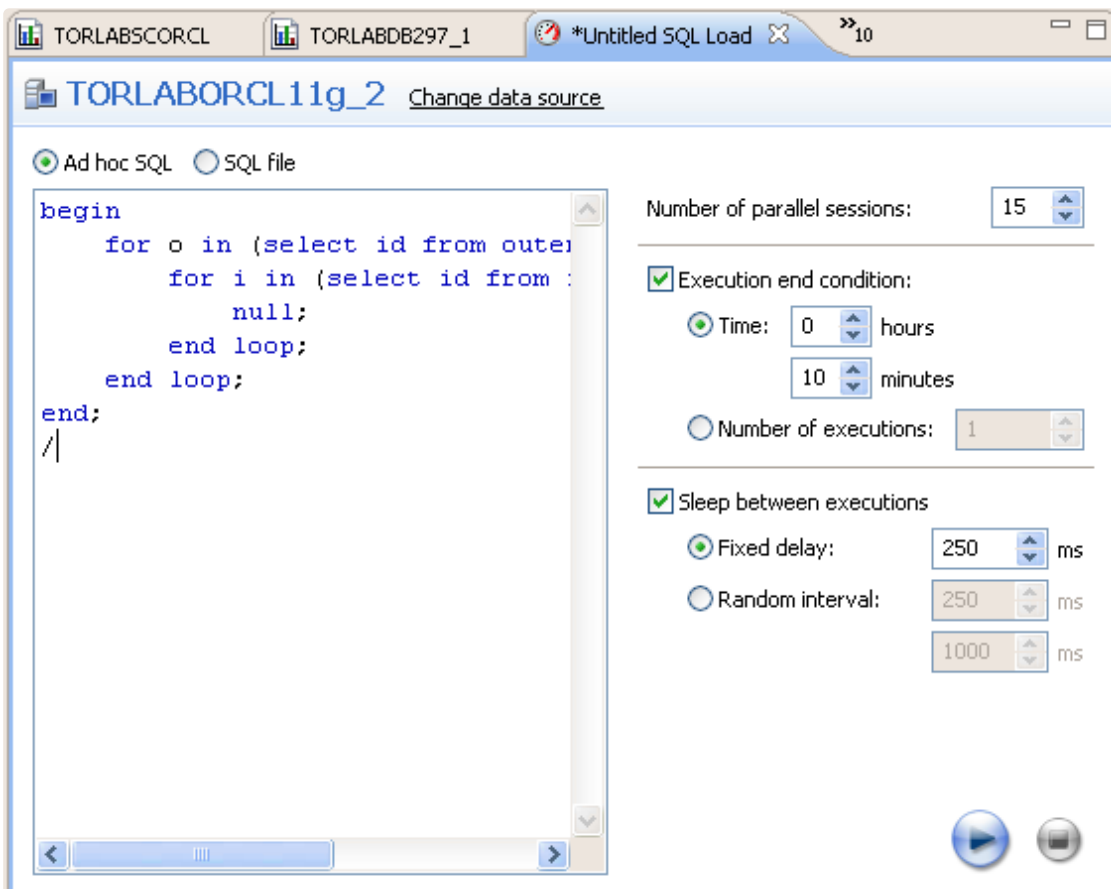
If you choose not to delete the profiling sessions, they will still be available from either the Profiling Repository node in the data source entry in Managed Data Sources or from the Profiling Repositories node. You will not be able to save profiling sessions to the unregistered repository until you register it.

Using SQL Load Editor/Tester



Run the load editor by using either **File > New > SQL Load** or with the **New SQL Load** icon shown above in the red square.

The icon depicts an RPM meter on a car with a red line. The idea behind the icon is that we can run a load on a database and stress the database with the load similar to red lining.



The load editor page has space on the left to show the SQL to be run. The SQL can be typed in or pasted in or read from a file if the **SQL File** option at the top right-hand side of the window is selected.

On the right are options on how to run the SQL:

- Number of parallel sessions
- Length of test
- Number of executions
- Sleep between executions
 - No sleep
 - Fixed sleep
 - Random sleep between a max and min

Methodology


1. Write SQL with Editor
2. Set up Load with Load Editor
3. Kick off profiling the database
4. Run the load in the Load Editor
5. Verify the database load profile to see if there are any major issues

The SQL Load Editor/Tester enables you to configure and execute SQL code against a data source.

This feature enables you to specify a data source against which the code will be executed, and then provides options that enable you to choose a period of time that you want the script to execute for, and at what intervals the execution "loop" occurs.

On execution, SQL Load Editor/Tester runs in the background. It can therefore be run in conjunction with a profiling session in order to analyze the effects of the executing load against the specified data source. Once you run a SQL script via Load Editor, you can start the SQL Profiling function and analyze the results of the load.

The SQL Load Editor/Tester is accessed via the **New SQL Load** icon on the Toolbar:



When you open SQL Load Editor/Tester, click **Select Data Source** to specify the data source against which you want the SQL script to run.

Choose **Ad hoc SQL** and manually type (or copy/paste) the SQL code into the window provided, or select **SQL file** and navigate to the SQL file you want to run. The window populates with the code from the selected file.

The following configuration parameters are set with SQL Load Tester/Editor prior to executing the SQL script:

| Configuration Parameter | Description |
|-----------------------------|--|
| Number of Parallel Sessions | Specifies the number of jobs that the execution script will operate on. |
| Execution End Condition | Specifies if the script execution process runs for a set amount of time or script executions. Choose Time if you want the script to execute over a specific period of time, or Number of Executions if you want the script to execute a specific number of times. |

| Configuration Parameter | Description |
|--------------------------|---|
| Sleep Between Executions | Specifies if Load Editor will wait before running the execution script again. Select the check box and choose Fixed Delay or Random Interval , depending on whether you want the script to execute at a specific time, or at random intervals within a specified range of time. |

To run Load Editor

1. Access SQL Load Editor/Tester by selecting the icon on the Toolbar. The **SQL Load Editor/ Tester** opens.
2. Click **Change Data Source** and choose a data source you want to run the SQL code against.
3. Choose **Ad hoc SQL** or **SQL file**, and then copy/paste or manually type the code you want to execute in the window provided, or navigate to the location of the file, respectively.
4. In the right-hand panel, choose the execution configuration parameters to specify how you want SQL Load Editor/Tester to handle the script.
5. Click the **Execute** icon in the lower right-hand corner of the screen. The script starts to execute against the specified data source, using the configuration parameters you selected.
6. If you are profiling a data source, start and run a new profiling session on the data source you specified in Load Editor. The session will reflect how your SQL script executes against the specified data source.

Using tuning

This section provides information on tuning, its functionality, and is structured so a user can follow the information provided to fully tune their enterprise in terms of more efficient query paths at the SQL statement level of individual data sources.

Tuner has three parts:

- Query rewrites and quick fixes
- Alternative execution plans generated via optimizer directives
- Analysis of Query showing:
 - Indexes used, not used, missing (suggested to create)
 - Graphic display of query

The SQL tuner takes a query and adds database optimizer directives to change the execution path of the query. A list of all the unique execution paths is generated with all duplicates eliminated from the list. The final list of alternative paths can be executed. Any path that takes more than 150% of the base case will be canceled because we are only interested on paths that could be faster than the base case so no need to waste time and resources continuing to run cases that are slower than the original. After the cases have been executed they can be sorted in order of elapsed time. If a better path is found then those optimizer directives can be included in the original query to achieve optimal response time.

You can save the entire content of a tuning job for later analysis or for sharing with other users. This section contains the following topics:

- [Understanding the tuner interface](#)
- [Tuning SQL statements](#)
- [Using platform-specific features](#)
- [Additional tuning commands](#)
- [Configuring tuning](#)
- [Examples of transformations and SQL query rewrites](#)
- [DBMS hints](#)

Understanding the tuner interface

In the application interface, tuning is composed of two tabs:

- [Overview](#)
- [Analysis](#)

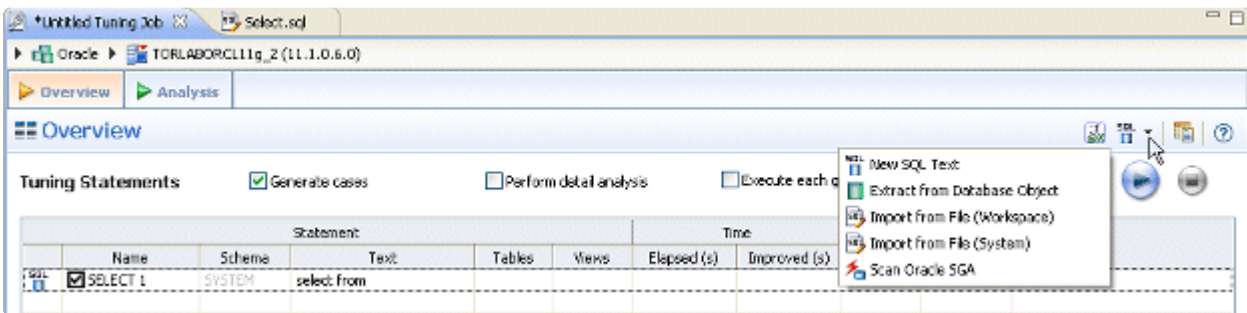


When using tuning on Oracle sources, several additional tabs appear on the Analysis and Outlines tabs. For more information on utilizing these extra features, see [Using platform-specific features](#).

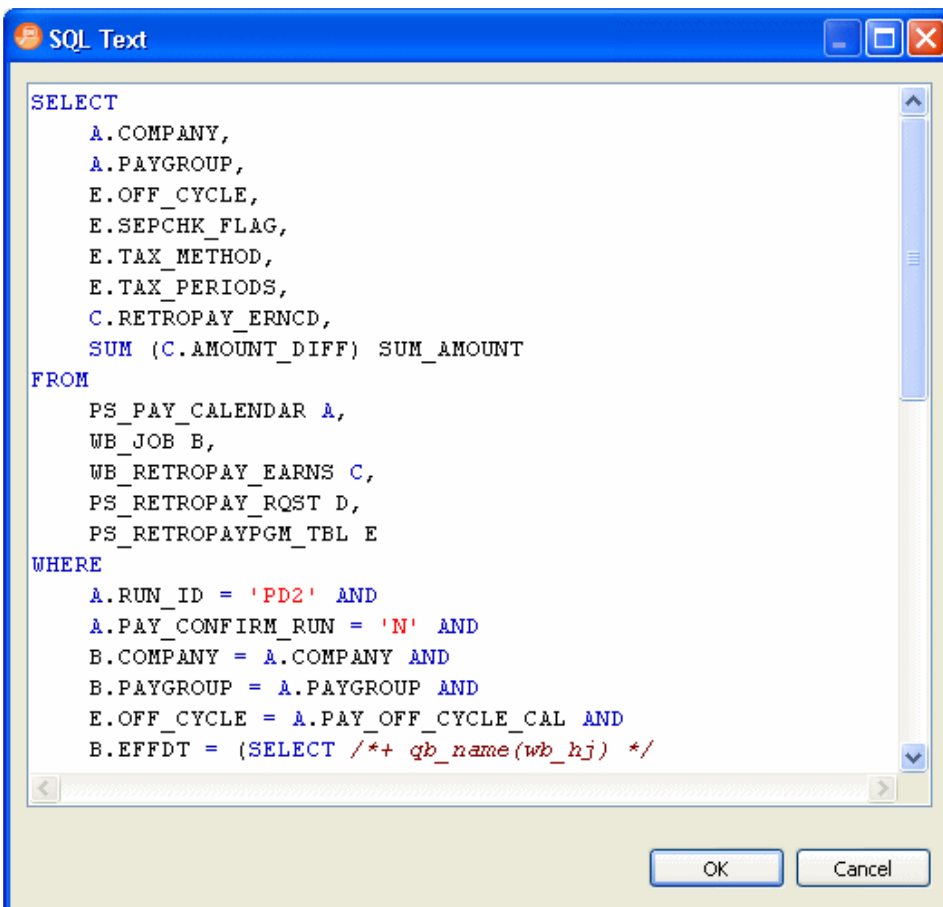
Understanding the Overview tab

Inputting SQL to tune

Click the SQL button on the **Overview** tab to specify the source of SQL statements you want to tune.



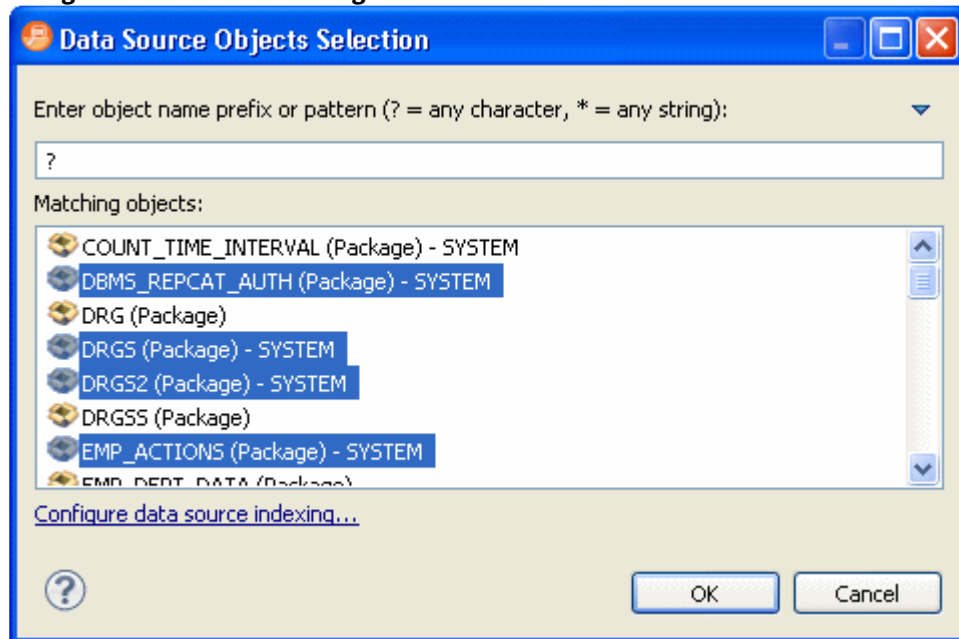
- **New SQL Text:** From the SQL button menu, select **New SQL Text**, and then copy/paste SQL statements to the **SQL Text** dialog or write queries by hand. Click **OK**.



i You can also input SQL by clicking anywhere in the Tuning Statements area and pressing Ctrl-V.

Once you have input the SQL and click **OK**, you can later edit the text by right-clicking an entry in the Tuning Statements area and selecting **Edit**.

- **Extract from Database Objects:** Search for and then select (Ctrl-click) data base objects containing SQL that you want to tune from the selected data source. SQL Query Tuner will search through the database to find objects matching your input and presents matches for you to choose. In order for this option to work, you must enable Data Source Indexing in the properties for the database. If the data source has not already been indexed you will receive a message indexing that no indexing information is available. You can configure the database Properties dialog from the **Data Source Objects Selection** dialog by clicking

Configure data source indexing.

For information on setting data source indexing properties, see [Specify Data Source Indexing Preferences](#).

- **Import from File (Workspace) and Import from File (System):** Browse the workspace or file system and select an SQL file from which to extract statements to tune.

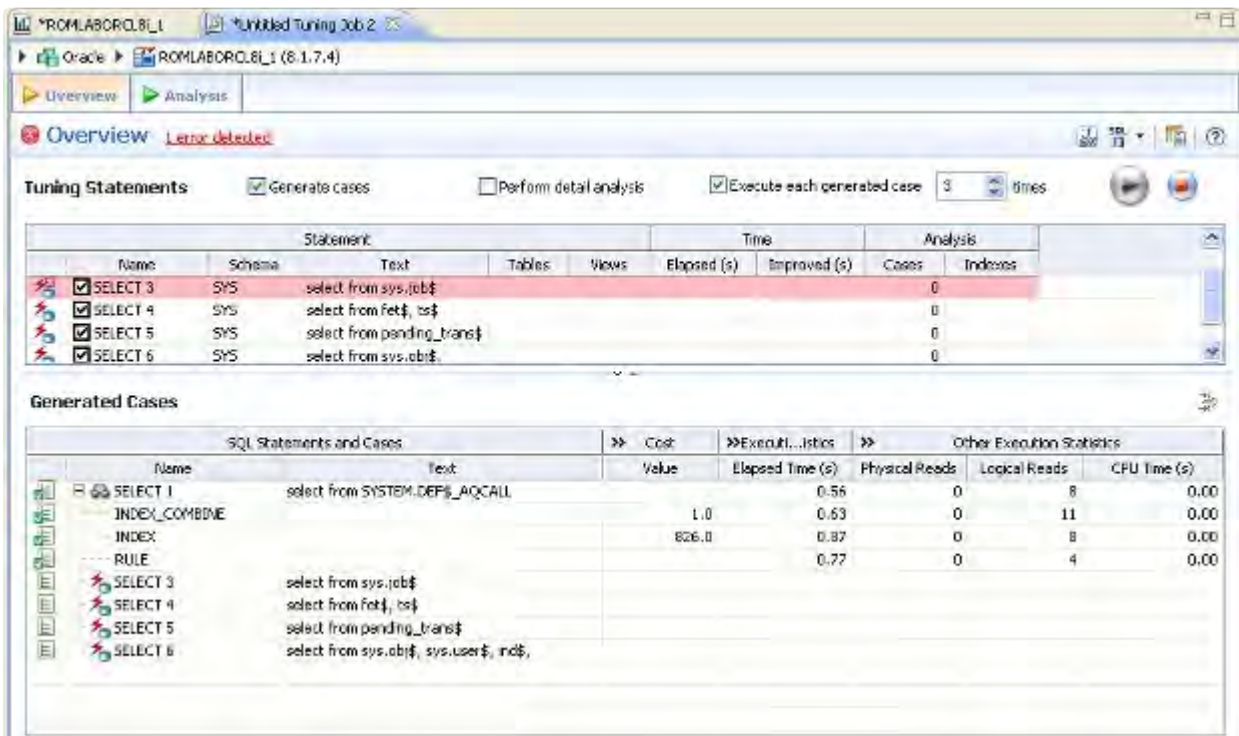
- **Scan Oracle SGA:** For the Oracle platform only, you can also scan the System Global Area (SGA) for statements to tune. Bind variables are extracted automatically.

- You can also drag and drop Materialized Views, Procedures, and Views from the Data Source Explorer to the Tuning Statements grid and they will be added to the list of statements to tune.

Running a tuning job

Once you click the Run Job icon on the top right-hand side of the Overview tab, the Overview tab provides the list of statements that were analyzed by the Tuner, as well as the cases suggested by the execution process to improve them. Additional information may include statement Name, Text, Source, Cost, and Elapsed Time values, depending on the platform.

Only the Elapsed Time statistic appears on all supported platforms. On Oracle and DB2 LUW platforms, Execution Statistics and Other Execution Statistics columns will appear. When determining the best possible path using the Overview tab, it is best to use the Elapsed Time value as the guideline. The faster the path, the more optimized the query will become.



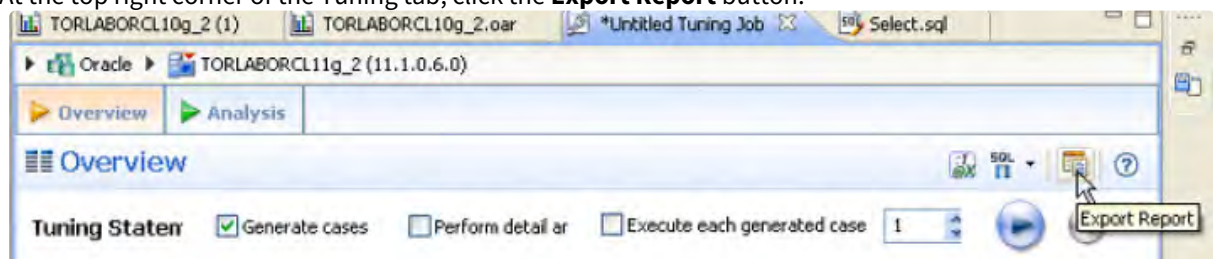
There are three tuning options to choose from before clicking Run Job:

- To analyze the SQL statement, click **Generate cases**.
- To perform the analysis that populates the Analysis tab now, click **Perform detail analysis**. Otherwise, the analysis tab is populated when you click the **Analysis** tab.
- To have the system generate execution statistics, click **Execute each generate case** and then select the number of time the system should execute each generated case. Multiple executions can verify that the case results are not skewed by caching. For example, the first time a query is run, data might be read off of disk, which is slow, and the second time the data might be in cache and run faster. Thus, one case might seem faster than another but it could be just benefiting from the effects of caching. Generally, you only need to execute the cases once, but it may be beneficial to execute the cases multiple times to see if the response times and statistics stay the same.

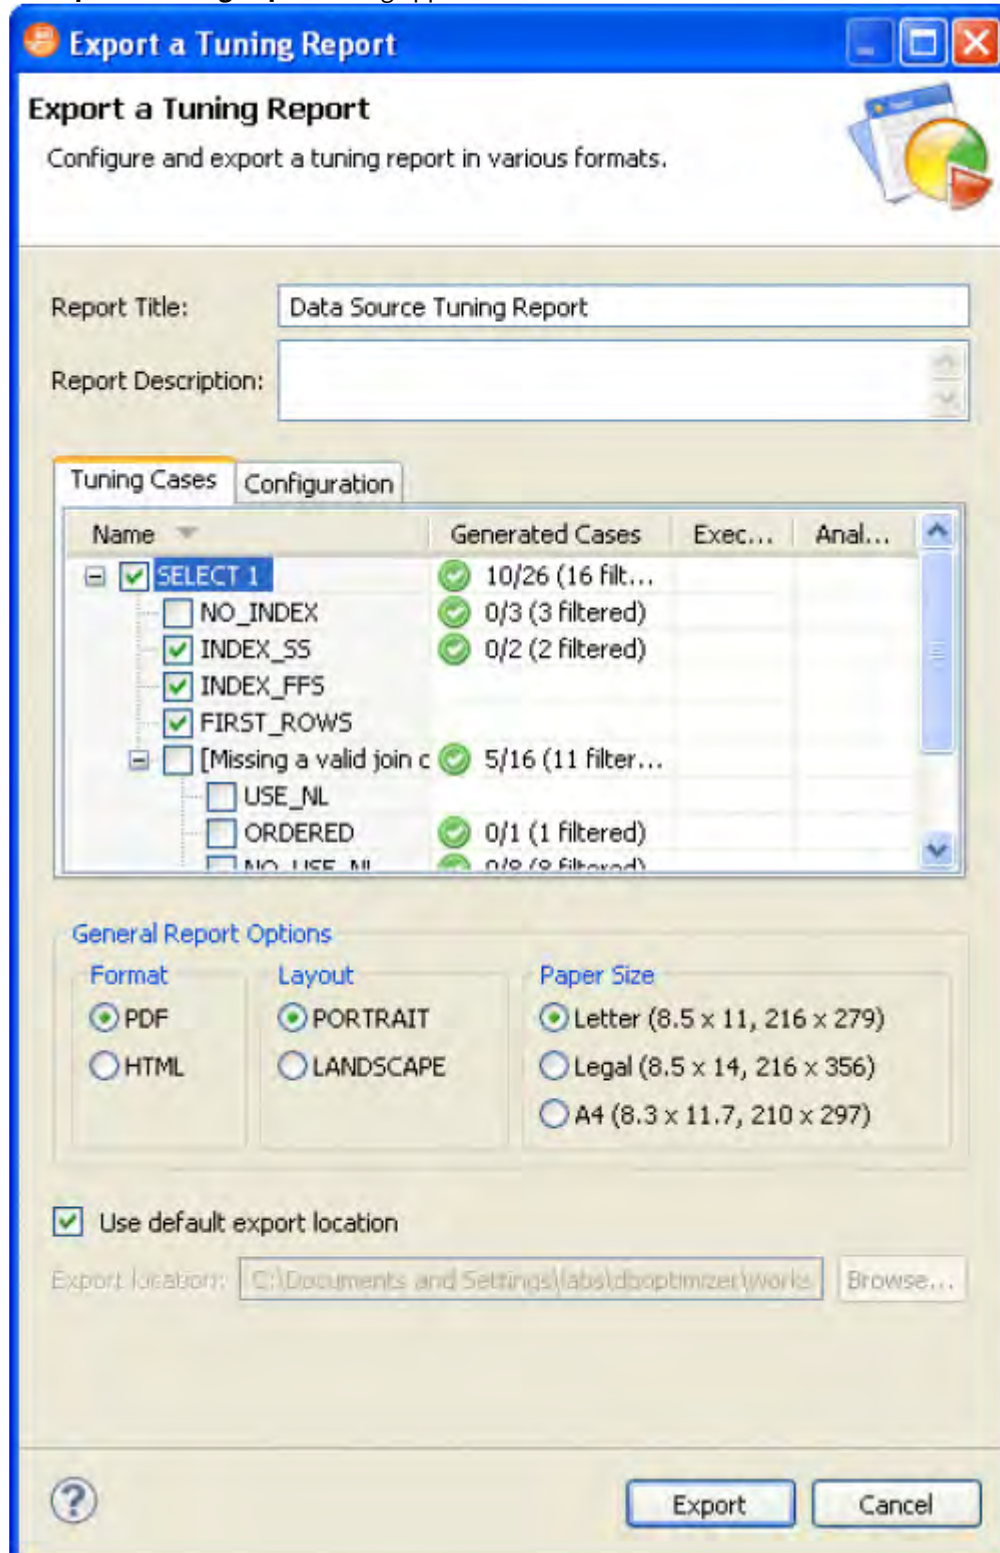
Creating tuning reports

After tuning SQL you can create an HTML or PDF Report of the tuning session. You can choose the details to include in the report.

1. At the top right corner of the Tuning tab, click the **Export Report** button.

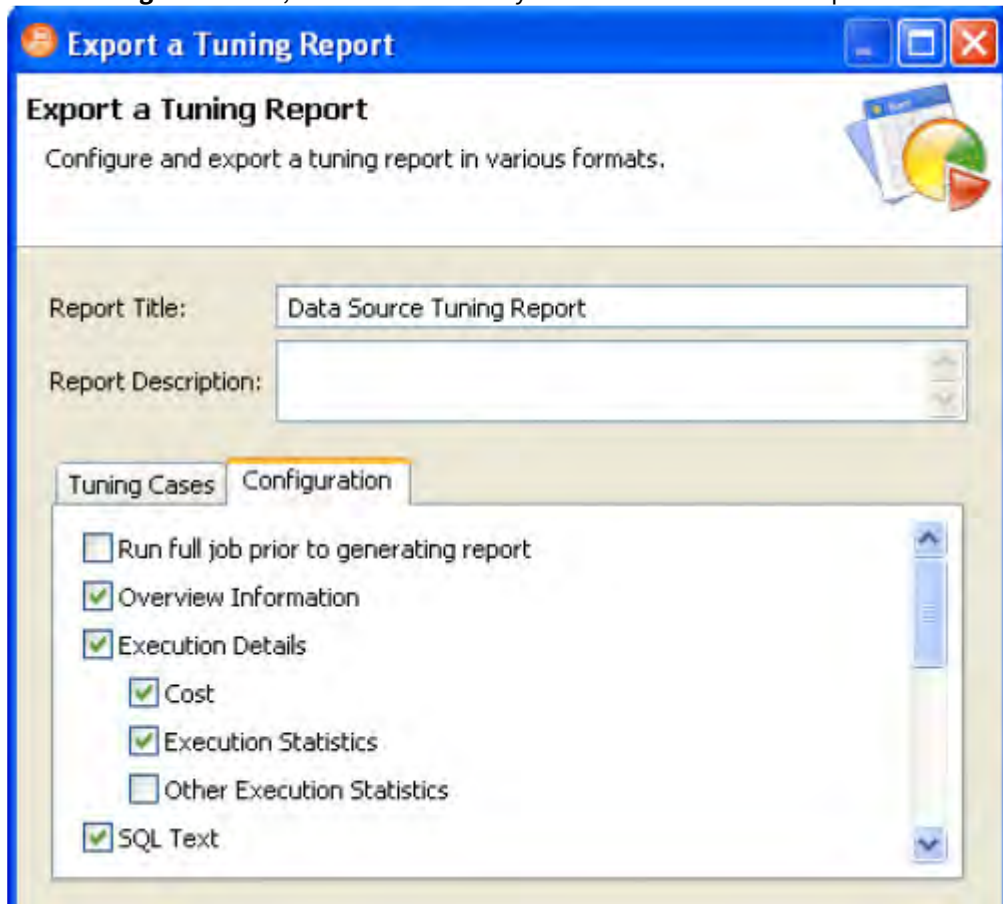


The **Export a Tuning Report** dialog appears.

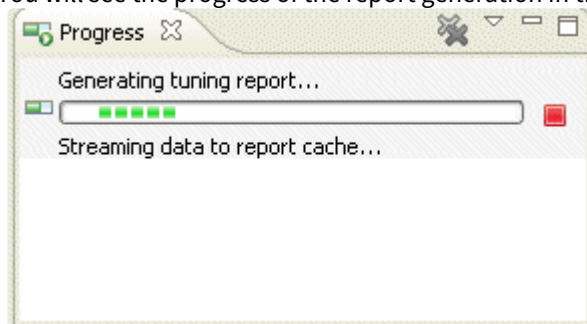


2. Enter a **Report Title** and **Description**.
3. On the Tuning Cases tab, choose the cases you want to report on. Click the + to expand the cases.

4. On the **Configuration** tab, choose the content you want to include in the report.

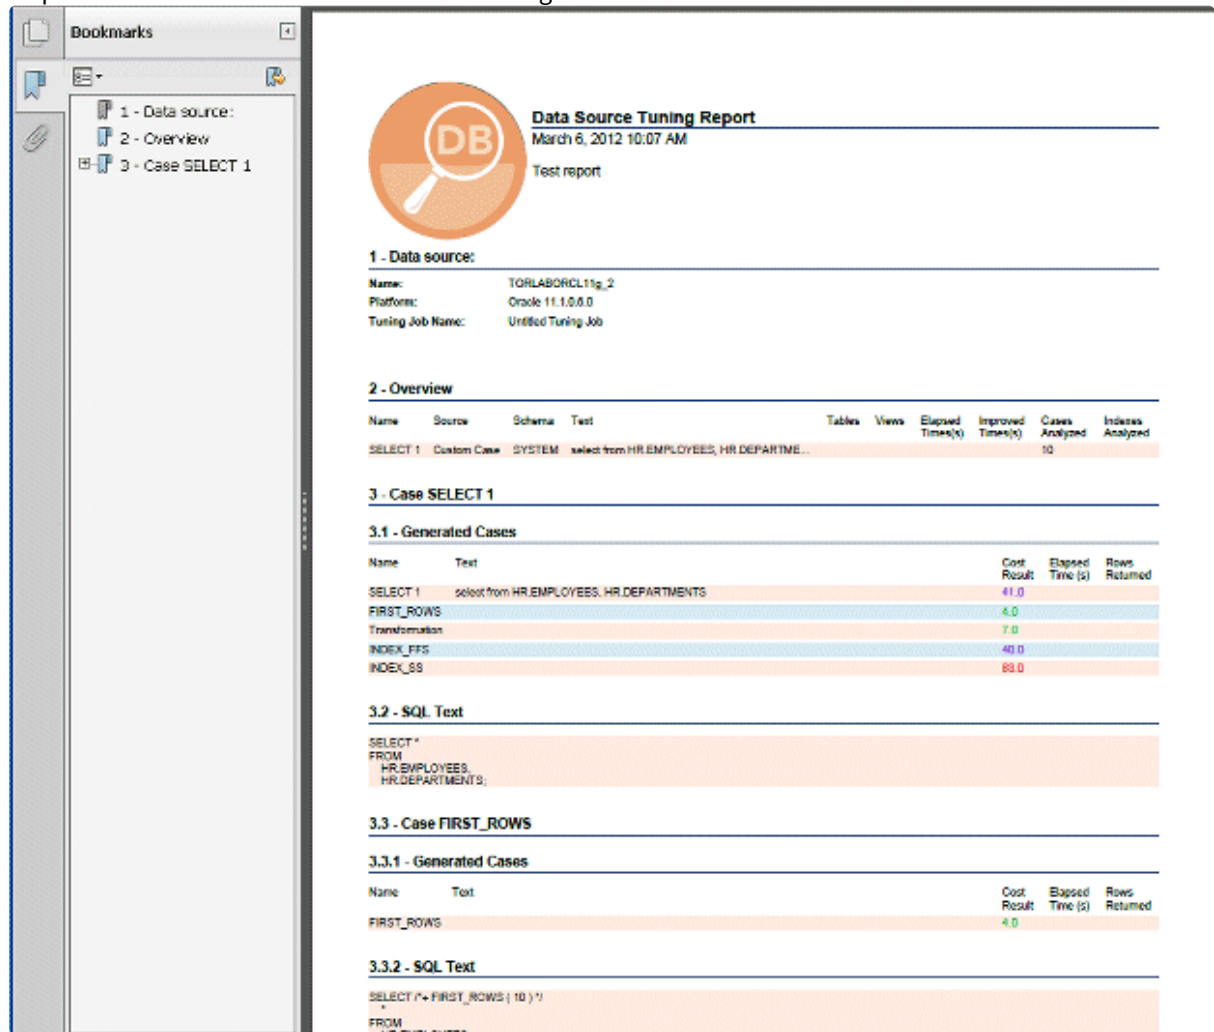


5. Select the **General Report** options, enter the location for the report, and then click **Export**. You will see the progress of the report generation in the Progress pane.



When complete, the report is stored at the top level of your workspace.





A report in PDF format will resemble the following:



Understanding the Analysis tab

Index analysis is started when you either generate cases with **Perform detail analysis** selected on the **Overview** tab, or when you click the **Analysis** tab. If any columns referenced in the WHERE clause of the tuning candidate are not the first column of an index, tuning will recommend that you create an index on that column.

The color-coded Index Analysis feature highlights missing indexes as well as shows which indexes are used and which are not used in the default execution path. The Index Analysis feature highlights issues where the database optimizer might not be using the preferred indexes. DB Optimizer also lists indexes on the tables that do not have fields in the WHERE clause helping the designer to see if adding an additional predicate in the WHERE clause might make use of an existing index.

| Text Color | Interpretation |
|---|---|
|  | Index is used in the query |
|  | Index is usable but not used in the current execution path. |
|  | This index is missing. DB Optimizer recommends that you create this index. |
|  | This index exists on the table but not usable in this query as it is written. |

In the **Collect and Create Indexes** table, orange-highlighted entries indicate missing indexes that DB Optimizer recommends be created to improve performance. Clicking on that index, displays text to the right outlining the rationale behind this recommendation.

For more information on using the Analysis tab, see [Using the Analysis tab](#).


Tuning SQL statements


A tuning job enables you to view the cost details of SQL statements on a registered data source and then select the best, or most efficient, array of execution path directives in order to make query execution faster, therefore improving the entire enterprise, overall.

A tuning job consists of a set of SQL statements and any analysis results you generate against a data source using tuning. The SQL statements and analysis results that compose a tuning job can be saved in a tuning file (.tun). This enables you to open a tuning job at a later time for inspection and analysis, to add, delete, or modify the SQL statements, or generate new execution statistics.

The following topics provide a high-level overview of the tuning process:

1. [Create a new tuning job](#)
2. [Specify a data source](#)
3. [Add SQL statements](#)
4. [Run a tuning job](#)
5. [Analyze tuning results](#)
6. [Modify tuning results](#)

 For additional commands that fall outside the general tuning workflow, but may still be helpful, see [Additional tuning commands](#).

 For information on working with data sources such as adding and browsing them, see [Working with data sources](#).

Create a new tuning job

You can create a new tuning job via the **File > New > Tuning Job** command, or by importing statements directly from profiling. The New Tuning Job icon is also available on the Toolbar.

To create a new tuning job via the Menu or Icon command

Select **File > New > Tuning Job**, or click the **New Tuning Job** icon on the Toolbar. Tuning opens.

You can now proceed to set up the parameters of the new job.

To create a new tuning job from profiling

After you have run a profiling session, in profiling's **Profiling Details** tab, select one or more statements, right-click, and select Tune from the context menu. Tuning opens, pre-populated with parameters based on the statements you selected.

To open an existing tuning job


Navigate to the **SQL Project** tab and double-click the name of the existing tuning job.

To name a job, save it


Ensure you specify a meaningful name that identifies the job in other views and dialogs. You can save the job by selecting **File > Save or File > Save All** from the Menu bar. Once a job is saved, it is added to the SQL Project view.

Specify a data source

The bread crumbs at the top of the tuning job window identify the data source where the SQL statements to be tuned reside. The default data source is the one that was selected when the new tuning job was initiated. For example, in the following image, we see that the data source is TORLABSCORCL, which is part of the Oracle data source group. The color of the bar at the top of the tuning window shows the category of the data source as defined in the data source configuration properties.



You can change the data source of a tuning job by clicking a bread crumb triangle and then navigating to the data source or using the filter to locate and then select a data source. In the following screenshot, Microsoft SQL Server was clicked and T was entered in the filter text area, which resulted in several matches.



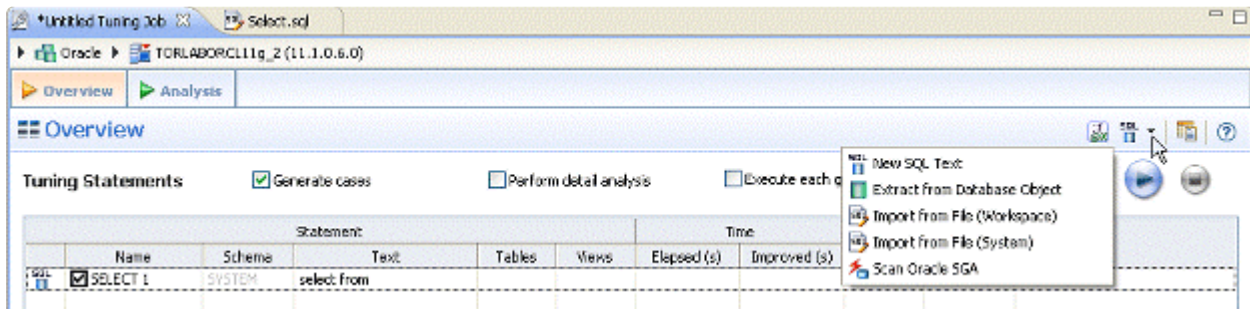
Click the name of the desired data source to affect the change.

i Multiple tuning jobs can be saved against the same data source. You can therefore set up your tuning jobs organizationally. You might for example, set up a tuning job to tune only SQL associated with procedures or a set of SQL sources that are functionally related. Alternatively, your tuning jobs may be organized by application.

Add SQL statements

Once you have created a tuning job and named it, using **File > Save As**, you need to add SQL statements to the job that are to be tuned. All standard DML statements can be tuned (SELECT, INSERT, DELETE, and UPDATE as well as MERGE on SQL Server 2008 and higher).

Statements are added to tuning via the Overview pane.



There are several different methods for adding SQL statements to a job, as reflected by the option in the New SQL text menu.

- **New SQL Text** enables tuning via manual entry, or cutting and pasting into the tuning window.
- **Extract from Database Objects** enables you to select stored SQL from the data source to which you are connected. You can either drag and drop objects from the Data Source Explorer or you can add database objects matching specified filters. For example, entering `t` in the filter area of the Data Source Objects Selection dialog, can match functions, materialized views, procedures, and views, whose name begins with `t`. You can then drag and drop the matches from the Data Source Explorer to the Tuning Statements grid.
- The **Import from File (Workspace)** and **Import from File (System)** options enable you to choose an SQL file saved in your workspace or elsewhere on your computer or network.
- The **Scan Oracle SGA** option is available for the Oracle platform only. It enables you to scan for and select active SQL in the System Global Area (SGA). For more information, see [Tuning SQL statements in the System Global Area \(Oracle\)](#).

To add an ad hoc SQL statement:

Select the **New SQL Text** option and manually type an SQL statement in the window, or copy/paste the statement from another source.

To add a database object:

1. Select the **Extract from Database Objects** option.
The Data Source Object Selection dialog appears where you can search for and then select the object you want to tune.
2. Type an object name prefix or pattern in the field provided. The Matching objects window automatically populates with all statements residing on the specified data source that match your criteria. Database objects include functions, materialized views, packages, package bodies, procedures, stored outlines, triggers, and views.

i In order to find matching objects, data source indexing must be enabled. To enable data source indexing, click Configure data source indexing, select **Enable indexing**, and then click **OK**.

3. Click the object you want to add. Ctrl-click to add more than one object to the job.
Click **OK**.

i Alternatively, after clicking the Database Objects tab, you can drag and drop objects from Data Source Explorer into the Database Objects window. As long as the dragged object is a valid object type, it will be added to the Database Objects tab.

To add an SQL *file*

1. From the New SQL Text menu, select either **Import from File (Workspace)** or **Import from File (System)**, depending on where the file you want to add is stored:

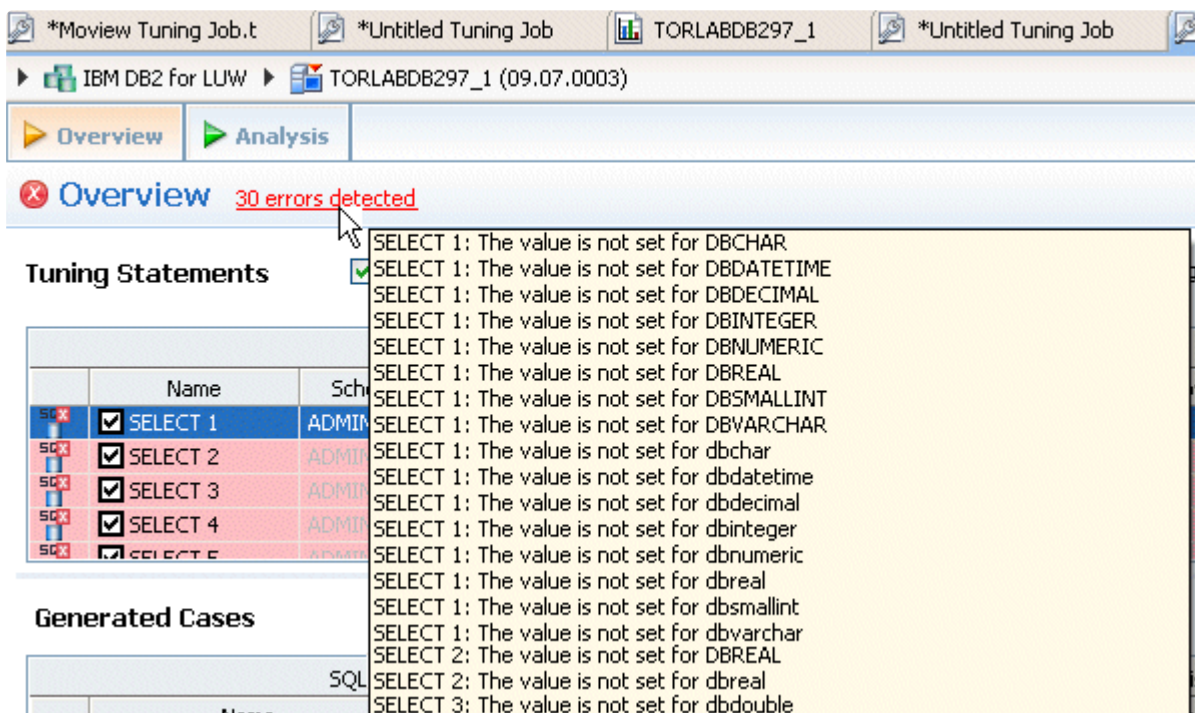
- Workspace files are files that reside in the application, meaning project files or other objects generated or stored in the system.
 - File System files are files that reside on your machine or the network.
2. Select a file from the dialog that appears. It is automatically added to the job.

To add SQL from the Oracle SGA

1. From the **New SQL Text** menu, select **Scan Oracle SGA**.
2. The system scans for SQL text which you can filter on the **Scan SGA** dialog that appears.
3. Choose the statement to be tuned and then click **Finish**.

Managing bind variable errors

When you try to tune a statement containing a bind variable you will be warned that either the type is not set or the value is not set. Mouse over the error to learn what problems were detected.

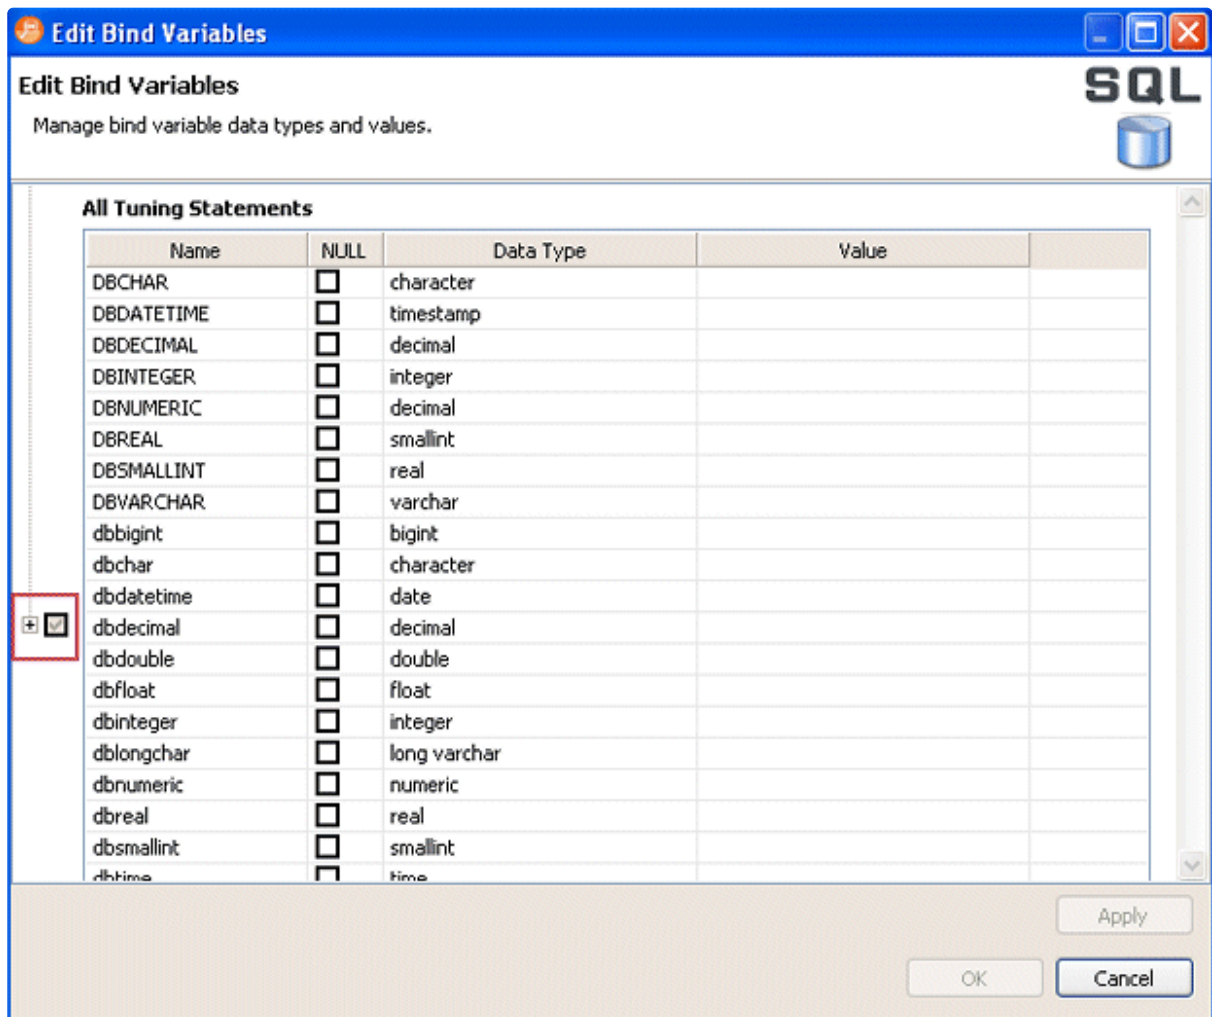


You can use the Bind Variable Editor to set the types or variables.

1. Click the **Edit Bind Variable** icon as shown below.

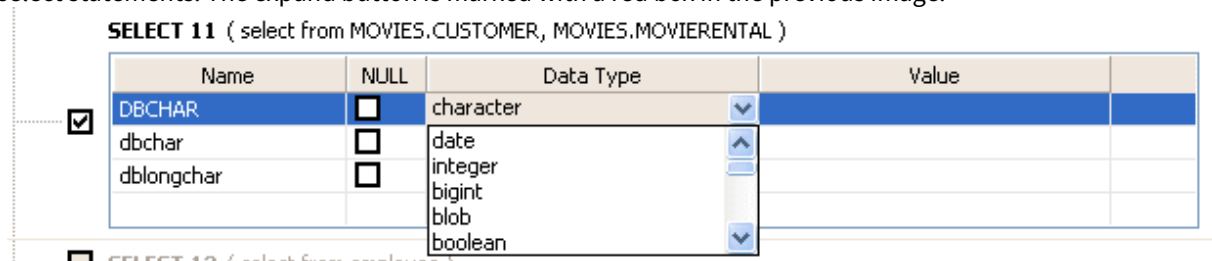


The **Edit Bind Variables** dialog appears.



You can either set the bind variables for all tuning statements at once or you can set the bind variables for each select statement individually. If you set the bind variable in the All Tuning Statement section, you can still override that setting for an individual select statement.

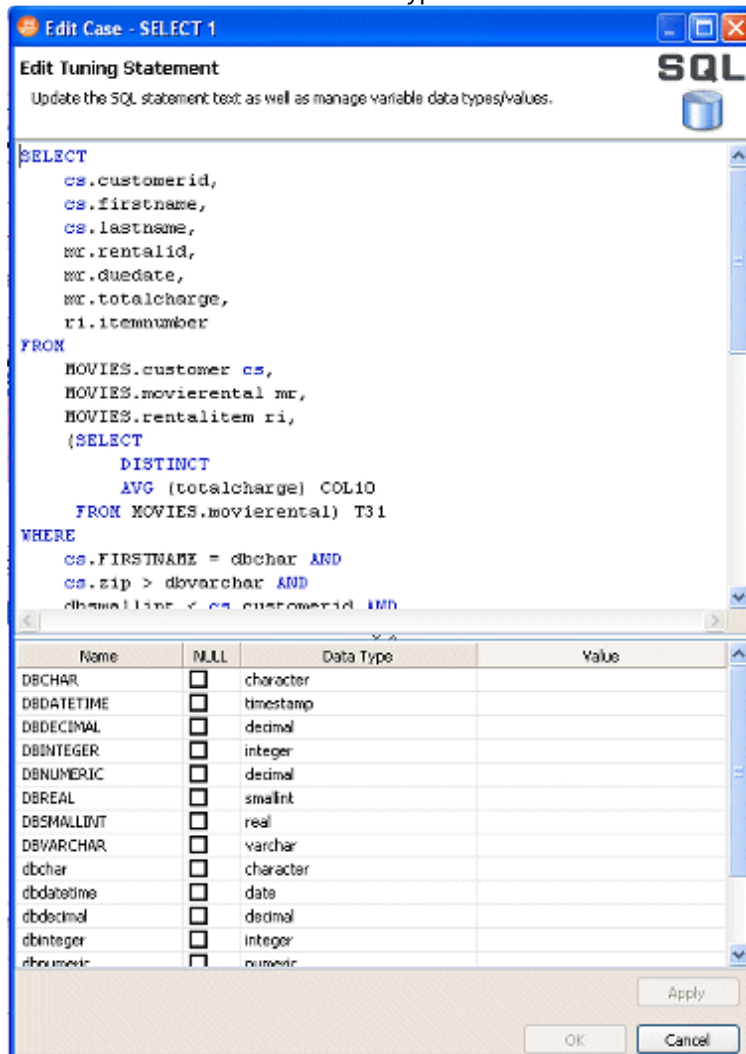
2. Set the bind variable by clicking the **NULL** box, clicking in the **Data Type** column and selecting the data type from the list, or by clicking in the **Value** column and entering the value.
3. If you want to set the bind variable for an individual select statement, click the expand button to see all the select statements. The expand button is marked with a red box in the previous image.



In this manner you can set the bind variables from all tuning statements and then override that setting by setting the bind variable for a specific select statement.

You may find it easier to set the bind variables when you can see the tuning statement. In the Generated Cases section, you can double click a statement and an editor appears where you can edit the SQL

statement and set the variable data types and values.



Run a tuning job

As you add SQL statements to the job on the Overview tab of the tuner, tuning-supported DML statements (SELECT, INSERT, DELETE, and UPDATE as well as MERGE on SQL Server 2008 and above) are parsed from the statements and added to the Overview tab in preparation for the tuning function execution.

Each tuning source statement is listed by Name, Schema, Text, Tables and Views. For SQL Server and Sybase platforms, there is also a Catalog column. Additionally, each statement will have Time and Analysis values that approximate how efficiently they execute on the specified data source.

In the Generated Cases area of the Overview tab of a tuning job, the Cost and Execution Statistics columns let you compare the relative efficiency of SQL statements or statement cases. While the explain plan Cost for a statement or case is calculated when you add SQL to a tuning job, the Elapsed Time and Execution Statistics (and Other Execution Statistics columns, if available) columns are not populated until you execute that statement or case.


If the Tuning Status Indicator indicates that a statement or case is ready to execute, you can execute one or more statements on the Overview tab. Alternatively, the Tuning Status Indicator may show that you have to correct the SQL or set bind variables before you can execute.

Once the tuning job has run, the Overview tab provides a series of cases, per statement, that you can select and modify based on your results.

In some cases, automatic case generation might be disabled (via the Preferences panel). If this is true, or if you have otherwise modified the Generated Cases table and can no longer generate a specific case, you can instead explicitly generate a case for specific statements.



To execute a tuning job

1. Ensure you have registered and selected a data source. For more information, see [Register data sources](#) and [Specify a data source](#).
2. Ensure you are connected to the database by double clicking the database name in the Data Source Explorer.
3. Click the tuning icon on the toolbar, or click **File > New > Tuning Job**.
4. On the **Overview** tab, specify the SQL you want to tune:
 - Modify the number of times to execute each statement in the **Execute each generated case** field at the top right of the tuner, as needed.
5. Click the execution button [] on the right side of the case generation field.
6. The tuning job runs, exacting and analyzing each statement and providing values in the appropriate columns.

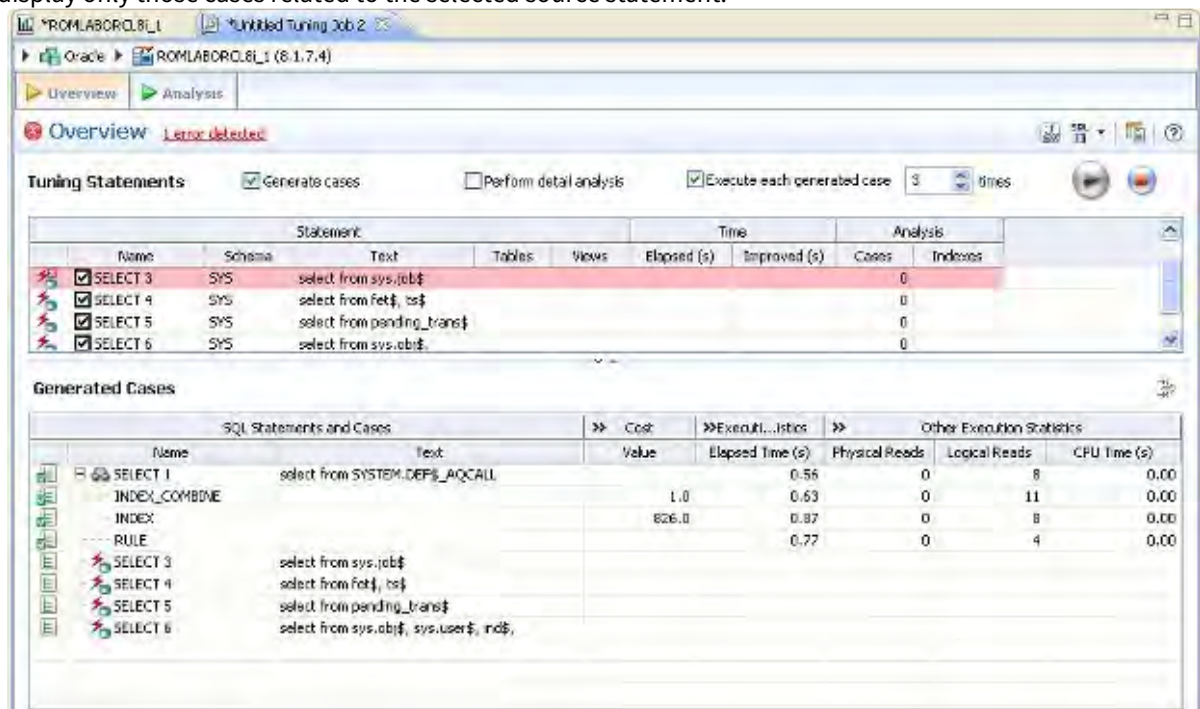
To explicitly generate a case for a specific statement

1. Ensure you are connected to the database by double clicking the database name in the Data Source Explorer.
2. Click the **Overview** tab.
3. In the Generated Cases area, right-click in the **Name** field of a statement or transformation case and select **Generate Cases** from the context menu, or click the **Overview Run Job** icon. The specified case is generated.

To view the generated cases for a specific statement

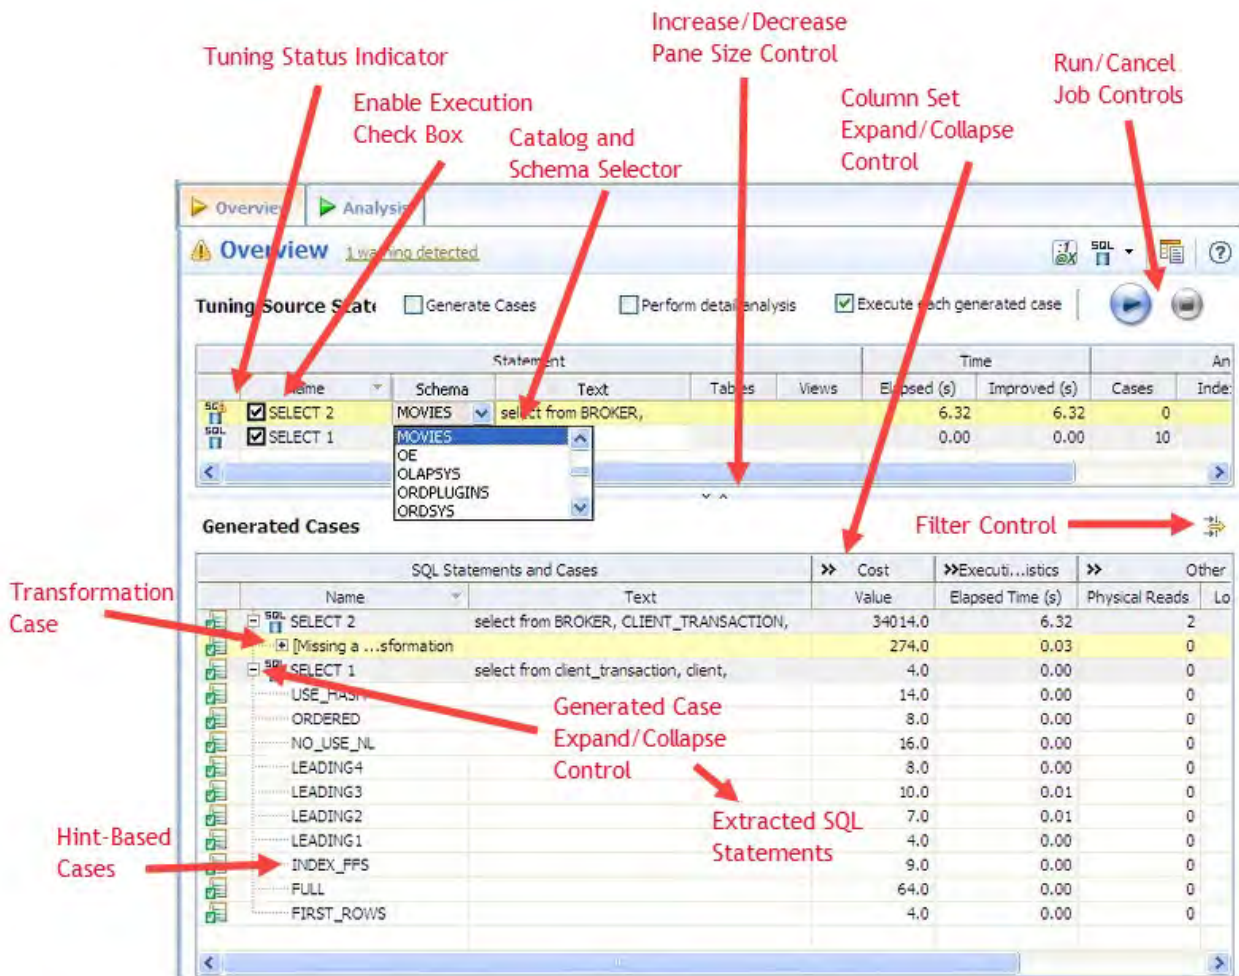
1. In the Tuning Statements area, click the checkbox to the left of the tuning source statement name. A check mark appears in the checkbox and the cases displayed in the Generated Cases area are filtered to

display only those cases related to the selected source statement.






Analyze tuning results



Once you have executed a tuning job, the **Overview** tab reflects tuning analysis of the specified statements. The **Analysis** tab shows the resulting analysis of the query, including indexes used, not used, and missing (or suggested to create). For more information on using the Analysis tab, see [Understanding the Analysis tab](#).



- The **Generated case Expand/Collapse** control lets you hide or display the hint-based cases and transformation-based case generated for a statement.
- The **Perform detail analysis** and **Execute each generated case** check boxes let you enable multiple statements or cases for simultaneous execution while the Run/Cancel Job controls let you start and stop simultaneous execution.
- Use the **Schema and Catalog Selectors** to select a schema and catalog for the tuning job. The catalog selector is available only for SQL Server and Sybase data sources. By specifying the schema and catalog, the tuner can use the paths of the schema and catalog selected to find the tables queried in the job rather than use the paths of the schema and catalog used to connect to the data source. If you change the schema or catalog used in a tuning statement you will need to refresh the tuning statements in order for new cases to be generated, which take into consideration the schema used. Right-click a tuning statement, and then select **Refresh Tuning Statements**.
- The **Column set Expand/Collapse** controls let you expand a column set to display more of the columns within the table.
- The **Tuning Status Indicator** indicates whether a statement or case is ready to execute or has successfully executed. The following table provides information on the Tuning Status Indicator states:


| Icon | Description |
|---|--|
|  | The case has not been executed. There are no errors or warnings and the case is ready to be executed. |
|  | The case has been successfully executed. |
|  | Execution for this case failed or was cancelled due to execution time exceeding 1.5 of original case time. |

Hovering the mouse over the Tuning Status Indicator displays a tip that notes the nature of a warning or error.


 If a warning  indicates that one or more tables do not have statistics, you can right-click the statement and select **Analyze Tables** to gather statistics. A warning may also indicate that the tuning statements are out of sync, in which case you can right-click a tuning statement and select **Refresh Tuning Statements**.

A warning can indicate an object caching error. For example, a table may not exist or not be fully qualified. Cases cannot be generated for the associated statement.

- The explain plan-based **Cost** field can be expanded to display a graphical representation of the values for statements or cases. Similarly, after executing a statement or case, the **Elapsed Time** field can be expanded to display a graphical representation. The bar length and colors used in the representation are intended as an aid in comparing values, particularly among cases. For example:



In the case of both **Cost** and **Elapsed Time**, the values for the original statement are considered the baseline values. With respect to color-coding for individual case variants, values within a degradation threshold (default 10%) and improvement threshold (default 10%) are represented with a neutral color (default light blue). Values less than the improvement threshold are represented with a distinctive color (default green). Values greater than the degradation threshold are shown with their own distinctive color (default red). With respect to bar length, the baseline value of the original statement spans half the width of the column. For child-cases of the original statement, if one or more cases show a degradation value, the largest degradation value spans the width of the column. Bar length for all other children cases is a function of the value for that case in comparison to the highest degradation value.

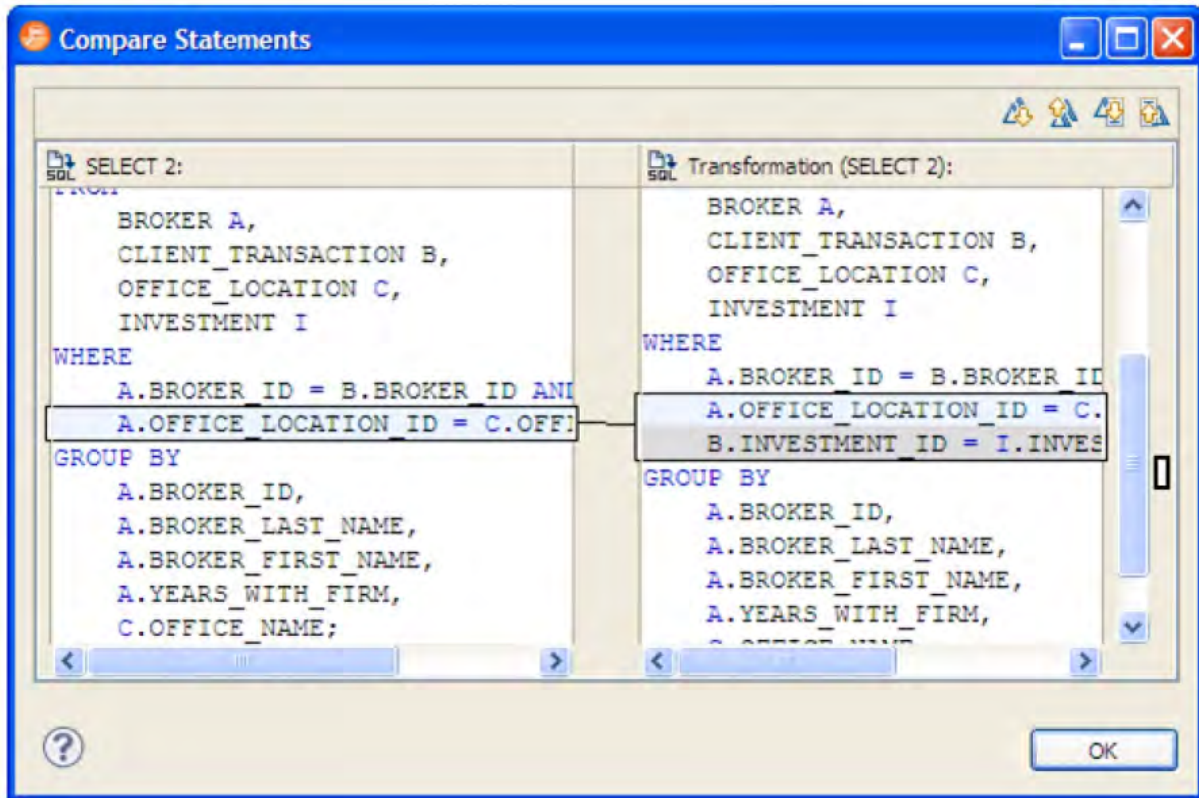
 For information on specifying colors, and the improvement threshold and degradation threshold values used in these graphical representations, see [Specify Tuning Job Editor preferences](#).

Additionally, once results have been generated you can:

- **Compare Cases.** For more information, see [Compare Cases](#).
- **Filter and Delete Cases.** For more information, see [Filter and Delete Cases](#).
- **Visual SQL Tuning.** For more information, see [Visual SQL tuning](#).
- **Create an Outline.** For more information, see [Create an Outline](#).

Compare cases

You can compare cases between an original statement and one of its tuning-generated statements, or another statement case via the Compare to Parent and Compare Selected commands, respectively.



To compare a case side-by-side with its parent

Right-click in the **Name** field of a case and select **Compare to Parent** from the context menu.

To compare two cases

Select the two cases, and then right-click in the **Name** field of either case. Select **Compare Selected** from the context menu.

Filter and delete cases

You filter cases from the Generated Cases table via the Filter icons on the Generated Cases Toolbar of the **Overview** tab.



Filter the cases on the **Overview** tab so that hints that are not improvements on the original statement are not displayed. You can filter:

- Non-optimizable statements
- Optimized statements
- Worst cost cases
- Worst elapsed time cases

When filtering, the criteria remain in effect until you change the criteria. That is, as new cases are generated, only those cases that do not satisfy the filtering criteria are displayed. To restore an unfiltered set of cases, open the **Filter** dialog and deselect the filtering options.

When removing cases, the criteria you set has no effect on cases subsequently generated.

To filter cases from the Overview table

1. Click the **Filter** button, respectively. A **Filters** dialog opens.
2. Use the check boxes to select your filtering and then click **OK**.

To delete cases from the Overview table

1. Right-click on the row of the case you want to delete and select **Delete**. A **Delete** dialog opens.
2. Use the check boxes to select your filtering and then click **OK**.

When removing cases, the criteria you set has no effect on cases subsequently generated.

Create an outline

If SQL is executed by an external application or If you cannot directly modify the SQL being executed but would like to improve the execution performance, you can create an outline on the Oracle platform. An outline instructs the Oracle database on the execution path that should be taken for a particular statement.

To create an outline for a change suggested by a case

1. On the Overview tab of a tuning job, right-click in the **Name** field of a case and select **CreateOutline** from the context menu.
A **NewOutline** wizard opens.
2. On the first panel, provide an **OutlineName**, select an **Outlinecategory**, and then click **Next**.
A **Preview Outline** panel opens previewing the SQL code to create the outline.
3. Select an **Actiontotake** option of **Execute** or **OpeninnewSQLEditor** and then click **Finish**.
For more information, see [Using the Outlines tab \(Oracle\)](#).

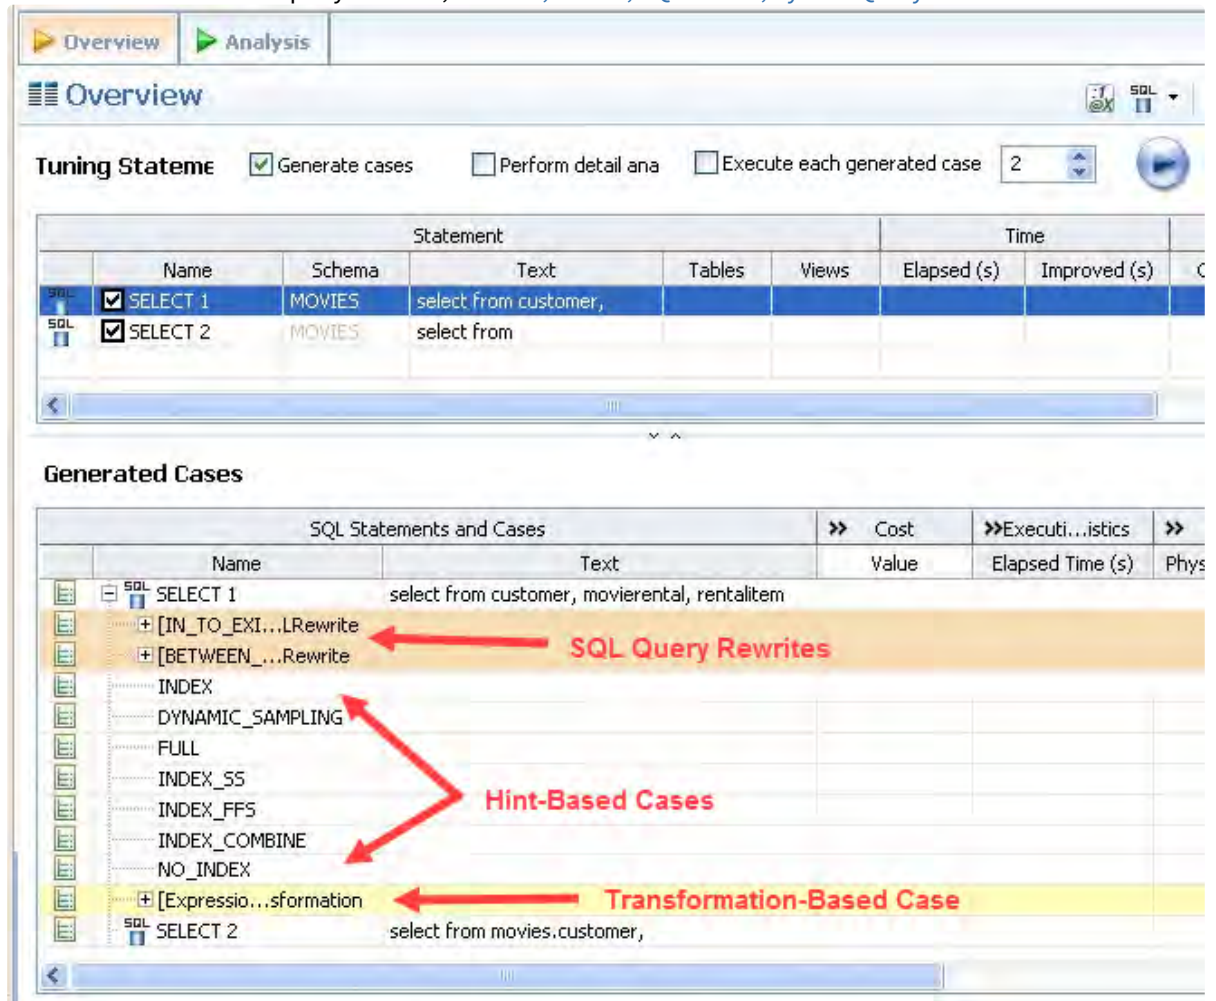
Modify tuning results

As you add SQL source to the Overview tab of a tuning job, the supported DML statements are automatically parsed out and a numbered statement record for each statement is added to the Overview tab.

Cases generated from tuning candidates are alternative forms of the original statement that have been optimized or otherwise "fixed" by the tuning function. Once you have executed a tuning job, tuning automatically generates all SQL optimizer hint-based variations that can be applied to the statement: If you change the schema of a statement

- All SQL Optimizer hint-based variations that can be applied to a statement.
- A transformation-based case, if any of the eight common quick fixes can be applied to an SQL statement. This feature leverages the DB Optimizer Code Quality Check functionality. See [Understanding Code Quality Checks](#) for more information on the eight quick fixes. A transformation case, in turn, has its own set of SQL Optimizer hint cases. For information on query rewrites, see [DB2, Oracle, SQL Server, Sybase Query Rewrites](#). For information on other transformations, see [Examples of transformations and SQL query rewrites](#).

- SQL Query Rewrites may be suggested when tuning. For example, a recommended rewrite for EXISTS may be IN. For information on query rewrites, see [DB2](#), [Oracle](#), [SQL Server](#), [Sybase Query Rewrites](#).



Hint-based cases and the transformation-based cases are a special case of the statement records added to the Overview tab as you add candidates to a tuning job. With the exception of the Text, Source, and Index Analysis fields, cases are identical to the standard statement record. Similarly, execution, statistics collection, and other options available for basic statement records are available for individual cases.

Once cases have been generated, if you have the required permissions on the specified data source, you can apply the changes suggested by hint and transformation based cases in the Overview table.

DB2, Oracle, SQL Server, and Sybase query rewrites

The following query rewrites or transformations may be recommended during tuning. The examples below are for Oracle data sources. The implementations for DB2, SQL Server, and Sybase data sources are slightly different. These rewrites are available on all four platforms except for those noted for ORACLE and DB2 only.

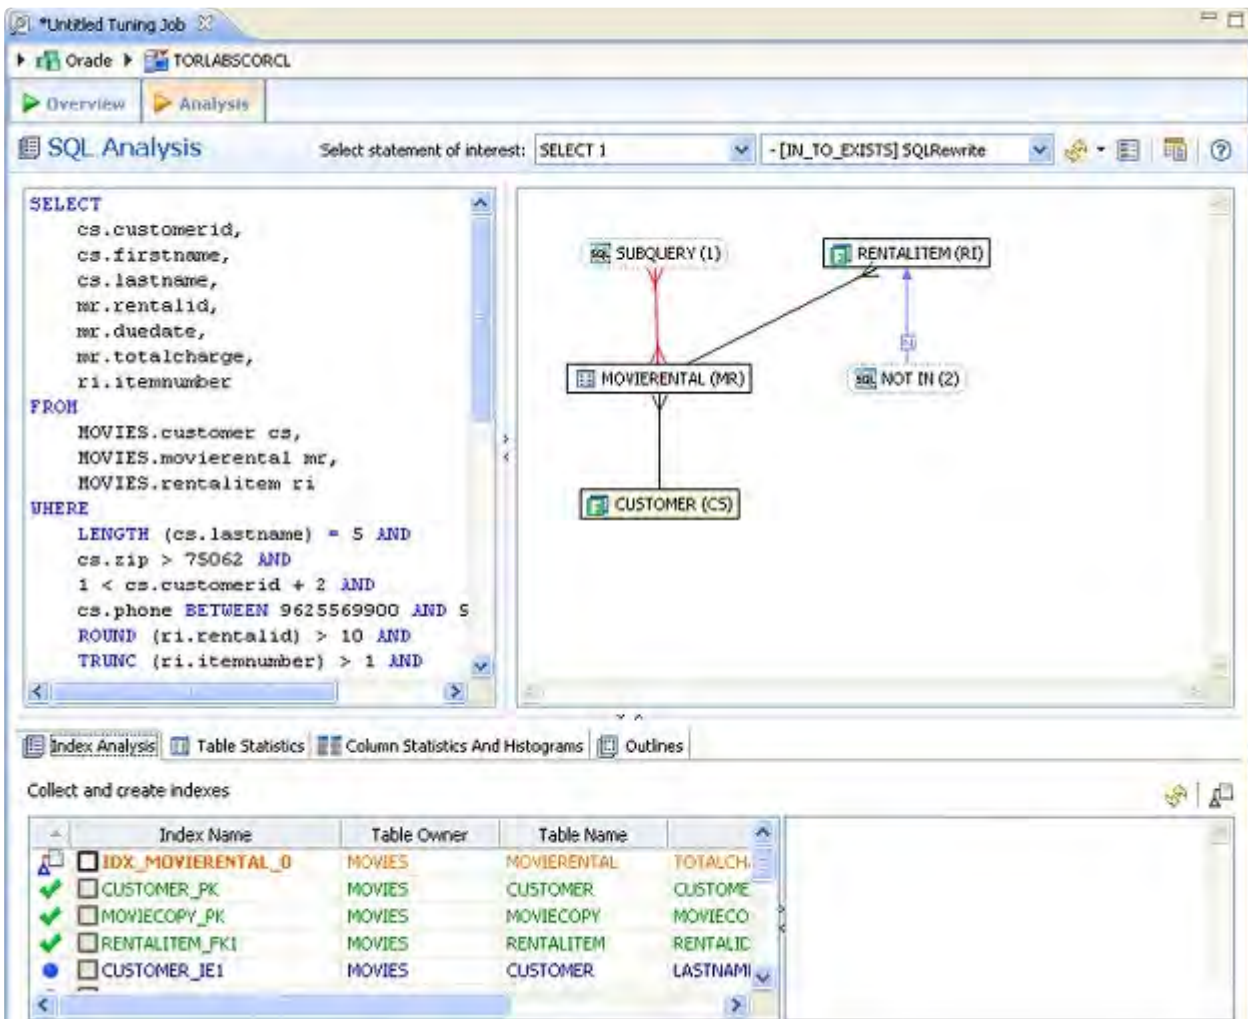
| Before | After |
|---|--|
| <pre>select * from t1 where EXISTS (select null from t2 where t2.key=t1.key);</pre> | <pre>select * from t1 where t1.key IN (select t2.key from t2);</pre> |

| Before | After |
|---|--|
| select * from t1 where NOT EXISTS (select null from t2 where t2.key=t1.key); | select * from t1 where t1.key NOT IN (select t2.key from t2 where t2.key is not null); |
| select * from t1 where t1.key IN (select t2.key from t2); | select * from t1 where EXISTS (select null from t2 where t2.key = t1.key); |
| select * from t1 where t1.key NOT IN (select t2.key from t2 where t2.key is not null); | select * from t1 where NOT EXISTS (select null from t2 where t2.key = t1.key); |
| select * from t1 where NOT EXISTS (select null from t2 where t2.key = t1.key); | select t1.* from t1, t2 where t1.key = t2.key(+) and t2.key is null |
| select * from t1 where t1.key NOT IN (select t2.key from t2 where t2.key is not null); | select t1.* from t1, t2 where t1.key = t2.key(+) and t2.key is null; |
| select column BETWEEN X AND Y | select (column <= X AND column >= Y) |
| select column NOT BETWEEN X AND Y | select (column < X AND column > Y) |
| select (column<= X AND column >= Y) | select column BETWEEN X AND Y |
| select (column < X AND column > Y) | select column NOT BETWEEN X AND Y |
| select t1.* from t1, t2 where t1.key = t2.key and t2.col = 10; | select t1.* from t1, (select * from t2 where t2.col = 10) inline_alias where t1.key= inline_alias.key; |
| <i>For DB2 and Oracle only</i> select t2.* from t1, t2 where t1.key = t2.key and t1.col is null | select * from t2 where t2.key IN (select t1.key from t1 where t1.col is null) |

Using the Analysis tab

The **Analysis** tab provides detailed information about statements and cases selected from the **Overview** tab, after a tuning job has been executed. It also shows filter ratio, and table and join sizes.

The **Analysis** tab contains information about the statement or case, its full SQL code, a diagram of the SQL statement, and Index Analysis.



Additionally, for the Oracle and SQL Server platforms there are Table Statistics, Column Statistics and Histograms, and Outlines/Plan Guides tabs. For more information, see [Using platform-specific features](#).

Statement analysis is performed when you click **Perform detail analysis** on the **Overview** tab and then click **Run Job** or when you click the **Analysis** tab. In order to view and analyze statement statistics, select the tab (Index Analysis, Table Statistics, Column Statistics and Histograms, or Outline) and the statements whose statistics you want to analyze.

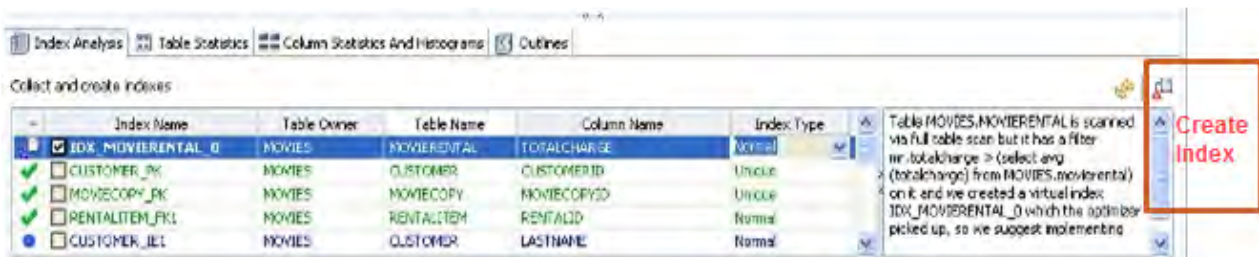
Next to the **Select statement of interest** list at the top, you choose to see an analysis of the **>ROOT** statement, or you can click the list and see an analysis of any one of the generated cases produced by running the tuning job from the **Overview** tab.

For more information, see [Visual SQL tuning](#).

Implementing index analysis recommendations

Once you have added tuning candidates to a tuning job, DB Optimizer can analyze the effectiveness of the indexes in the database and recommend the creation of new indexes where the new indexes can increase performance.

In the **Collect and create indexes** table, any indexes DB Optimizer recommends you create are marked in orange.




To accept the suggestion and have tuning automatically generate an index

1. For any recommended index, click the checkbox to the left of the index. Optionally, modify the Index type by clicking in the **Index Type** column and then selecting a type from the list.
2. Click the **Create Indexes** button. The **Index Analysis** dialog appears.

To view the index SQL in an editor for later implementation, click the statement and then click **Open in a SQL editor**.

To run the index SQL and create the index on the selected database, click **Execute**.

Visual SQL tuning

 Visual SQL Tuning is not available in DB Optimizer Developer.

DB Optimizer can now parse an SQL query and analyze the indexes and constraints on the tables in the query and display the query in graphical format on The Visual SQL Tuning (VST) diagram, which can be displayed in either Summary Mode or Detail Mode. This helps developers, designers and DBAs see flaws in the schema design such as Cartesian joins, implied Cartesian joins and many-to-many relationships. The VST diagram also helps the user to more quickly understand the components of an SQL query, thus accelerating trouble-shooting and analysis.

This section is comprised of the following topics:

- [Changing diagram detail display](#)
- [Interpreting the VST diagram graphics](#)

Changing diagram detail display


This section is composed of the following topics:

- [Choosing the tuning statement and generated case to analyze](#)
- [Viewing the VST diagram legend](#)
- [Viewing table counts and ratios](#)
- [Viewing the Explain Plan](#)
- [Viewing the VST diagram in Summary mode](#)
- [Viewing the VST diagram in Detail mode](#)
- [Changing detail level for a specific table](#)
- [Viewing all table fields](#)
- [Viewing diagram object SQL](#)
- [Expanding views in the VST diagram](#)
- [Viewing the Oracle explain plan overlay](#)
- [Expanding subqueries and nested subqueries](#)

Choosing the tuning statement and generated case to analyze


If, from the Overview tab, you have run the tuning job using more than one tuning statement, from the Analysis tab, you can choose to see the SQL analysis of any one of the tuning statements.

1. On the Analysis tab, click the **Select statement of interest** list and choose the tuning statement you want to see analyzed here.




Notice that next to the statement of interest box another list **>ROOT**. This shows that the statement being analyzed is the original statement, without any of the generated cases. This is the default selection.

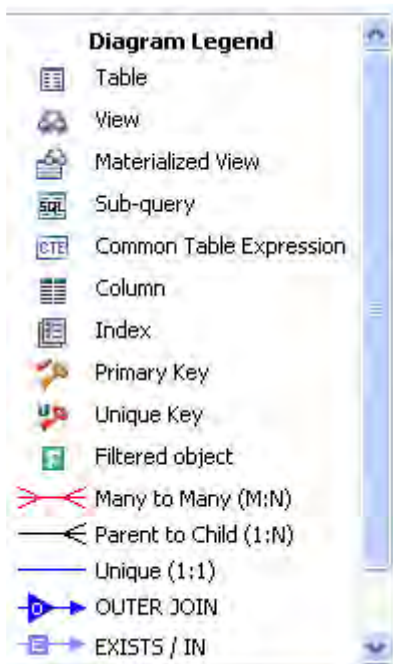
2. To choose the generated case to be analyzed, click the second **Select statement of interest** list and choose a generated case.




After you make your selections, a new analysis is performed taking into consideration the statement and case you chose. A new diagram is displayed and the Index Analysis, Table Statistics, Column Statistics And Histograms, and Outlines are recalculated and updated.

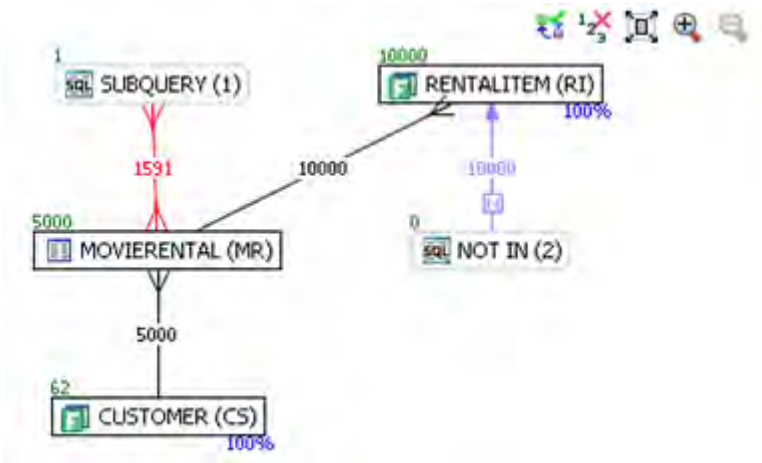
Viewing the VST diagram legend

Click the **Diagram Legend** toggle [] to view the legend and then click it again to hide it. All the icons used in the VST diagram are identified and in the Diagram Legend.



Viewing table counts and ratios

To view or hide table counts, two table join sizes and filtered result set ratios, click the Ratios and Counts icon [].

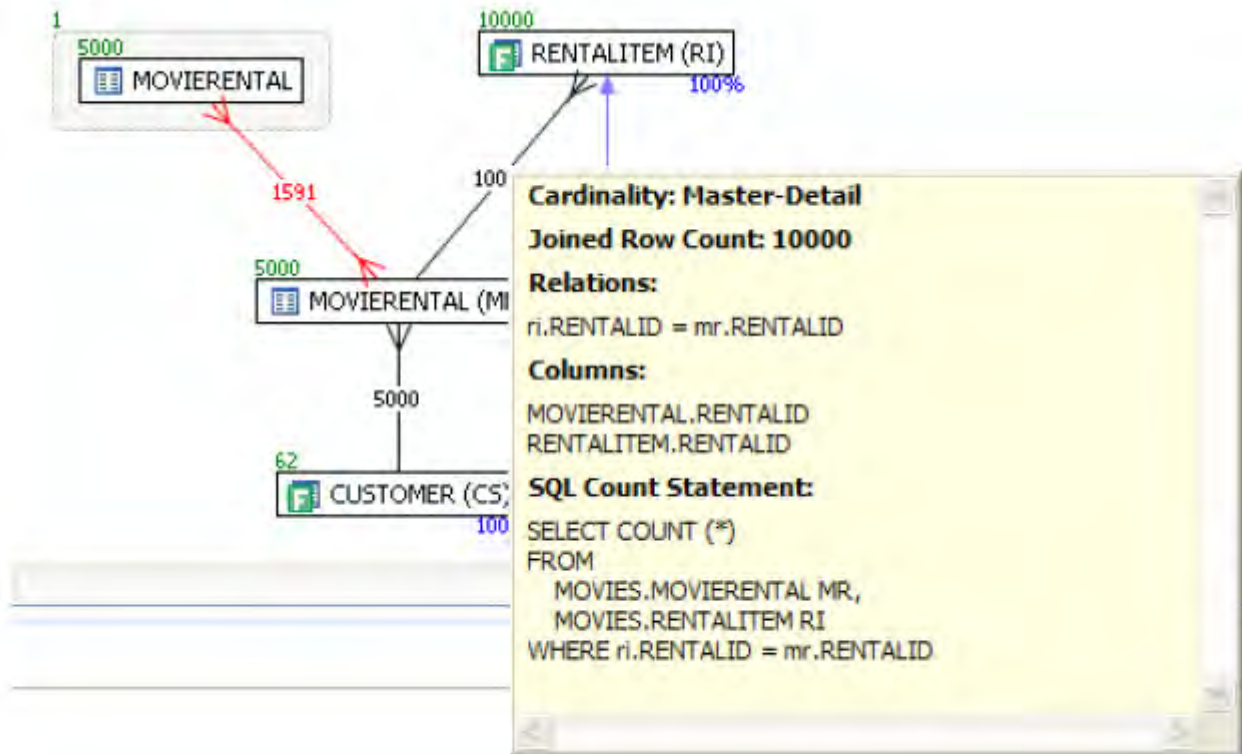


Green numbers at top left of table represent the total number of rows in that table. In the above the MOVIERENTAL (MR) table has 5000 rows.


Blue percentage at the bottom right of the table represent the percentage of rows in that table that meet the selection criteria. In the above example, 100 percent of the rows in the RENTALITEM (RI) table have met the selection criteria.


The numbers on the table joins indicate the total number of rows that meet the selection criteria for both tables.

You can also view the SQL Query that created a relationship by hovering over the relationship. If the tool tip content is larger than the size of the tool tip rectangle, you can hover the mouse on top of the tooltip for a second, and then it will turn into a dialog you can re-size, scroll in, and select text from to copy into the Clipboard.

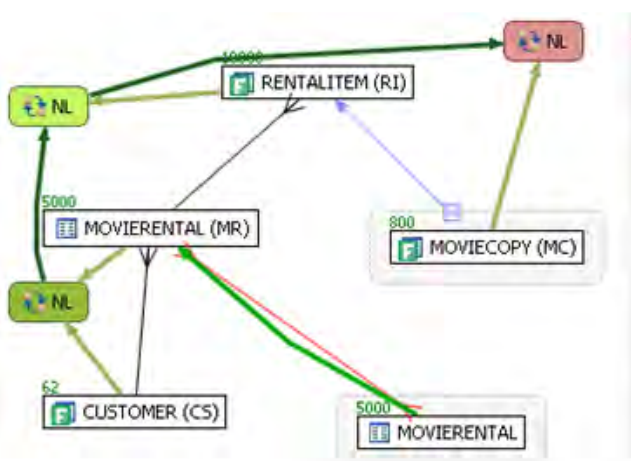


Viewing the Explain Plan

 The Explain Plan is available only for the Oracle 10g and later platforms.

Hover the mouse over the VST diagram to see the **Explain Plan** icon [] and then click it to view the Explain Plan Overlay. Click the **Explain Plan** icon again to hide the overlay.

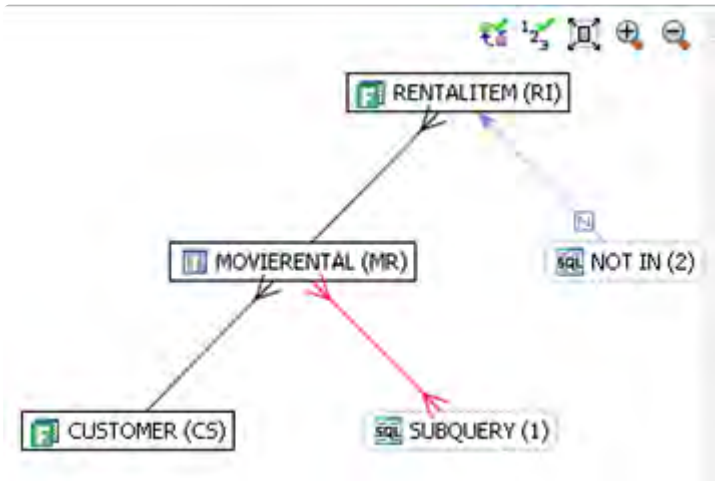
The additional nodes shown in the Explain Plan overlay provide details on the flow of the query plan, with operations (such as nested loops, sorts, and joins) showing connecting tables and other operations.



Hover the mouse over the objects or relationships in the overlay to view additional details.

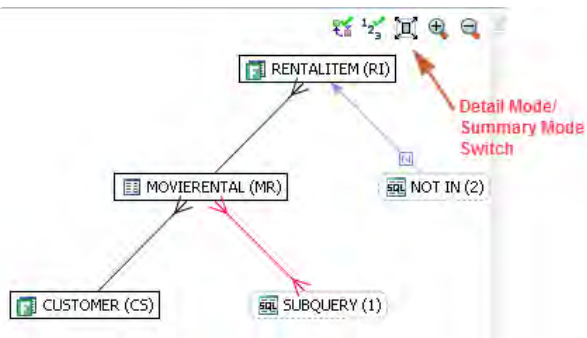
Viewing the VST diagram in Summary mode

By default the diagram displays Summary mode, showing only table names and joins, as seen in the following illustration

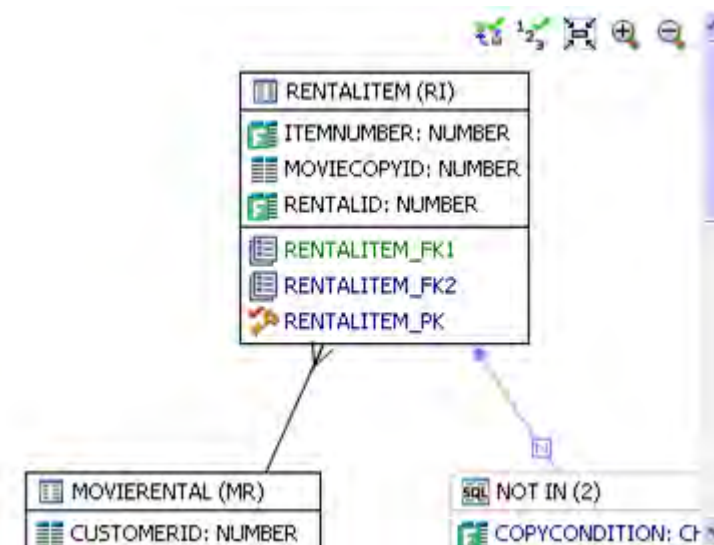


Viewing the VST diagram in Detail mode

By default, the VST diagram displays in Summary Mode, but by clicking the **Detail Mode/Summary Mode** switch.



Additional details of the tables display, including table columns and indexes

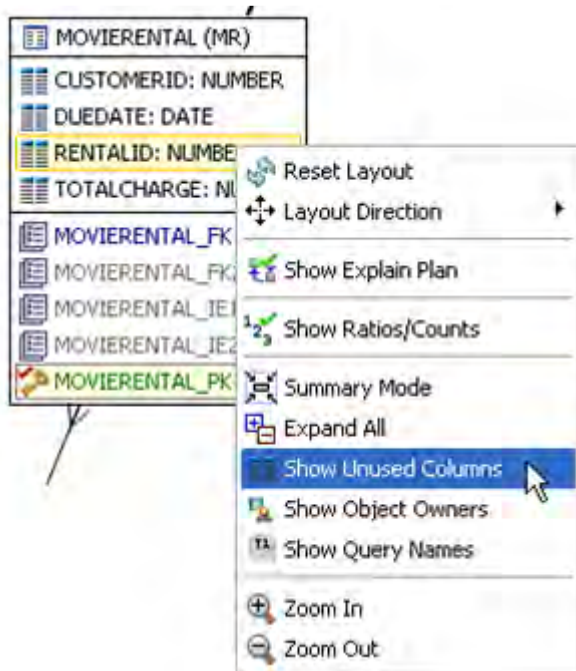


Changing detail level for a specific table

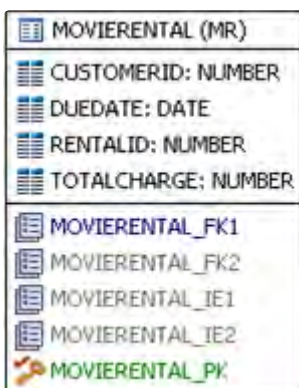
You can also switch between Summary Mode and Detail Mode for a specific table or view, by double-clicking the object name.

Viewing all table fields

By default, only fields that are used in the WHERE clause are displayed in detail mode; however, if you right-click the table you can choose to display even unused columns as follows:

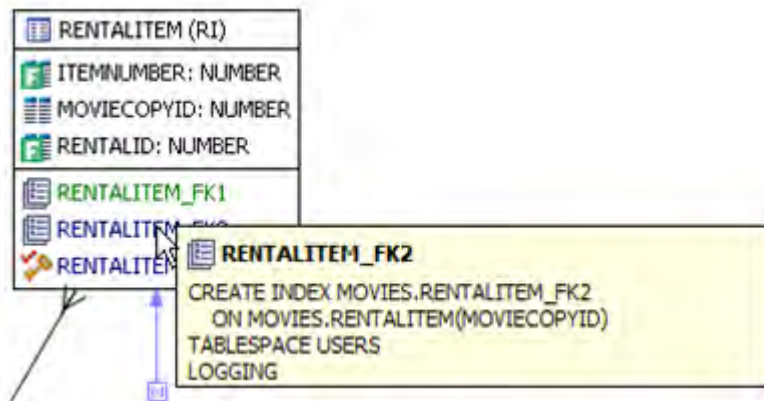


All the columns in the table are shown, and not just the ones used in the WHERE clause of the SQL statement.

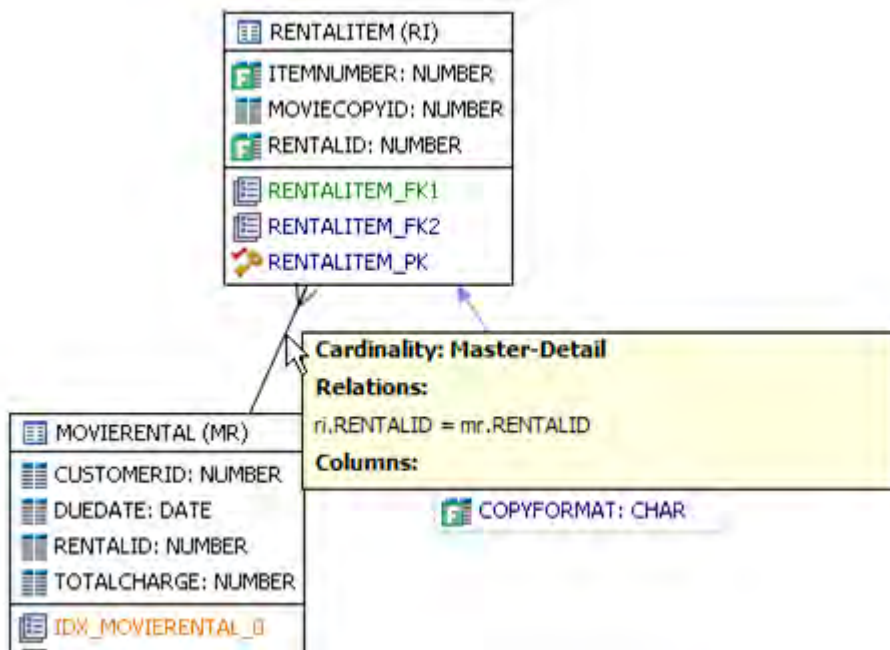


Viewing diagram object SQL

While in Detail Mode, hovering the mouse over the sub query, table name, field, or index displays the SQL required to create that object.



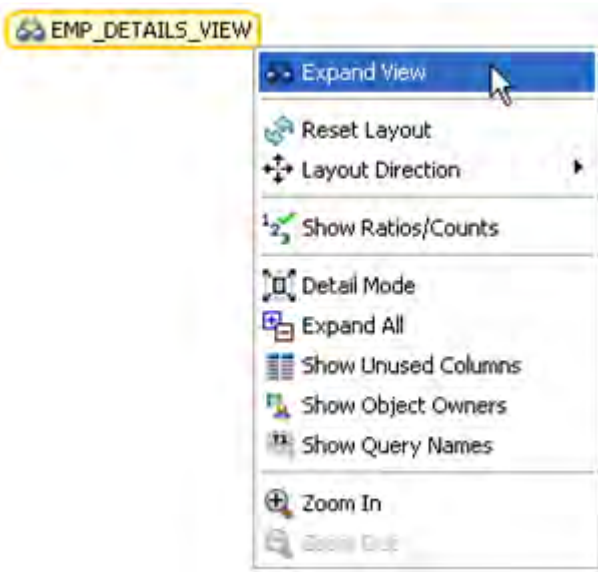
Hovering over the join between two tables displays the relationship between the two tables.



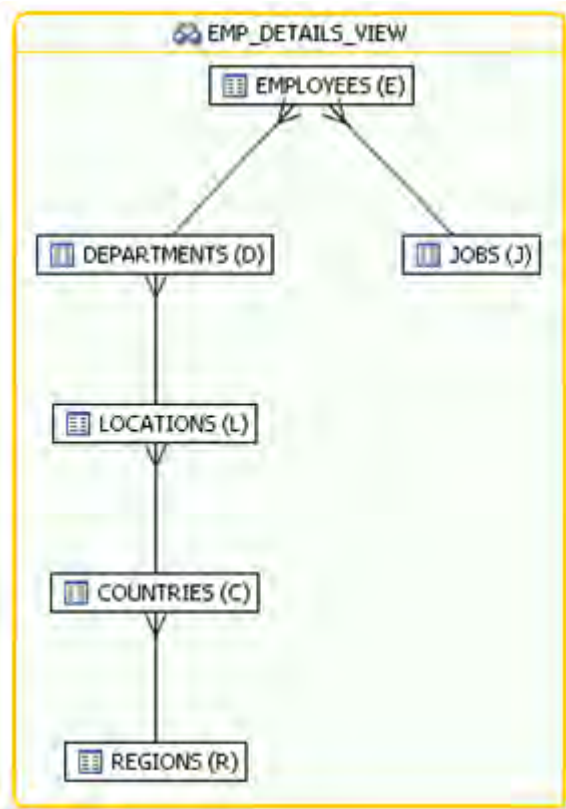
Expanding views in the VST diagram

If there are views in the Visual SQL Tuning diagram, they can be expanded by right clicking the view name and choosing **Expand View**.

Right click on the view, and choose **Expand View**.



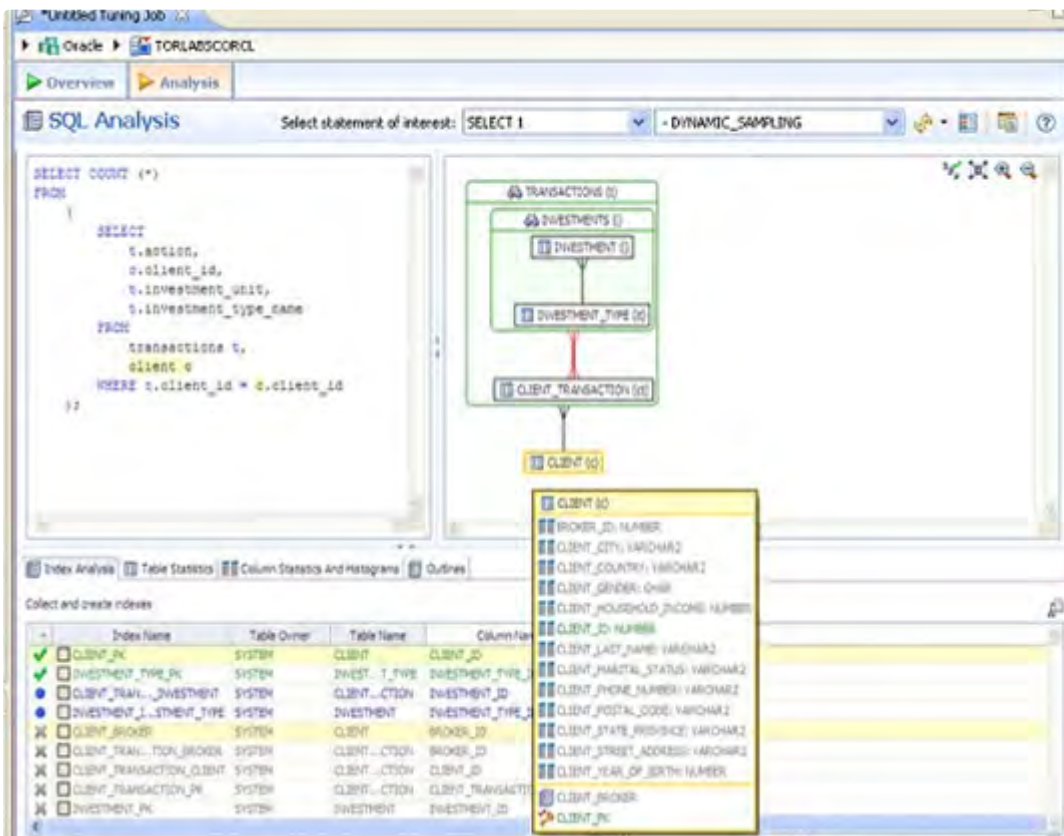
Now we can see the objects in the view:



You can further expand the sub-view within the original view also.


The following is an example of view expansion along with the Explain Plan to the left.


Notice in the view expansion a list of all the indexes on all the underlying tables in the views and sub views and which of those indexes is used in the default execution plan.



Viewing the Oracle explain plan overlay

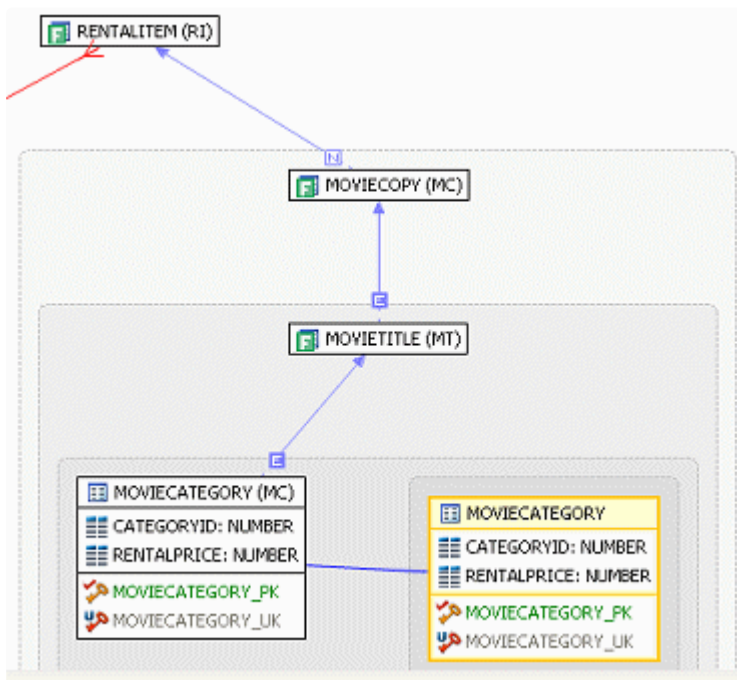
i This option is only available for Oracle versions 10 and higher.

Click the Explain Plan toggle [], you can choose to view or hide the Explain Plan details. The additional nodes shown in the Explain Plan overlay provide details on the flow of the query plan, with operations (such as nested loops, sorts, and joins) showing connecting tables and other operations. Hover the mouse over the objects or relationships to view additional details.



Expanding subqueries and nested subqueries

Double-click queries to expand them or right-click the query and select **Expand Query** from the menu that appears. The following shows several layers of nesting queries.



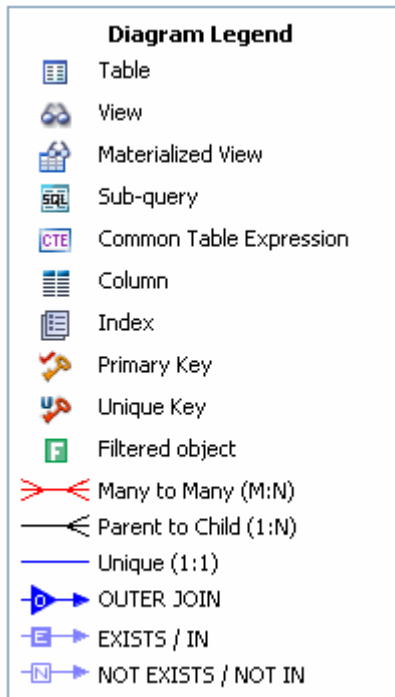
Interpreting the VST diagram graphics

This section contains the following topics that will help you understand the graphics in VST diagrams:

- [Viewing the Diagram Legend](#)
- [Colors](#)
- [Connecting lines/joins](#)
- [Viewing object SQL](#)
- [Refreshing tuning statements](#)
- [Refreshing the VST diagram](#)





Viewing the Diagram Legend

Click the **Diagram Legend Toggle** button as shown in the diagram below to see a description of the icons and relationship lines used in VST diagrams.




Colors









The color of the index entries in the **Collect and Create Indexes** table is interpreted as follows:

| Text Color | Interpretation |
|---|---|
|  | Index is used in the query. |
|  | Index is usable but not used by the current execution path. |
|  | This index is missing. SQL Query Tuner recommends that you create this index. |
|  | This index exists on the table but not usable in this query as it is written. |

Connecting lines/joins

Joins are represented with connecting lines between nodes. You can move tables in the diagram by clicking and dragging them to the desired location. The position of the connecting lines is automatically adjusted. The following describes when a particular type of connecting line is used and the default positioning of the line.

| Connecting Lines | When Used |
|---|---|
|  | One-to-One join relationships are graphed horizontally using blue lines. For more information, see One-to-One Join . |

| Connecting Lines | When Used |
|---|--|
|  | One-to-Many join relationships are graphed with the many table above the one table. For more information, see One-to-Many Join . |
|  | Cartesian Join shows the table highlighted in red with no connectors to indicate that it is joined in via a Cartesian join. For more information, see Cartesian Join . |
|  | Many-to-Many Join relationships are connected by a red line and the relative location is not restricted. For more information, see Many-to-Many Join . |
|  | Indirect Relationship. For more information, see Indirect Relationship . |
|  | Outer Join. For more information, see Outer Join . |
|  | Unique. For more information, see Unique . |
|  | Not Exists and Not in relationship lines connect the subquery to the table being queried. Notice that when you click this relationship line, the SQL text creating the relationship is also selected. For more information, see Not In or Not Exists Join . |
|  | Exists and In relationship lines connect the subquery to the table being queried. Notice that when you click this relationship line, the SQL text creating the relationship is also selected. For more information, see In or Exists Join . |


One-to-one join

If two tables are joined on their primary key, then graphically, these would be laid out side-by-side, with a one-to-one connector:



One-to-many join

This is the default positioning of a one-to-many relationship, where INVESTMENT_TYPE is the master table and INVESTMENT is the details table.




The following is an example of a query that consists of only many-to-one joins, which is more typical:

```

SELECT
    ct.action, c.client_id,
    i.investment_unit,
    it.investment_type_name
FROM
    client_transaction ct,
    client c,
    investment_type it,
    investment i
WHERE
    ct.client_id = c.client_id AND
    ct.investment_id = i.investment_id AND
    i.investment_type_id = it.investment_type_id and
    client_transaction_id=1

```



Cartesian join

A Cartesian join is described in the following example where the query is missing join criteria on the table INVESTMENT:


```

SELECT
    A.BROKER_ID BROKER_ID, A.BROKER_LAST_NAME
    BROKER_LAST_NAME, A.BROKER_FIRST_NAME
    BROKER_FIRST_NAME, A.YEARS_WITH_FIRM
    YEARS_WITH_FIRM, C.OFFICE_NAME OFFICE_NAME,
    SUM (B.BROKER_COMMISSION) TOTAL_COMMISSIONS
FROM
    BROKER A,
    CLIENT_TRANSACTION B,
    OFFICE_LOCATION C,
    INVESTMENT I
WHERE
    A.BROKER_ID = B.BROKER_ID AND
    A.OFFICE_LOCATION_ID = C.OFFICE_LOCATION_ID
GROUP BY

```

```
A.BROKER_ID,
A.BROKER_LAST_NAME,
A.BROKER_FIRST_NAME,
A.YEARS_WITH_FIRM,
C.OFFICE_NAME;
```

Graphically, this looks like:



INVESTMENT is highlighted in red with no connectors to indicate that it is joined in via a Cartesian join.

Possible missing join conditions are displayed in the **Overview** tab under **Generated Cases** in the transformations area. DB Optimize recommends that you create these joins.


| SQL Statements and Cases | | Cost | Execution Statistics | Other Execution Statistics | | |
|--|---------------------|---------|----------------------|----------------------------|---------------|-------|
| Name | Text | Value | Elapsed Time (s) | Physical Reads | Logical Reads | CPU T |
| SELECT 1 | select from BROKER, | 34014.0 | 6.22 | 0 | 170 | |
| [Missing a valid join criteria] transformation | | 274.0 | 0.04 | 0 | 167 | |
| ----- FULL | | 34019.0 | 6.29 | 0 | 173 | |
| ----- LEADING1 | | 34017.0 | 6.25 | 0 | 192 | |
| ----- ALL_ROWS | | 34014.0 | 6.35 | 0 | 170 | |
| ----- LEADING3 | | 34017.0 | 6.41 | 0 | 170 | |
| ----- INDEX | | 34392.0 | 6.58 | 0 | 414 | |
| ----- LEADING2 | | 38148.0 | 7.94 | 0 | 170 | |
| ----- ORDERED | | 38147.0 | 8.61 | 0 | 170 | |
| ----- USE_NL | | 38198.0 | 9.03 | 0 | 37518 | |

i Transformations are highlighted in yellow.

Implied Cartesian join

If there are different details for a master without other criteria then a Cartesian-type join is created:

```
SELECT *
FROM
  investment i,
  broker b,
  client c
WHERE
  b.manager_id=c.client_id and
  i.investment_type_id=c.client_id;
```



The result set of BROKER to CLIENT will be multiplied by the result set of INVESTMENT to CLIENT.


Many-to-many join

If there is no unique index at either end of a join then it can be assumed that in some or all cases the join is many-to-many; there are no constraints preventing a many-to-many join. For example, examine the following query:

```

SELECT *
FROM
  client_transaction ct, client c
WHERE
  ct.transaction_status=c.client_marital_status;
  
```

There is no unique index on either of the fields being joined so the optimizer assumes this is a many-to-many join and the relationship is displayed graphically as:



If one of the fields is unique, then the index should be declared as such to help the optimizer.

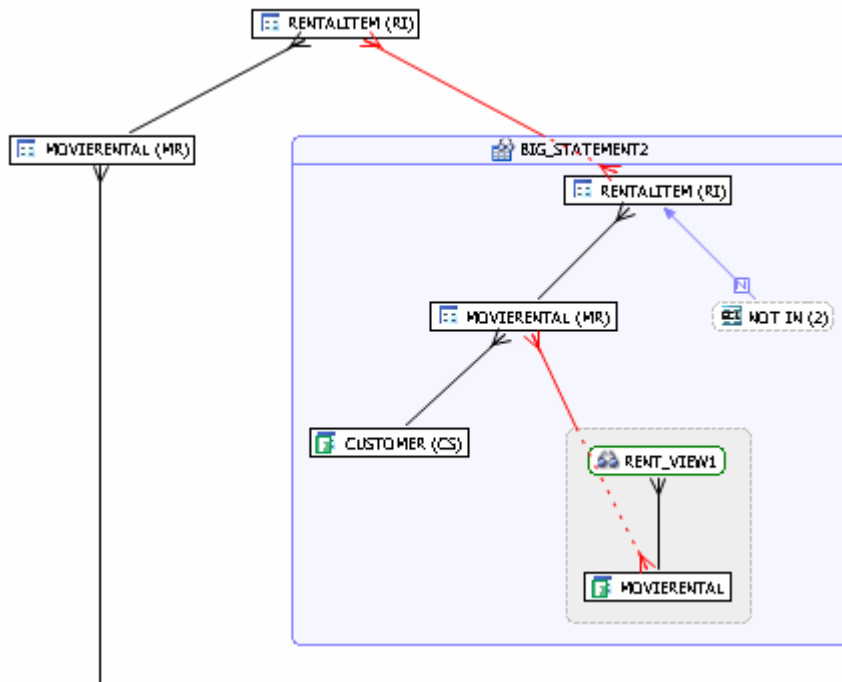
Indirect relationship

Indirect relationships are produced by the following SQL, where BIG_STATEMENT2 is a Materialized View.

```

SELECT CS.*
FROM
  MOVIES.CUSTOMER CS,
  MOVIES.MOVIERENTAL MR,
  MOVIES.RENTALITEM RI,
  OE.BIG_STATEMENT2
WHERE
  CS.ZIP > '75062' AND MR.RENTALID = RI.RENTALID
  AND RI.ITEMNUMBER = OE.BIG_STATEMENT2.ITEMNUMBER
  AND MR.CUSTOMERID = CS.CUSTOMERID;
  
```

The following diagram produced by the SQL above shows that an indirect relationship exists between the RENTALITEM(RI) tables inside and outside the materialized view, BIG_STATEMENT2. An indirect relationship also exists between MOVIERENTAL (MR) inside BIG_STATEMENT2 and MOVIERENTAL(MR) inside the RENT_VIEW1 view.



In or exists join

The following SQL contains a nest IN subquery (shown in bold text) that is graphically represented with the Subquery summary icon and the IN join.

```

SELECT
    cs.customerid, cs.firstname, cs.lastname,
    mr.rentalid, mr.duedate, mr.totalcharge, ri.itemnumber
FROM
    (
        SELECT
            c1.customerid, c1.firstname, c1.lastname, c1.phone
        FROM
            MOVIES.customer c1
        WHERE
            EXISTS (SELECT NULL) cs, (FROM MOVIES.customer c2
        WHERE
            c1.customerid <>c2.customerid AND c1.lastname=c2.lastname AND c1.phone
            BETWEEN 0 AND 9999569900)
        SELECT
            customerid, rentalid, duedate, totalcharge, rentaldate
        FROM
            MOVIES.movierental
        WHERE
    
```

```

totalcharge > 10
)
mr, MOVIES.rentalitem ri
WHERE
LENGTH (cs.lastname) = 10
AND 1 < cs.customerid
AND ROUND (ri.rentalid) > 10
AND TRUNC (ri.itemnumber) > 1
AND mr.totalcharge > (SELECT AVG (totalcharge)
FROM
MOVIES.movierental
WHERE
TOTALCHARGE >= 40)
AND ri.moviecopyid
NOT IN (SELECT mc.moviecopyid FROM MOVIES.moviecopy mc
WHERE
mc.copyformat = 'vhs'
AND mc.copycondition = 'new'
AND mc.movieid IN (SELECT mt.movieid FROM MOVIES.movietitle mt
WHERE
mt.year < 1990
AND mt.rating IN ('pg','r')
AND mt.categoryid IN (SELECT mc.categoryid
FROM
MOVIES.moviecategory mc
WHERE
mc.rentalprice=(SELECT MAX (rentalprice)
FROM
MOVIES.moviecategory
WHERE
categoryid=mc.categoryid)))
AND mr.CUSTOMERID=cs.CUSTOMERID
AND ri.RENTALID=mr.RENTALID

```

Outer join

The bold SQL predicate in the statement below defines the outer join between customer and movierental.

```

select cs.*
from MOVIES.customercs, MOVIES.movierentalmr
where

```

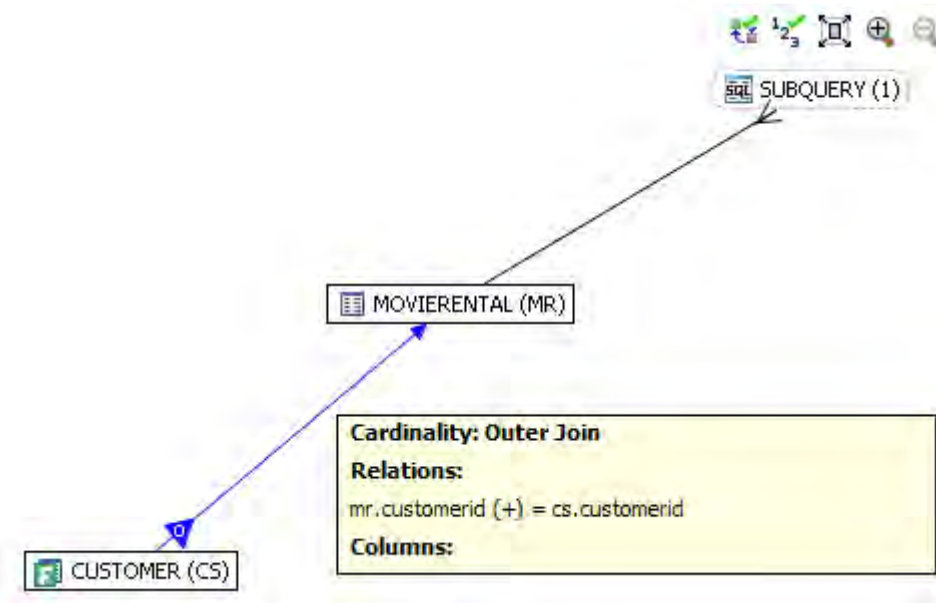
```
length (cs.lastname) = 8 and cs.zip > 75062 and 1 < cs.customerid + 2 and cs.phone between 9625569900
and 9999569900 and mr.rentalid = (select max (ri.rentalid)
```

```
from MOVIES.rentalitem ri, MOVIES.moviecopy mc
```

```
where
```

```
ri.itemnumber > 1 and mc.moviecopyid = 700) and mr.customerid(+)=cs.customerid;
```

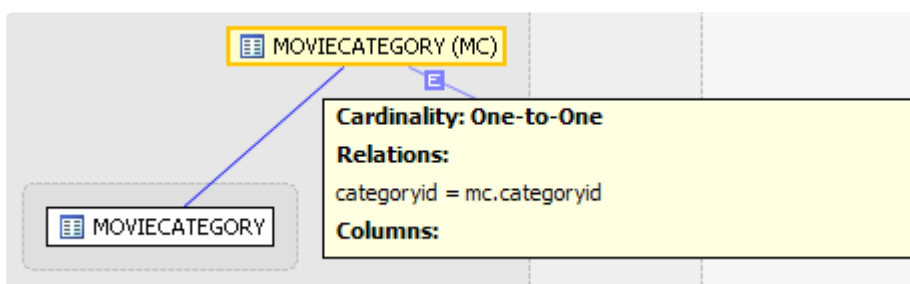
The following screen shot illustrates how the outer join is displayed in the VST diagram.



Unique

The subquery below illustrates a unique relationship between two primary keys.

```
...select max(rentalprice) from MOVIES.moviecategory where categoryid = mc.categoryid...
```



Not in or not exists join

The following SQL contains a NOT IN subquery (shown in bold below) that is graphically represented with the Subquery summary icon and the NOT IN join.

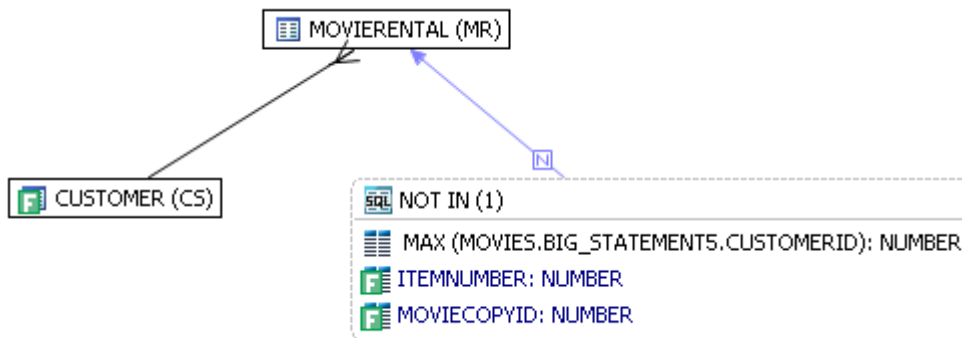
```
SELECT CS.*
FROM
```

```

MOVIES.CUSTOMERCS,
MOVIES.MOVIERENTALMR
WHERE
  CS.ZIP > '75062'
  AND MR.RENTALID NOT IN (SELECT MAX ( MOVIES.BIG_STATEMENTS5.CUSTOMERID)
FROM
  MOVIES.RENTALITEMRI,
  MOVIES.MOVIECOPYMC,
  MOVIES.BIG_STATEMENTS5
WHERE
  RI.ITEMNUMBER>1
  AND MC.MOVIECOPYID=700)
  AND MR.CUSTOMERID=CS.CUSTOMERID;


```

Graphically, this statement would look like this:



Viewing object SQL

Hover over the name of an object to view the object SQL as shown in the diagram below.

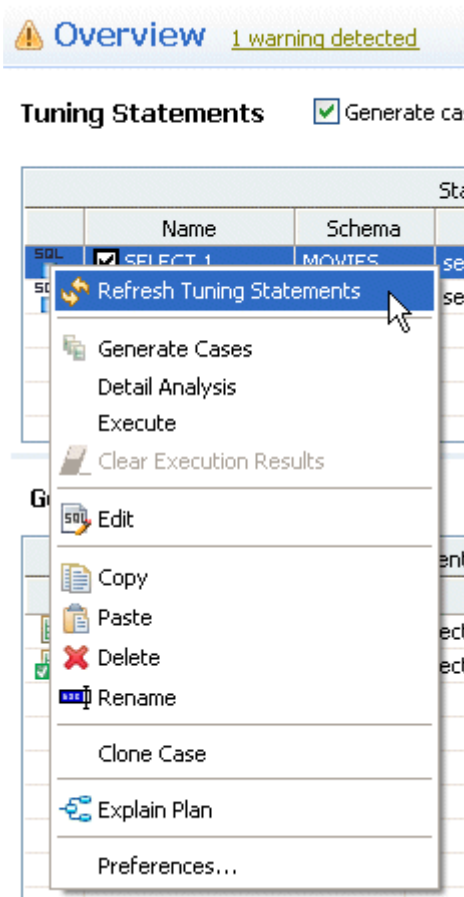


Refreshing tuning statements

At times you may see an error on the Overview page, which when you mouse over it, indicates that the tuning statements are out of sync and need to be refreshed. This can happen, for example, if you tune a statement, then delete it, and insert another SQL query for tuning.


To refresh the tuning statements

In the Tuning Statements area of the Overview tab, right-click the tuning statement and select **Refresh Tuning Statements**.



Refreshing the VST diagram

There are two refresh options available: **Refresh** and **Refresh All**. Click the **Refresh** list as shown below to gain access to these options.



- **Refresh:** Regenerates the Analysis tab including the VST diagram. Any changes made on the tab are reflected in the diagram.
- **Full Refresh:** Re-caches all objects used in (or related to) the query, then regenerates the Analysis tab including the VST diagram. This option is typically used when the underlying objects have been recently changed.

Using platform-specific features

This section describes the tuning features specific to individual platforms.

- [Using the Table Statistics tab \(Oracle and SQL Server\)](#)
- [Using the Column Statistics and Histograms tab \(Oracle and SQL Server\)](#)
- [Using the Outlines tab \(Oracle\)](#)
- [Using the Plan Guides tab \(SQL Server\)](#)




- [Tuning SQL statements in the System Global Area \(Oracle\)](#)


Using the Table Statistics tab (Oracle and SQL Server)

The Table Statistics area of the Analysis tab indicates when and if table statistics were last taken. Using the Table Statistics you can view the information the optimizer uses to choose a path and assess the validity of the various hints presented on the Overview tab.


| Object | | Statistics | | Monitoring | Attributes |
|---------------------------------|--------------------|-------------------|------------------------|------------|------------|
| Table Owner | Table Name | Statistics Status | Days Since Stats Taken | Monitoring | Cache |
| <input type="checkbox"/> SYSTEM | PS_RETROPAYPGM_TBL | Statistics OK | 200 YES | YES | N |
| <input type="checkbox"/> SYSTEM | PS_PAY_CALENDAR | Statistics OK | 200 YES | YES | N |
| <input type="checkbox"/> SYSTEM | WB_JOB | Statistics OK | 200 YES | YES | N |
| <input type="checkbox"/> SYSTEM | WB_RETROPAY_EARN5 | Statistics OK | 200 YES | YES | N |
| <input type="checkbox"/> SYSTEM | PS_RETROPAY_DICT | Statistics OK | 200 YES | YES | N |

This table draws attention to:

- **Missing statistics:** Missing statistics can cause the optimizer to choose the wrong path because the optimizer uses table statistics to make decisions. If the statistics are missing, you can click the select a table and then click **Collect Statistics** () on the far right of the tab. This sends a request to the database to analyze the table and calculate the statistics.
- **Out-of-date statistics:** Like missing statistics, out-of-date statistics can also cause the optimizer to choose the wrong path. You can update the statistics by selecting a table, and then clicking **Display Statistics** (), which refreshes the statistics from the database or by clicking **Collect Statistics** (), which requests the database to analyze the table and calculate the statistics.

 Collecting Statistics may be time-consuming, depending on how many tables the database is analyzing and the number of rows in each table.

- **Useful statistics:** The number of rows in a table and whether the table has been modified since the statistics were last collected can help you to determine which hints you should implement in the SQL code. These statistics can help the DB Administrator to better understand the database.

 You can right-click anywhere in a row and choose options such as Collect Statistics, Display Statistics, and Copy from the short-cut menu.

Using the Column Statistics and Histograms tab (Oracle and SQL Server)

Histograms are special statistics that exist for a limited number of columns and are created by the database administrator. Column histograms should be created only when there are highly-skewed values in a column, such as is the case of an order details table with an Order Status column where the number of closed orders for a business operating for several years is far greater than the number of open orders. The Order Status column therefore meets the criteria of a useful target for a histogram because the data is highly skewed. Using histograms the optimizer determines that a full-scan is recommended when searching for closed orders, but an index scan is more useful when searching for open orders.

Oracle-specific Column Statistics and Histograms tab example

DB Optimizer looks at the columns that have histograms and using statistics tries to determine whether the column is a good or bad candidate for a histogram and presents this information on the Column Statistics and Histograms tab.

| Object | | Histograms | | | | | | | |
|--------------------------|--------------------------|-----------------|-----------------|-----------|-----------|------------------|---------|------------------------|------------------|
| Gather | Drop | Table Name | Column Name | Histogram | # Buckets | Filter Type | Indexed | Median Value Deviation | # Distinct Value |
| <input type="checkbox"/> | <input type="checkbox"/> | PS_PAV_CALENDAR | PAY_CONFIRM_RUN | NONE | 1 | Literal Equality | Yes | | |
| <input type="checkbox"/> | <input type="checkbox"/> | PS_RETROPAYQST | RETROPAY_SEQ_NO | NONE | 1 | Join Equality | Yes | | 1 |
| <input type="checkbox"/> | <input type="checkbox"/> | PS_RETROPAYQST | OFF_CYCLE | NONE | 1 | Join Equality | No | | |
| <input type="checkbox"/> | <input type="checkbox"/> | PS_PAV_CALENDAR | PAYGROUP | NONE | 1 | Join Equality | Yes | | |
| <input type="checkbox"/> | <input type="checkbox"/> | PS_RETROPAYQST | RETROPAY_SEQ_ID | NONE | 1 | Join Equality | Yes | | |

The row shading indicates the following:

- **Green:** Good histogram candidate
- **Red:** Bad histogram candidate
- **No shading:** Not determined to be a good or bad histogram candidate

Median value deviation

For columns that have histograms, the median value deviation is presented. Understanding the median value deviation can help you determine whether an index scan or a full-table scan would be more efficient.

The median value deviation represents the number of values that have duplicates away from the median. In the case of the Order Status column, there are only three possible values, open, processing, and closed. Consider the following:

- 10 open orders
- 100,000 closed orders
- 1 order in processing

In this case the median is the middle value, 10. The number of closed orders is 10,000 times the median which indicates that the column data is highly skewed. In this case the value in the Median Value Deviation column would be presented as:

1, 0, 0, 0, 1, 0, 0, 0

There are 1's at the first and 5th spot in the median value deviation field indicating one column value (value of orders in the *processing* state which appears once) is 1 factor of 10 away from the median and there is a 1 at the 5th position indicating there is a column value (orders in the *closed* state) that appears 5 factors of 10 more often (10,000) than the median value of 10.

A column with a median value deviation of 0, 0, 0, 0, 0, 0, 0, 0 indicates that the column data is not skewed and it is a bad candidate for a histogram, and therefore a full scan of the table would more efficiently satisfy a query than an index scan.

To update the statistics of any object, you can select **Gather** for that column and then click **Display Statistics** or **Collect Statistics**.

To stop gathering statistics for an object, such as a bad candidate for a histogram, select **Drop** for that column and then click **Display Statistics** or **Collect Statistics**.

i If you are gathering statistics for a column for which the statistics were missing or out-of-date, then once the statistics collection is complete, you should return to the Overview tab and rerun the cases, because the characteristics of the column may have changed, so the hints to improve performance would also change.

Using the Outlines tab (Oracle)

The Outlines tab provides detailed information about outlines created by the query during the statement execution process on the **Overview** tab.

It provides information including the SQL statement name, whether the outline is enabled, and the Name, Category, and Hints associated with the outline. Additionally, the Drop parameter specifies whether it is dropped at execution time.

| Enabled | Drop | Name | Category | Hints |
|-------------------------------------|--------------------------|---------------|----------|------------------------------|
| <input type="checkbox"/> | <input type="checkbox"/> | SYS_O...20684 | DEFAULT | FULL(@\"SEL\$2\" |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> | FULL | TEST | SWAP_JOIN_INPUTS(@\"SEL\$1\" |

In order to view outlines, the session needs to have

USE_STO

RED_OUTL

```
INES=tru
```

```
e
```

set prior to execution. Outlines in tuning are created for the DEFAULT category, by default.

Use the following commands to enable outlines with the default settings:

alter

system

set

USE_STO

RED_OUT

ILNES=t

```
rule;
```

```
alter
```

```
system
```

set

USE_STO

RED_OUT

LINES= ‘

DEFAULT

’ ;

alter

session

set

USE_STO

RED_OUT

```
LINES=t
```

```
ruce;
```

Additionally, in order for a session to

```
USE_STO
```

RED_OUTL

INES

, the user requires the create any outline role. Use

the following command to set up the proper permissions:

grant

create

any

```
outline
```

```
to
```

```
[user];
```

Using the Plan Guides tab (SQL Server)

This tab displays any existing plan guides on the server for the statements being tuned. This can be useful when trying to determine whether a plan guide is, or should be, used for a query, or conversely, finding plan guides that have become obsolete.

Index Analysis | Table Statistics | Column Statistics And Histograms | Plan Guides

View Plan Guides

| Enabled | Drop | Name | Statement | Type | Module / Batch | Parameters | Hints |
|-------------------------------------|--------------------------|------------------|---------------------|------|----------------|------------|-------|
| <input checked="" type="checkbox"/> | <input type="checkbox"/> | TUNER_PLAN_GUIDE | SELECT ...[ShiftID] | SQL | | | |

In addition to the **Enabled** and **Drop** controls, each plan guide entry on the tab shows columns corresponding to the

sp_crea

te_plan_

guide


call arguments the plan guide was

created with.

| Column | Description |
|-----------|--|
| Name | The name given the plan guide when it was created. |
| Statement | The text of the Transact-SQL statement |
| Type | <p>The type of entity in which the statement occurs:</p> <ul style="list-style-type: none"> • OBJECT. In the context of a stored procedure, scalar function, multi-statement table-valued function, or DML trigger in the current database. • SQL. In the context of a stand-alone statement or batch. • TEMPLATE. Specifies that this plan guide applies to any query that parameterizes to the specified statement form. Only the <div data-bbox="325 983 1370 1323" style="text-align: center; background-color: #f0f0f0; padding: 20px; margin: 20px 0;"> <h1>PARAMETERIZED</h1> </div> <div data-bbox="325 1438 1291 1778" style="text-align: center; background-color: #f0f0f0; padding: 20px; margin: 20px 0;"> <h1>STATEMENT</h1> </div> |


| Column | Description |
|--------|---|
| | <p data-bbox="327 488 968 826">TION</p> <p data-bbox="327 943 1291 1281">{ FORC</p> <p data-bbox="327 1397 968 1736">ED </p> |

| Column | Description |
|------------------|--|
| | <div data-bbox="325 483 1289 826" style="text-align: center; background-color: #f0f0f0; padding: 20px;"> <h1>SIMPLE</h1> </div> <div data-bbox="325 943 564 1283" style="display: inline-block; vertical-align: middle; background-color: #f0f0f0; padding: 20px; margin-right: 20px;"> <h1>}</h1> </div> <p data-bbox="571 1182 1193 1211">hint can be specified if a Type of TEMPLATE is specified.</p> |
| Module/ Batch | Specifies either the name of the object in which the statement occurs or the batch text in which the statement appears. |
| Parameters | Specifies the definitions of all parameters in the statement. |
| Hints | <p>Can be one of the following:</p> <ul style="list-style-type: none"> • An OPTION clause containing any valid sequence of query hints • A query plan in XML format to be applied as a hint • NULL |


 For a detailed understanding of the **Parameters** and **Hints** values, particularly when using the edit plan guide functionality described below, see the Microsoft SQL Server documentation.

You can modify any plan guide listed on the **Plan Guides** tab. You can enable or disable the plan guide, drop it, or specify parameters or hints.

To drop a plan guide

- Select the **Drop** check box for that plan guide entry and then click the Apply button () on the far right of the tab.

To enable or disable a plan guide

- Enable by selecting the **Enabled** check box (or deselect **Enabled** to disable) for the plan guide entry and then click the Apply button () on the far right of the tab.

To edit a plan guide

1. Right-click a plan guide entry and select **Edit** to open the **Edit Plan Guide**.
2. In the **Parameters** field, provide a valid `sp_create_plan_guide` stored procedure call

@params

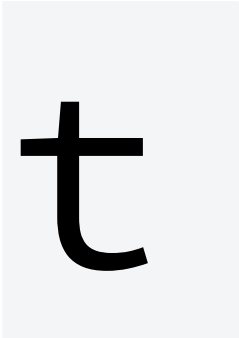
=

argument

t

3. In the **Hints** field, provide a valid `sp_create_plan_guide` stored procedure call

@hints


=
argument
t

4. Click **Next** to preview the generated stored procedure calls.
5. Click **Finish**.

Tuning SQL statements in the System Global Area (Oracle)

On Oracle platforms, SQL statements that reside in the SGA can also be tuned. When you create a tuning job and specify an Oracle source, an additional tab appears in the Tuning Candidates section of tuning, named Active SQL in SGA.

The SGA contains all the SQL since the database has been started up, except for those that have been purged when the system runs out of memory. When analyzing the causes of a database bottleneck, it is perhaps more useful to view and tune the SQL statements most recently run, than those that have run in the last month, for example. DB Optimizer cannot tell you which statements have most recently run by looking in the SGA. However, by profiling the database using DB Optimizer Profiling and then optimizing the code by executing and running the generated cases, you will be able to see which paths are most likely causing a bottleneck and can be altered to enhance performance. Also, you can use IDERA Performance Center to continually monitor a database over a longer period of time to help you analyze and optimize database performance.

To add a statement active in the SGA

1. From the **Overview** tab, click the SQL icon and select **Scan Oracle SGA**. The **Scan SGA** wizard appears.
2. Set the filtering criteria for an SGA scan and then run the wizard. It returns all active statements on the Oracle source.
3. Choose the specific statements and add them to the tuning job.

Additional tuning commands

In addition to tuning, the interface provides additional commands and functionality that enables you to view source code, statements, and other information regarding the data source.

- [View the source code of tuning candidates](#)
- [View statement or case code in SQL Viewer](#)
- [Open an explain plan for a statement or case](#)
- [Executing a session from the command line](#)
- [Saving a tuning job](#)

View the source code of tuning candidates

You can view the source code of a tuning candidate as follows:

- On the **Ad hoc SQL** tab of the **Input** tab, you can see the SQL statements you typed or pasted into that tab.
- On the **Database objects**, **SQL Files**, and **Active SQL in SGA** tabs of the **Input** tab, you can double-click the name of any object added to that tab and an SQL session will open that displays the SQL of that database Object. The SQL editor in use is actually Rapid SQL, an IDERA product that is integrated with DB Optimizer.

View statement or case code in SQL Viewer

The Tuning job's **Overview** tab let you open a statement in an SQL Viewer if you want to perform either of the following tasks:

- View the entire SQL statement.
- Set bind variables. If the Tuning Status Indicator indicates a statement or case has invalid bind variables, you must set those variables before executing the statement or case.

To view or set bind variables in a statement or case

1. Right-click a statement or case and select **Edit**.
2. Use the **Data Type** and **Value** (or **NULL**) controls to specify the type and value for each bind variable. After setting bind variables, you can execute a case.

i Setting the bind variables in a parent statement sets the bind variables in all generated cases for that statement.

Open an explain plan for a statement or case

Any valid SQL statement added to the **Overview** tab shows a calculated explain plan cost in the **Cost** field of the statement or case record. You can open an explain plan on these statements to view the sequence of operations used to execute the statement and the costs and other explain plan details for each operation.

On Oracle version 9 and higher, DB Optimizer attempts to get the Explain Plan from V\$SQL_PLAN when possible. Otherwise, the Explain Plan is generated by the Oracle EXPLAIN PLAN command.

To initially open an explain plan on a valid SQL statement on the Overview tab

1. Right-click in the **Name** field of any statement record showing a value in the **Cost** field.
2. Select **Explain Plan** from the context menu.

An Explain Plan tab opens below the **Overview** tab.

| Operation | Plan Cost | Cost | Operation Cost | Result | Cardinality | Bytes | CPU Cost | IO Cost | Optimizer | Actual Statistics |
|------------------------|-----------|------|----------------|--------|-------------|--------|-----------|---------|-----------|-------------------|
| SELECT STATEMENT | 253.0 | 0.0 | | | 37 | 2183 | 102...107 | 235 | ALL...WS | |
| PX COORDINATOR | | | | | | | | | | |
| PX SEND - SYS::TQ10004 | 253.0 | 0.0 | | | 37 | 2183 | 102...107 | 235 | | |
| HASH | 253.0 | 0.0 | | | 37 | 2183 | 102...107 | 235 | | |
| PX RECEIVE | 253.0 | 0.0 | | | 37 | 2183 | 102...107 | 235 | | |
| PX SEND - SYS::TQ10003 | 253.0 | 0.0 | | | 37 | 2183 | 102...107 | 235 | | |
| HASH | 253.0 | 1.0 | | | 37 | 2183 | 102...107 | 235 | | |
| HASH JOIN | 252.0 | 1.0 | | | 7981 | 470879 | 95195351 | 235 | | |
| PX RECEIVE | 67.0 | 0.0 | | | 37 | 703 | 24604129 | 63 | | |
| PX SEND -...:TQ10001 | 67.0 | 0.0 | | | 37 | 703 | 24604129 | 63 | | |
| PX BLOCK | 67.0 | 0.0 | | | 37 | 703 | 24604129 | 63 | | |
| TABLE A...JENT11 | 67.0 | 67.0 | | | 37 | 703 | 24604129 | 63 | ANA...ED | |
| HASH JOIN | 184.0 | 1.0 | | | 11706 | 468240 | 67387172 | 173 | | |

Explain plan operations are shown in a typical tree structure showing parent-child relationships. The following table describes the column groups shown for each operation on the **Explain Plan** tab.

With the **Explain Plan** tab open, you can quickly switch the view to an explain plan for another SQL statement.

To change to another SQL statement

- Click in the **Name** field of another statement record showing a value in the **Cost** field.

Executing a session from the command line

You can launch a tuning job from the command line using the following syntax:

```
dboptim
```

```
izer.exe
```

e tune

ds : ROML

ABORCL1

OG_1

sqlfile

: C :

\dbopti

mizer\w

orkspac

```
e\test.
```

```
sql
```

In the above command, the user has specified

```
ROMLABO
```

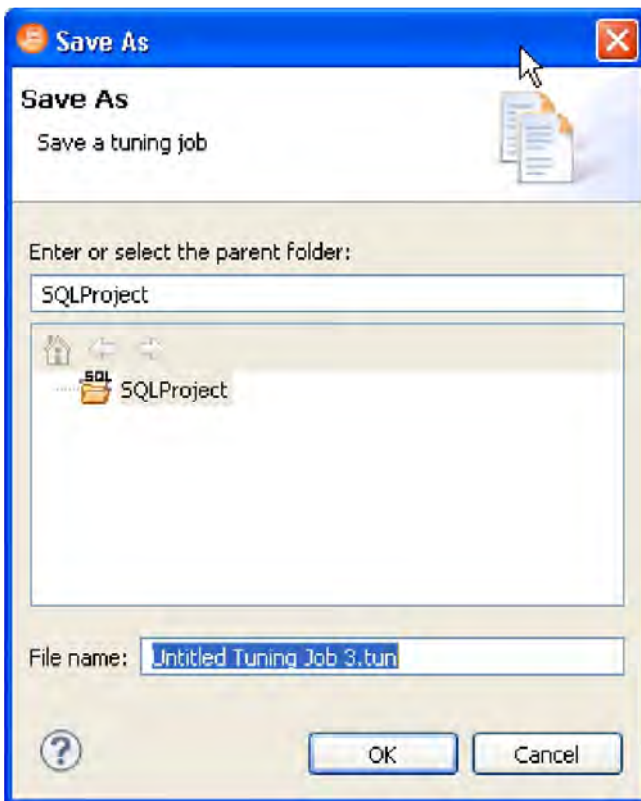
RCL10G_

1

as the data source, and indicates a tuning session using the test.SQL script.

Saving a tuning job

A tuning session can be saved to a file with a .tun suffix. This enables you to open the file at a later time for analysis and to share the tuning job results with other users.



Tuning sessions can be saved as .tun files for use at a later time.

Once you have saved a tuning session to disk as a .tun file, it appears in the SQL Project Explorer under the name you saved it as. It can be opened again by double-clicking the project name.

To save a tuning session

Select the tuning session and then choose **File > Save As....** Specify the project location you want to save the file in and modify the file name, as needed. Click **OK**. The tuning job project is added to SQL Project Explorer.

Configuring tuning

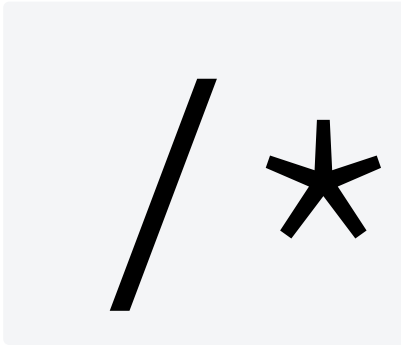
This section contains information on configuring tuning. It provides information on setting up your data sources to work with tuning functionality, as well as information regarding preferences within the application for the customization of various features and functionality.

This section is comprised of the following topics:

- [Set roles and permissions on data sources](#)
- [Specify Tuning Job Editor preference](#)
- [Specify Case Generation preferences](#)
- [Specify VST Diagrams Tuning preference](#)

Set roles and permissions on data sources

In order to take advantage of all tuning features, each user must have a specific set of permissions. The code below creates a role with all required permissions. To create the required role, execute the SQL against the target data source, modified according to the specific needs of your site:



Create

the

```
role */
```

```
CREATE
```

```
ROLE
```

SQLTUNI

NG NOT

IDENTIF

IED

/ GRANT

SQLTUNI

NG TO

" CONN EC

T"

/ GRANT

SQLTUNI

NG TO

SELECT_

CATALOG

_ROLE

/ GRANT

ANALYZE

ANY TO

SQLTUNI

NG

/ GRANT

CREATE

ANY

OUTLINE

TO

SQLTUNI

NG

/ GRANT

CREATE

ANY

PROCEDU

RE TO

SQLTUNI

NG

/ GRANT

CREATE

ANY

TABLE

TO

SQLTUNI

NG

/ GRANT

CREATE

ANY

TRIGGER

TO

SQLTUNI

NG

/ GRANT

CREATE

ANY

VIEW TO

SQLTUNI

NG

/ GRANT

CREATE

PROCEDU

RE TO

SQLTUNI

NG

/ GRANT

CREATE

SESSION

TO

SQLTUNI

NG

/ GRANT

CREATE

TRIGGER

TO

SQLTUNI

NG

/ GRANT

CREATE

VIEW TO

SQLTUNI

NG

/ GRANT

DROP

ANY

OUTLINE

TO

SQLTUNI

NG

/ GRANT

DROP

ANY

PROCEDU

RE TO

SQLTUNI

NG

/ GRANT

DROP

ANY

TRIGGER

TO

SQLTUNI

NG

/ GRANT

DROP

ANY

VIEW TO

SQLTUNI

NG

/ GRANT

SELECT

ON

SYS.V_

SESSION

TO

SQLTUNI

NG

/ GRANT

SELECT

ON

SYS.V_\$

SESSTAT

TO

SQLTUNI

NG

/ GRANT

SELECT

ON

SYS.V_\$

SQL TO

SQLTUNI

NG

/ GRANT

SELECT

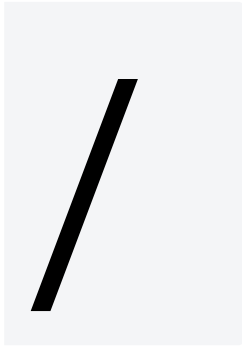
ON

SYS.V_\$


STATNAM

E TO

SQLTUNI

A light gray square containing the letters "NG" in a large, bold, black sans-serif font.A light gray square containing a single, thick, black diagonal slash (/).

Once complete, you can assign the role to users who will be running tuning jobs:



Create

a

sample

user* /

CREATE

USER

TUNINGU

SER

IDENTIF


IED BY

VALUES

' 05FFD2

6E95CF4

A4B'




DEFAULT

TABLESP

ACE

USERS



TEMPORA

RY

TABLESP

ACE


TEMP

QUOTA

UNLIMIT

ED ON


USERS



PROFILE

DEFAULT

ACCOUNT



UNLOCK

/ GRANT

SQLTUNI

NG TO

TUNINGU

SER

/ ALTER

USER


TUNINGU

SER

DEFAULT

ROLE

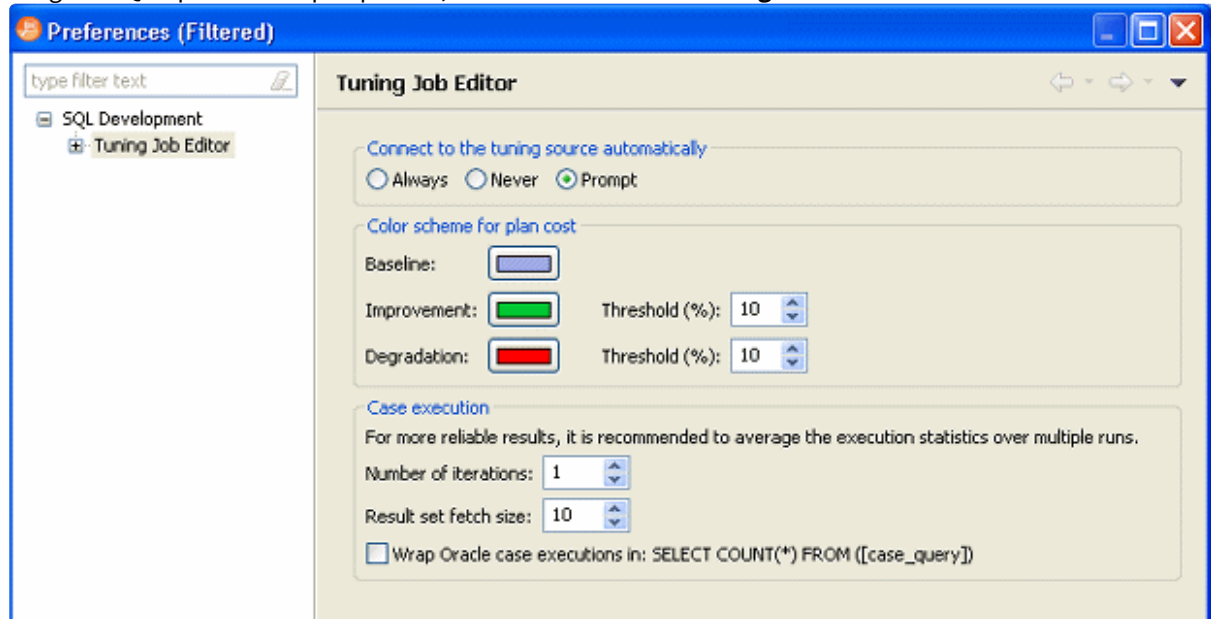
SQLTUNI



Specify Tuning Job Editor preference

Tuning job editor preferences let you control certain aspects of the appearance of items in the tuning job editor as well as default behaviors.

- Using the SQL Optimization perspective, select **Preferences > Tuning Job Editor**.



- Make your changes and then to save your changes, click **Apply**.

The following describes the options available:

General Preferences

Connect to the tuning source automatically: When you open a tuning perspective, it automatically opens the last saved tuning jobs that were open when you closed the application. This option lets you specify whether, in addition, you want to automatically connect to the data sources associated with these tuning jobs. If you typically review existing tuning job archives rather than run new tuning jobs, you may wish to explicitly connect to a data source rather than connect automatically. The options are:

- **Always:** Automatically connects to data sources associated with tuning jobs that were open last time you shut down tuning.
- **Never:** Automatically opens tuning job archives that were open last time you shut down the application but does not automatically connect to the associated data sources.
- **Prompt:** Prompts you to connect to data sources associated with tuning jobs that were open last time you shut down the application.

Color scheme for plan cost: In the graphical representations of explain plan cost and elapsed time, tuning uses a color scheme to highlight differences among generated cases. Values for the original statement are treated as a baseline, and values for individual cases that are within a specified threshold range of the baseline value are represented with a **Baseline** color. For cases whose values are outside the threshold range, **Improvement** and **Degradation** colors are used to represent values in those cases.

i You can set the threshold in the application preferences, by selecting **Window > Preferences > Tuning Job Editor** and then changing the threshold levels as required.

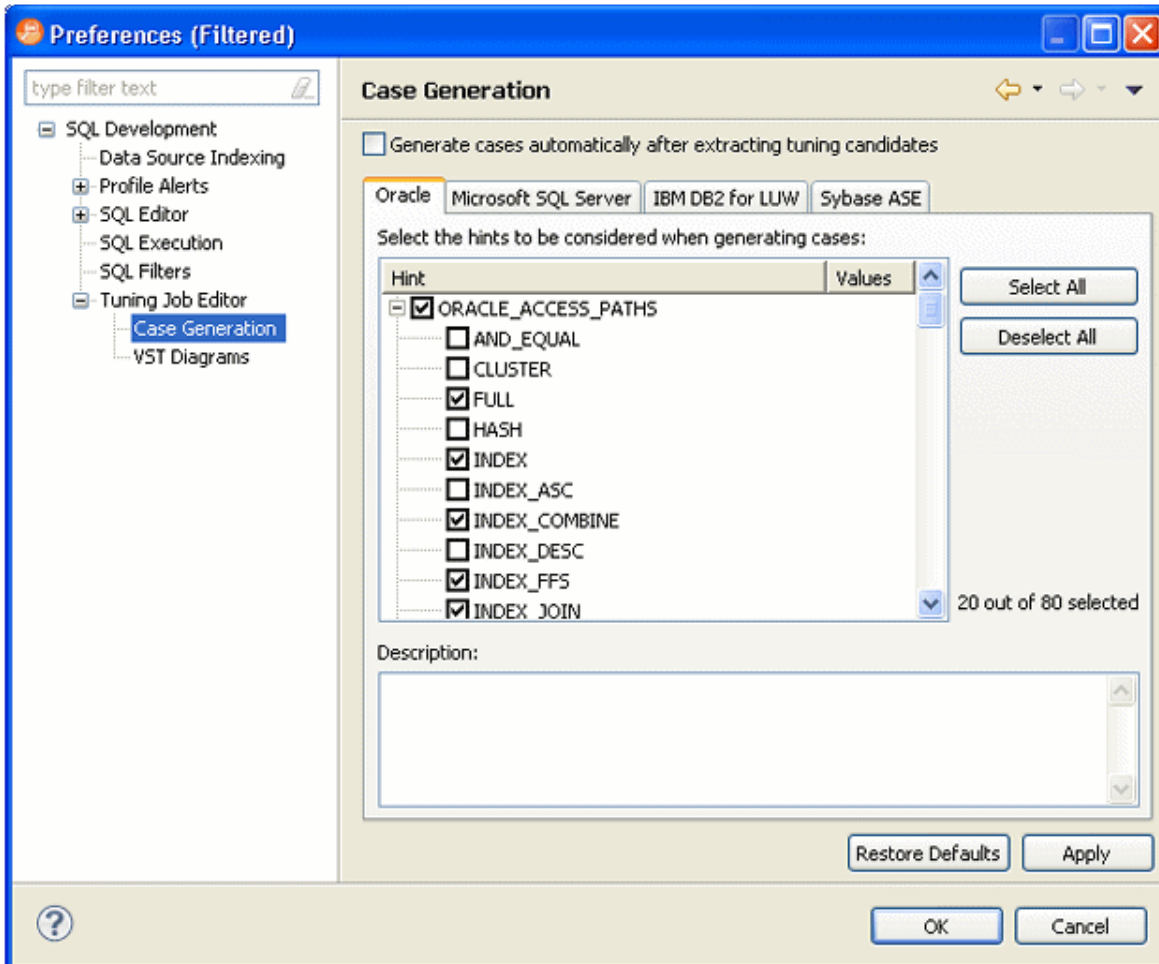
Case execution: Lets you dictate how execution statistics are gathered.

See Also

- [Specify Case Generation preferences](#)
- [Specify VST Diagrams Tuning Preferences](#)

Specify Case Generation preferences

Additionally, the Generated Case preference page lets you enable or disable the automatic generation of SQL Optimizer hint-based cases of SQL statements added to a tuning job. It also lets you indicate which specific hint types are generated when the feature is enabled.



Using the SQL Optimization perspective, select **Preferences > Tuning Job Editor > Case Generation**.

Use the Case Generation option automatically after extracting tuning candidates control to enable or disable automatic generation of hint-based cases, and then select the check boxes to specify the hint-based cases that are generated for a statement added to a tuning job.


About statement records

| Column or Column Set | Description |
|----------------------------|---|
| SQL Statements and Cases | <p>Identifiers for the generated statement or case:</p> <ul style="list-style-type: none"> • Name: Statements are assigned a numbered identifier based on the order in which they were added to a tuning job. • Text: An excerpt of the statement or case based on the statement type. For details on how to view the entire statement or case. |
| Cost | <p>An explain plan-based cost estimate. This field is populated as soon as the statement is added to the Overview tab.</p> <p>This column set can be expanded to display a graphical representation of the cost to facilitate comparisons among cases.</p> |
| Index Analysis | <p>Tuning automatically detects indexes that require optimization and offers you the option to automatically optimize the index. For more information, see Implementing Index Analysis Recommendations.</p> |
| Elapsed Time | <p>The execution time during the most recent execution. This column set is not populated until you execute the statement or case.</p> <p>This column set can be expanded to display a graphical representation of the elapsed time to facilitate comparisons among cases.</p> |
| Other Execution Statistics | <p>Against Oracle datasources, the default, collapsed view has Physical Reads and Logical Reads columns. Expanded, there are also Consistent Gets, Block Gets, and Rows Returned, CPU time(s), Parse CPU Time(s), Row Sorts, Memory Sorts, Disk Sorts, and Open Cursors columns.</p> <p>For details on these statistics, refer to your DBMS documentation.</p> <p>Against DB2 LUW datasources, the default, collapsed view has Pool Physical Reads and Pool Logical Reads columns. Expanded, there are also Direct Reads, CPU Time(s), Sorts, Sort Time, and Sort Overflows columns.</p> <p>This column set is not populated until you execute the statement or case.</p> |

Specify VST Diagrams Tuning preference

The preferences on this page allow you to change the default presentation of count information and sub-query names in Visual SQL Tuning diagrams.

Using the SQL Optimization perspective, select **Preferences > Tuning Job Editor > VST Diagrams**.

- **Show count information.** If enabled, shows the ratios and count information when the VST diagram is generated. If not enabled, you must click the **Ratios/Counts** icon () on the VST diagram to view ratio and count information.

- **Hide sub-query names when expanded.** If enabled, shows the sub-query name when the VST diagram is generated. If not enabled, you must right-click anywhere in the VST and from the menu that appears, choose Show Query Names.

Examples of transformations and SQL query rewrites

Cartesian product elimination

Detects Cartesian Joins and propose corrections based on analysis of statement, for example suggesting dept.deptno = emp.deptno if emp and dept had no join criteria.

Expression transformation

Identifies actions on predicates that might suppress index usage such as "where empid + 1 = 1 ", should be "where empid=0"

Invalid outer join

Identifies invalid outer joins and suggests more efficient alternatives.

| Before | After |
|---|--|
| <pre>SELECT * FROM employee e, customer c WHERE e.employee_id = c.salesperson_id (+) AND c.state = 'CA'</pre> | <pre>SELECT * FROM employee e, customer c WHERE e.employee_id = c.salesperson_id (+) AND c.state(+) = 'CA'</pre> |

Transitivity

| Before | After |
|---|--|
| <pre>SELECT * FROM item i, product p, price pr WHERE i.product_id = p.product_id AND p.product_id = pr.product_id</pre> | <pre>SELECT * FROM item i, product p, price pr WHERE i.product_id = p.product_id AND p.product_id = pr.product_id AND i.product_id = pr.product_id</pre> |

Move expression to WHERE clause

| Before | After |
|---|--|
| <pre>SELECT col_a, SUM(col_b) FROM table_a GROUP BY col_a HAVING col_a > 100</pre> | <pre>SELECT col_a, SUM(col_b) FROM table_a WHERE col_a > 100 GROUP BY col_a</pre> |

NULL column

| Before | After |
|--|--|
| <pre>SELECT * FROM employee WHERE manager_id != NULL</pre> | <pre>SELECT * FROM employee WHERE manager_id IS NULL</pre> |

Push subquery

| Before | After |
|--|---|
| <pre>SELECT * FROM employee WHERE employee_id = (SELECT MAX(salary) FROM employee)</pre> | <pre>SELECT employee.* FROM employee, (SELECT DISTINCT MAX(salary) col1 FROM employee) t1 WHERE employee_id = t1.col1</pre> |

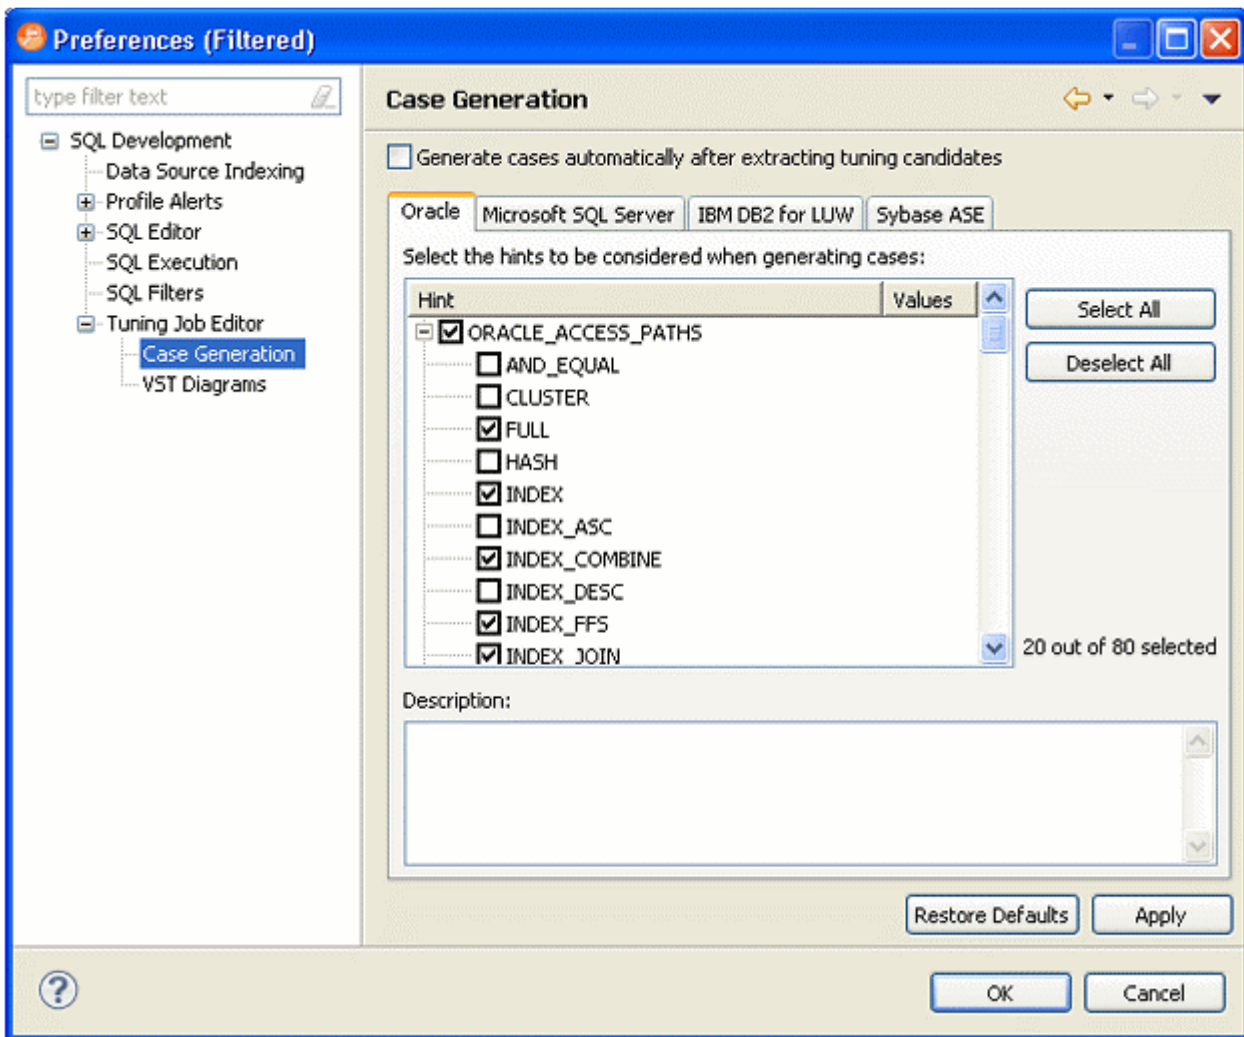
Mismatched column types

Identify joins type mismatch such as number = character which might suppress use of Index.

DBMS hints

Users can provide hints to a specified platform in order to instruct data source optimizer on the best way to execute SQL statements. Tuning automatically generates cases using these hints.

Hints can be enabled or disabled when cases are being generated by tuning on the **Window > Preferences > Tuning Job Editor > Case Generation** pane. Choose a tab as it pertains to the platform you want to modify and use the check boxes to select and de-select the hints you want to enable or disable, respectively.



The following topics describe platform hints that are packaged in tuning to provide optimal efficiency when executing jobs:

- [SQL Server hints](#)
- [DB2 hints](#)
- [Sybase hints](#)

SQL Server hints

The following table highlights SQL hints based on MS SQL Server hints optimization.

| Category | Hint | Available For | Notes |
|----------|---------------------------------------|-----------------------------------|--|
| QUERY | IGNORE_NONCLUSTERED_COLUMNSTORE_INDEX | SELECT/INSERT/UPDATE/DELETE/MERGE | This hint lets you disable use of a nonclustered xVelocity memory optimized columnstore index. |

| Category | Hint | Available For | Notes |
|-----------------|----------------|---------------------------|--|
| JOIN | LOOP | SELECT/UPDATE/ DELETE | Not applicable for RIGHT OUTER or FULL joins. |
| JOIN | HASH | SELECT/UPDATE/ DELETE | -- |
| JOIN | MERGE | SELECT/UPDATE/ DELETE | -- |
| JOIN | REMOTE | SELECT/UPDATE/ DELETE | Only for INNER JOINS. Not applicable with COLLATE. |
| QUERY | RECOMPILE | SELECT/UPDATE/ DELETE | -- |
| QUERY | FORCE ORDER | SELECT/UPDATE/ DELETE | -- |
| QUERY | ROBUST PLAN | SELECT/UPDATE/ DELETE | -- |
| QUERY | KEEP PLAN | SELECT/UPDATE/ DELETE | -- |
| QUERY | KEEPFIXED PLAN | SELECT/UPDATE/ DELETE | -- |
| QUERY | EXPAND VIEWS | DML Statements | Only for statement containing views. |
| QUERY | HASH GROUP | SELECT | Only when GROUP BY, COMPUTE and DISTINCT clauses are used. |
| QUERY | ORDER GROUP | SELECT/UPGRADE/ DELETE | Only when GROUP BY, COMPUTE and DISTINCT clauses are used. |
| QUERY | MERGE UNION | SELECT | Only for statements chained using UNION. |
| QUERY | HASH UNION | SELECT | Only for statements chained using UNION. |

| Category | Hint | Available For | Notes |
|-----------------|--------------------|-------------------------|---|
| QUERY | CONCAT UNION | SELECT | Only for statements chained using UNION. |
| QUERY | LOOP JOIN | SELECT/UPGRADE/DELETE | -- |
| QUERY | MERGE JOIN | SELECT/UPGRADE/DELETE | -- |
| QUERY | HASH JOIN | SELECT/UPGRADE/DELETE | -- |
| Table | FORCESCAN | SELECT/UPGRADE/COMPLETE | Forces the optimizer to use an index scan operation as the access path to the referenced table or view. |
| Table | FORCESEEK | SELECT/UPDATE/COMPLETE | Forces the optimizer to use an index seek operation as the access path to the referenced table or view. |
| Table | INDEX() | DML Statements | Only for tables and views with indexes. |
| Table | KEEPIDENTITY | INSERT | Only for INSERT statements using OPENROWSET clause with BULK option. |
| Table | KEEPDEFAULTS | INSERT | Only for INSERT statements using OPENROWSET clause with BULK option. |
| Table | HOLDLOCK | DML Statements | Not applicable for SELECT statements using FOR BROWSE clause. |
| Table | IGNORE_CONSTRAINTS | INSERT | Only for INSERT statements using OPENROWSET clause with BULK option. |
| Table | IGNORE_TRIGGERS | INSERT | Only for INSERT statements using OPENROWSET clause with BULK option. |
| Table | NOLOCK | SELECT/UPDATE/COMPLETE | Not applicable for the target table in UPDATE/DELETE statements. |
| Table | NOWAIT | DML Statements | -- |

| Category | Hint | Available For | Notes |
|----------|------------------------------|----------------------------|---|
| Table | PAGELock | DML Statements | -- |
| Table | READCOMMITTED | DML Statements | -- |
| Table | READCOMMITTEDLOCK | SELECT/UPDATE/ COMPLETE | -- |
| Table | READPAST | SELECT/UPDATE/ COMPLETE | Not applicable for the target table in UPDATE/ DELETE statements. |
| Table | READUNCOMMITTED | SELECT/UPDATE/ COMPLETE | Not applicable for the target table in UPDATE/ DELETE statements. |
| Table | REPEATABLEREAD | DML Statements | -- |
| Table | ROWLOCK | DML Statements | -- |
| Table | SERIALIZABLE | DML Statements | Not applicable for SELECT statements using FOR BROWSE clause. |
| Table | SPATIAL_WINDOW_MAX_ CELLS | DML Statements | Specifies the maximum number (1 - 8192) of cells to use when tessellating a geometry or geography object. |
| Table | TABLOCK | DML Statements | -- |
| Table | TABLOCKX | DML Statements | -- |
| Table | UPDLOCK | DML Statements | -- |
| Table | XLOCK | DML Statements | -- |

DB2 hints

The following table highlights SQL hints based on IBM DB2 hints optimization.

| Category | Hint | Notes |
|----------|------------------------|---------------------------------------|
| Command | SET OPTIMIZATION LEVEL | For top-level SELECT statements only. |

| Category | Hint | Notes |
|----------|---------------------------|---------------------------------------|
| Clause | optimize for <n> rows | For top-level SELECT statements only. |
| Clause | fetch first <n> rows only | For SELECT statements only. |

Sybase hints

The following table highlights SQL hints based on Sybase hints optimization.

| Category | Hint | Notes |
|----------|-----------|---|
| Logical | distinct | No explicit implementation. |
| Logical | group | No explicit implementation. |
| Logical | g_join | No explicit implementation. |
| Logical | nl_g_join | Not applicable for: statements with chained queries; select statements with group by clause and having clause or group by clause and order by clause. |
| Logical | m_g_join | Not applicable for: statements with chained queries; select statements with group by clause and having clause or group by clause and order by clause. |
| Logical | join | No explicit implementation. |
| Logical | nl_join | Not applicable for: select statements with group by clause and having clause or group by clause and order by clause. |
| Logical | m_join | Not applicable for: select statements with group by clause and having clause or group by clause and order by clause. |
| Logical | h_join | Not applicable for: select statements with group by clause and having clause or group by clause and order by clause. |
| Logical | union | No explicit implementation. |

| Category | Hint | Notes |
|-----------------|------------------|--|
| Logical | scan | No explicit implementation. |
| Logical | scalar_agg | Only used in combination with other operators. It does not change the execution plan itself. |
| Logical | sequence | Is a keyword that will be used in the implementation of scalar_agg operator. |
| Logical | hints | We don't support a combination of hints. |
| Logical | prop | Uses a set of pre-defined values. |
| Logical | table | Used only in combination with other operators, when referring tables from subqueries. |
| Logical | work_t | This operator is applicable only together with store operator. |
| Logical | in | Used only in combination with other operators, when referring tables from subqueries. |
| Logical | subq | Used only in combination with other operators, when referring tables from subqueries. |
| Physical | distinct_sorted | Only for SELECT statements containing DISTINCT, and only for tables. |
| Physical | distinct_sorting | Only for SELECT statements containing DISTINCT, and only for tables. |
| Physical | distinct_hashing | Only for SELECT statements containing DISTINCT, and only for tables. |
| Physical | group_sorted | Only for SELECT statements (not working for views) with no having and no order by clause. |
| Physical | group_hashing | Only for SELECT statements (not working for views) with no having and no order by clause. |

| Category | Hint | Notes | | | | | | | | |
|-----------------------|-------------------------------------|---|-----------------------|-----------|------|----|-----------------------|-----------|-------|--|
| Physical | group_in terting | Not implemented. | | | | | | | | |
| Physical | append_ union_ all | Not applicable for: UNION chained clauses, nested sub-selects in a from clause, if a group by clause is present or if scalar aggregation is present. | | | | | | | | |
| Physical | merge_ union_ all | Not applicable for: UNION ALL chained clauses, nested sub-selects in a from clause, or if a group by clause is present. | | | | | | | | |
| Physical | merge_ union_ dis tinct | Not applicable for: UNION ALL chained clauses, nested sub-selects in a from clause, or if a group by clause is present. | | | | | | | | |
| Physical | hash_ uni on_ dis tinct | Not applicable for: UNION ALL chained clauses, nested sub-selects in a from clause, if a group by clause is present, or if scalar aggregation is present. | | | | | | | | |
| Physical | i_scan | Applied to all table references in the from clause of the main select and of the sub select statements except: 1. statement has sub-selects. 2. table references has no indexes. | | | | | | | | |
| Physical | t_scan | Applied to all the table references in the from clause of the main select and of the sub select statements except: On Sybase 12.5 not applied for tables in the main query if: <ol style="list-style-type: none"> 1. statement has chained queries. 2. Sub queries have group by and having clauses; and not applied to the tables in sub selects if: <ol style="list-style-type: none"> a. has select statements in from clause of the main select. b. sub queries have group by and having clauses. c. statement has select statements in select clause. d. statement has parent statement and insert statement; On Sybase 15 not applied for tables in sub selects if: <ol style="list-style-type: none"> 1. has select statements in from clause of the main select. 2. statement has chained queries. | | | | | | | | |
| Physical | m_scan | Applied for all tables if in the where clause there is a condition like: <table style="margin-left: 40px; border: none;"> <tr> <td style="padding-right: 10px;">table1.indexedColumn1</td> <td>condition</td> <td>body</td> <td>OR</td> </tr> <tr> <td style="padding-right: 10px;">table1.indexedColumn2</td> <td>condition</td> <td>body;</td> <td></td> </tr> </table> Not applied if the LIKE operator is used. For columns that belong to a primary key only the first column is considered. | table1.indexedColumn1 | condition | body | OR | table1.indexedColumn2 | condition | body; | |
| table1.indexedColumn1 | condition | body | OR | | | | | | | |
| table1.indexedColumn2 | condition | body; | | | | | | | | |

| Category | Hint | Notes |
|-----------------|-------------|--------------|
| Physical | store | -- |
| Physical | store_index | -- |
| Physical | sort | -- |
| Physical | xchg | -- |

Reference

The following topics provide reference details:

- [Database objects](#)
- [DBMS connection parameters by platform](#)

Database objects

The following table describes the database objects displayed in DB Optimizer and contains information regarding each one, including object name, DBMS platform, and any notes pertaining to the specified object.

In DB Optimizer, database objects are stored in Data Source Explorer as subnodes of individual, pertinent databases.

| Data base Object | DB MS Plat forms | Notes |
|-------------------------|-------------------------|--|
| Alias es | DB2 | <p>An alias is an alternate name that references a table, view, and other database objects. An alias can also reference another alias as long as the aliases do not reference one another in a circular or repetitive manner.</p> <p>Aliases are used in view or trigger definitions in any SQL statements except for table check-constraint definitions. (The table or view name must be referenced in these cases.)</p> <p>Once defined, an alias is used in query and development statements to provide greater control when specifying the referenced object. Aliases can be defined for objects that do not exist, but the referenced object must exist when a statement containing the alias is compiled.</p> <p>Aliases can be specified for tables, views, existing aliases, or other objects. Create Alias is a command available on the shortcut menu.</p> |
| Check Constraints | All | <p>A check constraint is a search condition applied to a table.</p> <p>When a check constraint is in place, Insert and Update statements issued against the table will only complete if the statements pass the constraint rules.</p> <p>Check constraints are used to enforce data integrity when it cannot be defined by key uniqueness or referential integrity restraints.</p> <p>A check condition is a logical expression that defines valid data values for a column.</p> |
| Clusters | Oracle | <p>A cluster is a collection of interconnected, physical machines used as a single resource for failover, scalability, and availability purposes.</p> <p>Individual machines in the cluster maintain a physical host name, but a cluster host name must be specified to define the collective as a whole.</p> <p>To create a cluster, you need the CREATE CLUSTER or CREATE ANY CLUSTER system privilege.</p> |

| Data base Object | DB MS Plat forms | Notes |
|------------------|------------------|--|
| Data base Links | Oracle | <p>A database link is a network path stored locally, that provides the database with the ability to communicate with a remote database.</p> <p>A database link is composed of the name of the remote database, a communication path to the database, and a user ID and password (if required).</p> <p>Database links cannot be edited or altered. To make changes, drop and re-create.</p> |
| Foreign Keys | All | <p>A foreign key references a primary or unique key of a table (the same table the foreign key is defined on, or another table and is created as a result of an established relationship). Its purpose is to indicate that referential integrity is maintained according to the constraints.</p> <p>The number of columns in a foreign key must be equal to the number of columns in the corresponding primary or unique key. Additionally, the column definitions of the foreign key must have the same data types and lengths.</p> <p>Foreign key names are automatically assigned if one is not specified.</p> |
| Functions | DB2, Oracle | <p>A function is a relationship between a set of input data values and a set of result values.</p> <p>For example, the <code>TIMESTAMP</code> function passes input data values of type <code>DATE</code> and <code>TIME</code>, and the result is <code>TIMESTAMP</code>.</p> <p>Functions can be built-in or user-defined. Built-in functions are provided with the database. They return a single value and are part of the default database schema. User-defined functions extend the capabilities of the database system by adding function definitions (provided by users or third-party vendors) that can be applied in the database engine itself.</p> <p>A function is identified by its schema, a function name, the number of parameters, and the data types of its parameters.</p> <p>Access to functions is controlled through the <code>EXECUTE</code> privilege. <code>GRANT</code> and <code>REVOKE</code> statements are used to specify who can or cannot execute a specific function or set of functions.</p> |
| Groups | All | <p>Groups are units that contain items. Typically, groups contain the result of a single business transaction where several items are involved.</p> <p>For example, a group is the set of articles bought by a customer during a visit to the supermarket.</p> |

| Data base Object | DB MS Plat forms | Notes |
|------------------|------------------|--|
| Indices | All | <p>An index is an ordered set of pointers to rows in a base table.</p> <p>Each index is based on the values of data in one or more table columns. An index is an object that is separate from the data in the table. When an index is created, the database builds and maintains it automatically.</p> <p>Indexes are used to improve performance. In most cases, access to data is faster with an index. Although an index cannot be created for a view, an index created for the table on which a view is based can improve the performance of operations on that view.</p> <p>Indexes are also used to ensure uniqueness. A table with a unique index cannot have rows with identical keys.</p> <p>DB2: Allow Reverse Scans, Percent Free (Lets you type or select the percentage of each index page to leave as free space when building the index, from 0 to 99), Min Pct Used (Lets you type or select the minimum percentage of space used on an index leaf page. If, after a key is removed from an index leaf page, the percentage of space used on the page is at or below integer percent, an attempt is made to merge the remaining keys on this page with those of a neighboring page. If there is sufficient space on one of these pages, the merge is performed and one of the pages is deleted. The value of integer can be from 0 to 99.</p> <p>Oracle: The Logging, No Sort, Degrees, and Instances properties are documented in the editor.</p> |
| Java Classes | Oracle | <p>A model or template, written in Java language, used to create objects with a common definition and common properties, operations and behavior.</p> <p>Java classes can be developed in Eclipse (or another Java development environment such as Oracle JDeveloper) and moved into an Oracle database to be used as stored procedures.</p> <p>Java classes must be public and static if they are to be used in this manner.</p> <p>When writing a class to be executed within the database, you can take advantage of a special server-side JDBC driver. This driver uses the user's default connection and provides the fastest access to the database.</p> <p>Java classes become full-fledged database objects once migrated into the database via the loadjava command-line utility or the SQL CREATE JAVA statement.</p> <p>A Java class is published by creating and compiling a call specification for it. The call spec maps a Java method's parameters and return type to Oracle SQL types.</p> <p>Once a Java class is developed, loaded, and published – the final step is to execute it.</p> |
| Java Resources | Oracle | <p>A Java resource is a collection of files compressed in a .jar file.</p> |

| Data base Object | DB MS Plat forms | Notes |
|------------------------|------------------|--|
| Libraries | Oracle | <p>A library is a configurable folder for storing and sharing content with an allocated quota. Multiple libraries may exist in the same database environment.</p> <p>A library is a special type of folder in Oracle Content Services. Unlike Containers and regular folders, each library has a Trash Folder and an allocated amount of disk space.</p> <p>A library is composed of a name (mandatory), description, quota, path, and library members.</p> <p>The library service allows you to create folders, list quotas, and manage categories, workflow, trash folders, and versioning. The Library service does not allow you to create or upload files.</p> |
| Materialized Views | Oracle | <p>A database object that contains the results of a query. They are local copies of data located remotely, or are used to create summary tables based on aggregations of table data. Materialized views are also known as snapshots.</p> <p>A materialized view can query tables, views, and other materialized views. Collectively, these are called master tables (a replication term) or detail tables (a data warehouse term).</p> <p>For replication purposes, materialized views allow you to maintain copies of remote data on your local node. These copies are read-only. If you want to update the local copies, you need to use the Advanced Replication feature. You can select data from a materialized view as you would from a table or view.</p> <p>For data warehousing purposes, the materialized views commonly created are aggregate views, single-table aggregate views, and join views.</p> |
| Materialized View Logs | Oracle | <p>Because Materialized Views are used to return faster queries (a query against a materialized view is faster than a query against a base table because querying the materialized view does not query the source table), the Materialized View often returns the data at the time the view was created, not the current table data.</p> <p>There are two ways to refresh data in Materialized Views, manually or automatically. In a manual refresh, the Materialized View is completely wiped clean and then repopulated with data from the source tables (this is known as a complete refresh). If source tables have changed very little, however, it is possible to refresh the Materialized View only for changed records – this is known as a fast refresh.</p> <p>In the case of Materialized Views that are updated via fast refresh, it is necessary to create Materialized View Logs on the base tables that compose the Materialized View to reflect the changes.</p> <p>If the number of entries in this table is too high, it is an indication that you might need to refresh the Materialized Views more frequently to ensure that each update does not take longer than it needs.</p> <p>Select owner, then select from tables with Materialized Views, etc.</p> |

| Data base Object | DB MS Plat forms | Notes |
|------------------|------------------|---|
| Oracle Job Queue | Oracle | <p>The Oracle Job Queue allows for the scheduling and execution of PL/SQL stored procedures at predefined times and/or repeated job execution at regular intervals, as background processes.</p> <p>For example, you could create a job in the Oracle Job Queue that processed end-of-day accounting – a job that must run every weekday, but can be run unattended, or you could create a series of jobs that must be run sequentially – such as jobs that might be so large, that in order to reduce CPU usage, only one is run at a time.</p> <p>Runs PL/SQL code at specified time or on specified schedule, can enable/disable.</p> |
| Outlines | Oracle | <p>Oracle preserves the execution plans of "frozen" access paths to data so that it remains constant despite data changes, schema changes, and upgrades of the database or application software through objects named stored outlines.</p> <p>Outlines are useful for providing stable application performance and benefit high-end OLTP sites by having SQL execute without having to invoke the cost-based optimizer at each SQL execution. This allows complex SQL to be executed without the additional overhead added by the optimizer when it performs the calculations necessary to determine the optimal access path to the data.</p> |

| Data base Object | DB MS Plat forms | Notes |
|------------------|------------------|---|
| Pack ages | All | <p>A package is a procedural schema object classified as a PL/ SQL program unit that allows the access and manipulation of database information.</p> <p>A package is a group of related procedures and functions, together with the cursors and variables they use, stored together in the database for continued use as a unit. Similar to standalone procedures and functions, packaged procedures and functions can be called explicitly by applications or users.</p> <p>DB applications explicitly call packaged procedures as necessary with privileges granted, a user can explicitly execute any of the procedures contained in it.</p> <p>Packages provide a method of encapsulating related procedures, functions, and associated cursors and variables together as a unit in the database. For example, a single package might contain two statements that contain several procedures and functions used to process banking transactions.</p> <p>Packages allow the database administrator or application developer to organize similar routines as well as offering increased functionality and database performance.</p> <p>Packages provide advantages in the following areas: encapsulation of related procedures and variables, declaration of public and private procedures, variables, constraints and cursors, separation of the package specification and package body, and better performance.</p> <p>Encapsulation of procedural constructs in a package also makes privilege management easier. Granting the privilege to use a package makes all constructs of the package assessable to the grantee.</p> <p>The methods of package definition allow you to specify which variables, cursors, and procedures are: public, directly accessible to the users of a package, private, or hidden from the user of the package.</p> |
| Pack age Bodi es | Orac le | <p>A package body is a package definition file that states how a package specification will function.</p> <p>In contrast to the entities declared in the visible part of a package, the entities declared in the package body are only visible within the package body itself. As a consequence, a package with a package body can be used for the construction of a group of related subprograms in which the logical operations available to clients are clearly isolated from the internal entities.</p> |

| Data base Object | DB MS Plat forms | Notes |
|------------------|------------------|--|
| Primary Keys | All | <p>A key is a set of columns used to identify or access a row or rows. The key is identified in the description of a table, index, or referential constraint. The same column can be part of more than one key.</p> <p>A unique key is a key that is constrained so that no two of its values are equal. The columns of a unique key cannot contain NULL values.</p> <p>The primary key is one of the unique keys defined on a table, but is selected to be the key of the first importance. There can only be one primary key on a table.</p> <p>Oracle: If an index constraint has been defined for a table, the constraint status for the table's primary key cannot be set to Disabled.</p> |
| Procedures | All | <p>A procedure is an application program that can be started through the SQL CALL statement. The procedure is specified by a procedure name, which may be followed by arguments enclosed within parenthesis.</p> <p>The argument or arguments of a procedure are individual scalar values, which can be of different types and can have different meanings. The arguments can be used to pass values into the procedure, receive return values from the procedure, or both.</p> <p>A procedure, also called a stored procedure, is a database object created via the CREATE PROCEDURE statement that can encapsulate logic and SQL statements. Procedures are used as subroutine extensions to applications, and other database objects that can contain logic.</p> <p>When a procedure is invoked in SQL and logic within a procedure is executed on the server, data is only transferred between the client and the database server in the procedure call and in the procedure return. If you have a series of SQL statements to execute within a client application, and the application does not need to do any processing in between the statements, then this series of statements would benefit from being included in a procedure.</p> |
| Profiles | Oracle | <p>Profiles are a means to limit resources a user can use by specifying limits on kernel and password elements. Additionally, Profiles can be used to track password histories and the settings of specific profiles may be queried.</p> <p>The following kernel limits may be set: maximum concurrent sessions for a user, CPU time limit per session, maximum connect time, maximum idle time, maximum blocks read per session, maximum blocks read per call, and maximum amount of SGA.</p> |

| Data base Object | DB MS Plat forms | Notes |
|------------------|------------------|---|
| Roles | Oracle | <p>A role is a set or group of privileges that can be granted to users to another role.</p> <p>A privilege is a right to execute a particular type of SQL statement or to access another user's object. For example: the right to connect to a database, the right to create a table, the right to select rows from another user's table, the right to execute another user's stored procedure.</p> <p>System privileges are rights to enable the performance of a particular action, or to perform a particular action on a particular type of object.</p> <p>Roles are named groups of related privileges that you grant users or other roles. Roles are designed to ease the administration of end user system and object privileges. However, roles are not meant to be used for application developers, because the privileges to access objects within stored programmatic constructs needs to be granted directly.</p> |
| Sequences | DB2, Oracle | <p>A sequence generates unique numbers.</p> <p>Sequences are special database objects that provide numbers in sequence for input into a table. They are useful for providing generated primary key values and for the input of number type columns such as purchase order, employee number, sample number, and sales order number.</p> <p>Sequences are created by use of the CREATE SEQUENCE command.</p> |
| Structured Types | DB2 | <p>Structured Types are useful for modeling objects that have a well-defined structure that consists of attributes. Attributes are properties that describe an instance of the type.</p> <p>A geometric shape, for example, might have as attributes its list of Cartesian coordinates. A person might have attributes of name, address, and so on. A department might have a name or some other attribute.</p> |
| Synonyms | Oracle | <p>A synonym is an alternate name for objects such as tables, views, sequences, stored procedures, and other database objects.</p> <p>A synonym is an alias for one of the following objects: table, object table, view, object view, sequence, stored procedure, stored function, package, materialized view, java class, user-defined object type or another synonym.</p> |
| Tables | All | <p>Tables are logical structures maintained by the database manager. Tables are composed of columns and rows. The rows are not necessarily ordered within a table.</p> <p>A base table is used to hold persistent user data.</p> <p>A result table is a set of rows that the database manager selects or generates from one or more base tables to satisfy a query.</p> <p>A summary table is a table defined by a query that is also used to determine the data in the table.</p> |

| Data base Object | DB MS Plat forms | Notes |
|------------------|------------------|---|
| Tablespaces | DB2, Oracle | <p>A tablespace is a storage structure containing tables, indexes, large objects, and long data. Tablespaces reside in database partition groups. They allow you to assign the location of database and table data directly onto containers. (A container can be a directory name, a device name, or a file name.) This can provide improved performance and more flexible configuration.</p> |
| Triggers | All | <p>A trigger defines a set of actions that are performed when a specified SQL operation (such as delete, insert, or update) occurs on a specified table. When the specified SQL operation occurs, the trigger is activated and starts the defined actions.</p> <p>Triggers can be used, along with referential constraints and check constraints, to enforce data integrity rules. Triggers can also be used to cause updates to other tables, automatically generate or transform values for inserted or updated rows, or invoke functions to perform tasks such as issuing alerts.</p> |
| Undo Segments | Oracle | <p>In an Oracle database, Undo tablespace data is an image or snapshot of the original contents of a row (or rows) in a table. The data is stored in Undo segments in the Undo table space.</p> <p>When a user begins to make a change to the data in a row in an Oracle table, the original data is first written to Undo segments in the Undo tablespace. The entire process (including the creation of the Undo data is recorded in Redo logs before the change is completed and written in the Database Buffer Cache, and then the data files via the database writer (DBW) process.)</p> |
| Unique Keys | All | <p>A unique key is a key that is constrained so that no two of its values are equal. The columns of a unique key cannot contain null values. The constraint is enforced by the database manager during the execution of any operation that changes data values, such as INSERT or UPDATE. The mechanism used to enforce the constraint is called a unique index. Thus, every unique key is a key of a unique index. Such an index is said to have the UNIQUE attribute.</p> <p>A primary key is a special case of a unique key. A table cannot have more than one primary key.</p> <p>A foreign key is a key that is specified in the definition of a referential constraint.</p> <p>A partitioning key is a key that is part of the definition of a table in a partitioned database. The partitioning key is used to determine the partition on which the row of data is stored. If a partitioning key is defined, unique keys and primary keys must include the same columns as the partitioning key, but can have additional columns. A table cannot have more than one partitioning key.</p> <p>Oracle: You cannot drop a unique key constraint that is part of a referential integrity constraint without also dropping the foreign key. To drop the referenced key and the foreign key together, check the Delete Cascade option for the foreign key.</p> <p>Clustered: A cluster composes of a group of tables that share the same data blocks, and are grouped together because they share common columns and are often used together.</p> <p>Filegroup: Lets you select the filegroup within the database where the constraint is stored.</p> <p>Fill Factor: Lets you specify a percentage of how large each constraint can become.</p> |

| Data base Object | DB MS Plat forms | Notes |
|------------------|------------------|---|
| Views | All | <p>A view provides an alternate way of looking at the data in one or more tables.</p> <p>A view is a named specification of a result table and can be thought of as having columns and rows just like a base table. For retrieval purposes, all views can be used just like base tables. You can use views to select certain elements of a table and can present an existing table in a customized table format without having to create a new table.</p> |

DBMS connection parameters by platform

The following topics provide connection details:

- [Using Multi-Factor Authentication](#)
- [IBM DB2 LUW](#)
- [Microsoft SQL Server](#)
- [JDBC connection parameters](#)
- [Oracle connection parameters](#)
- [Sybase connection parameters](#)

Using Multi-Factor Authentication

DB PowerStudio users can leverage multi-factor authentication (MFA) using the identity provider Microsoft Entra, formerly Azure Active Directory (AD), when connecting to Oracle. When logging into Oracle, you can gain the inherent security policies and procedures from Microsoft Entra by entering your password only once.

Users can choose from authentication types for the following databases:

- **Microsoft Entra authentication with Azure Database**
 - PostgreSQL
- **Microsoft Entra ID**
 - Oracle
- **Microsoft SQL Server**
 - Active Directory - Managed Identity
 - Azure Active Directory - Password
 - Azure Active Directory - Service Principal
 - Azure Active Directory - Universal with MFA
 - SQL Server Authentication
 - Windows Authentication

IBM DB2 LUW

| Connection Parameter | Description |
|---|--|
| Use Alias from IBM Client or Generic JDBC Configuration | If you choose to use the alias from the IBM client, select the appropriate alias name. Otherwise, choose Generic JDBC Configuration and enter the connection parameters, as specified. |
| Schema ID | Optional. The name of the database schema. |
| Function Path | Optional. Enter an ordered list of schema names to restrict the search scope for unqualified function invocations. |
| Security Credentials | The log on information required by SQL Query Tuner to connect to the data source. |
| Auto Connect | Automatically attempts to connect to the data source when selected in Data Source Explorer, without prompting the user for connection information. |
| JDBC Driver (Advanced) | The name of the JDBC Driver utilized by SQL Query Tuner to initiate a JDBC standard access connection. |
| Connection URL (Advanced) | Used by the JDBC Driver to connect with a data source. Often contains host and port numbers, as well as the name of the data source to which it connects. For example: <code>jdbc:postgresql://host:port/database</code> <code>jdbc:derby://host:port/database</code> |
| Custom JDBC Driver Properties (Advanced) | The name and property value of any custom JDBC drivers associated with the data source. |

Microsoft SQL Server

| Connection Parameter | Description |
|------------------------------------|---|
| Use Network Library Configuration | If the data source utilizes a network library, select this parameter. The corresponding connection parameter fields become available. Otherwise, choose Generic JDBC Configuration and enter the connection parameters, as specified. |
| Host/Instance (JDBC Configuration) | The name of the data source. |

| Connection Parameter | Description |
|--|--|
| Port (JDBC Configuration) | Optional. The listening port used in TCP/IP communications between DB Optimizer and the data source. |
| Protocol (JDBC Configuration) | The communication mechanism between DB Optimizer and the data source. Choose TCP/IP or Named Pipes. |
| Default Database (Optional) | The default SQL database name, as defined by the schema. |
| Security Credentials | The log on information required by DB Optimizer to connect to the data source. |
| Auto Connect | Automatically attempts to connect to the data source when selected in Data Source Explorer, without prompting the user for connection information. |
| Allow Trusted Connections | Enables trusted connections to the data source from DB Optimizer. |
| JDBC Driver (Advanced) | The name of the JDBC Driver utilized by DB Optimizer to connect and communicate with the database. |
| Connection URL (Advanced) | <p>Used by the JDBC Driver to connect with a database. Often contains host and port numbers, as well as the name of the database to which it connects.</p> <p>For example:</p> <pre>jdbc:postgresql://host:port/database jdbc:derby://host:port/database</pre> |
| Custom JDBC Driver Properties (Advanced) | The name and property value of any custom JDBC drivers associated with the data source. |

JDBC connection parameters

| Connection Parameter | Description |
|----------------------|--|
| Connect String | Used by the JDBC Driver to connect with a database. Often contains host and port numbers, as well as the name of the database to which it connects. For example: <code>jdbc:postgresql://host:port/database</code> <code>jdbc:derby://host:port/database</code> |
| Data Source Name | The name of the data source to which you want DB Optimizer to connect. |

Oracle connection parameters

| Connection Parameter | Description |
|---------------------------------------|---|
| Use TNS Alias | If the data source is mapped to a net service name via tnsnames.ora, select this parameter. Otherwise, choose Generic JDBC Configuration and enter the connection parameters, as specified. |
| Host/Instance (JDBC Configuration) | The name of the host machine on which the data source resides. Port (JDBC Connection)The listening port used in TCP/IP communications between DB Optimizer and the data source. |
| Type (JDBC Configuration) | Indicates if the data source is defined via a system identifier (SID) or a service name. |
| Service/SID Name (JDBC Configuration) | The name of the system identifier (SID) or service name that identifies the data source. |
| Security Credentials | The log on information required by DB Optimizer to connect to the data source. |
| Auto Connect | Automatically attempts to connect to the data source when selected in data source Explorer, without prompting the user for connection information. |
| Allow Trusted Connections | Enables trusted connections to the data source from DB Optimizer. |

| Connection Parameter | Description |
|--|--|
| JDBC Driver (Advanced) | The name of the JDBC Driver utilized by DB Optimizer to connect and communicate with the database. |
| Connection URL (Advanced) | Used by the JDBC Driver to connect with a database. Often contains host and port numbers, as well as the name of the database to which it connects. For example: <code>jdbc:postgresql://host:port/database</code> <code>jdbc:derby://host:port/database</code> |
| Custom JDBC Driver Properties (Advanced) | The name and property value of any custom JDBC drivers associated with the data source. |

Sybase connection parameters

| Connection Parameter | Description |
|--|--|
| Use Alias Information from your SQL.INI File | If the data source is mapped to a name via SQL.INI, select this parameter to use that name for connection. Otherwise, choose Generic JDBC Configuration and enter the connection parameters, as specified. |
| Host/Instance (JDBC Connection) | The name of the host machine on which the data source resides. |
| Port (JDBC Connection) | The listening port used in TCP/IP communications between DB Optimizer and the data source. |
| Default Database (JDBC Connection) | Optional. The default database name, as defined by the schema. |
| JDBC Driver (Advanced) | The name of the JDBC Driver utilized by DB Optimizer to connect and communicate with the database. |
| Connection URL (Advanced) | Used by the JDBC Driver to connect with a database. Often contains host and port numbers, as well as the name of the database to which it connects. For example: <code>jdbc:postgresql://host:port/database</code> <code>jdbc:derby://host:port/database</code> |

| Connection Parameter | Description |
|--|---|
| Custom JDBC Driver Properties (Advanced) | The name and property value of any custom JDBC drivers associated with the data source. |