



Product Documentation

DB Optimizer™

Quick Start Guide

Version 3.5

Published April 23, 2012

© 2012 Embarcadero Technologies, Inc. Embarcadero, the Embarcadero Technologies logos, and all other Embarcadero Technologies product or service names are trademarks or registered trademarks of Embarcadero Technologies, Inc. All other trademarks are property of their respective owners.

Embarcadero Technologies, Inc. is a leading provider of award-winning tools for application developers and database professionals so they can design systems right, build them faster and run them better, regardless of their platform or programming language. Ninety of the Fortune 100 and an active community of more than three million users worldwide rely on Embarcadero products to increase productivity, reduce costs, simplify change management and compliance and accelerate innovation. The company's flagship tools include: Embarcadero® Change Manager™, CodeGear™ RAD Studio, DBArtisan®, Delphi®, ER/Studio®, JBuilder® and Rapid SQL®. Founded in 1993, Embarcadero is headquartered in San Francisco, with offices located around the world. Embarcadero is online at www.embarcadero.com.

April 23, 2012

Contents

- Introduction 7
 - About this Document 7
 - Major Interface Components 7
 - Limitations of the Trial Version 8
 - Limitations of DB Optimizer XE Developer 8
 - New Features of DB Optimizer 9
 - New Features of DB Optimizer 3.5 10
 - Profiling 10
 - Tuning 12
 - New Features of DB Optimizer 3.0 19
- Installation 25
 - Technical Requirements 25
 - Installation Issues 25
 - Installing DB Optimizer 25
 - Installation Issues 26
 - Installing DB Optimizer on Linux 26
 - Licensing DB Optimizer 27
- Getting Started with DB Optimizer 31
 - Start DB Optimizer 31
 - Overview of the User Interface 32
 - Working with Data Source Explorer 34
 - Adding Data Sources 35
 - Add a Data Source 36
 - Browsing Data Sources 37
 - Profiling a Data Source 39
 - Starting a Profiling Session 40
 - Analyzing Session Data 40
 - Load Chart 41
 - Top Activity 41
 - Profiling Details 43
 - View Session Details 44
 - View SQL 45
 - Saving a Profiling Session 46
 - Save a Profiling Session as a .oar File 46

Contents

Save a Profiling Session in the Profiling Repository	46
Importing Statements to SQL Tuner	47
Tuning SQL Statements	48
Creating a New Tuning Job	49
Create a New Tuning Job	50
Adding SQL Statements to a Tuning Job	50
Add New SQL Text	51
Add a Database Object	52
Add a Saved SQL File	52
Add Active SQL in SGA	53
Running a Tuning Job	53
Run a Tuning Job	55
Analyzing Tuner Results on the Overview Tab	55
Determine the Best Cases for Statement Execution Path Time	56
Finding Missing Indexes and SQL Problems	57
Finding Missing Indexes	58
Exploring the VST Diagram	59
Displaying or Hiding the Diagram Legend	60
Displaying or Hiding the Explain Plan	61
Displaying or Hiding Table Counts and Ratios	61
Displaying or Hiding Table Columns and Indexes	62
Viewing All Fields in a Table	63
Viewing Object SQL	64
Viewing Relationships	64
Expanding Views in the VST Diagram	64
Finding Problematic SQL or Schema	66
Implementing Recommendations on the Index Analysis Tab	67
Accept the Suggestion and Automatically Generate an Index	67
SQL Code Assist and Execution	68
Code Extraction	69
Code Highlighting	70
Automatic Error Detection	70
Code Complete	72
Hyperlinks	72
Code Formatting	73
Code Folding	74
Code Quality Checks	74
SQL Execution	74

Execute a File	75
Execute a Transaction.....	76
Commit a Transaction.....	76
Configuring SQL Execution Parameters	76
Modify SQL Session Options	77

Contents

INTRODUCTION

Embarcadero DB Optimizer simplifies SQL optimization and development 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 offers native support for IBM® DB2® for LUW, Oracle®, Microsoft® SQL Server and Sybase® as well as JDBC support for other DBMS. It is available as a stand-alone application, or as an Eclipse plug-in. DB Optimizer is structured and composed into three main interface parts. This design provides a comprehensive workflow that enables development, query analysis, and tuning capabilities. This workflow, in turn, leads to more efficient task management in terms of time and efficiency, overall.

ABOUT THIS DOCUMENT

This document helps you install DB Optimizer and introduces you to the basic functionality of DB Optimizer.

NOTE: For the latest versions of DB Optimizer documents and for information that will help you install and upgrade, including technical requirements and databases supported, see the Embarcadero documentation website. Always make sure to consult the README for late breaking updates at docs.embarcadero.com/products/db_optimizer.

MAJOR INTERFACE COMPONENTS

The three major interface components of DB Optimizer are as follows:

SQL Profiler: Provides continuous data source monitoring that builds a statistical model or profile, of the specified data source, and highlights top SQL, event, and session activity. This component is used to locate and diagnose problematic SQL code and event-based bottlenecks via its graphical interface, which is used to identify problem areas and drill down to individual, problematic statements. Additionally, Profiler enables the investigation of execution and wait time event details for individual stored routines. Profiling Details of the SQL Profiler have been expanded to show Session Details for Sybase, and SQL Server and the SQL that ran in the selected session for Sybase, SQL Server, and DB2.

SQL Tuner: 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. Tuner enables the optimization of poorly performing SQL code through the detection and modification of execution paths in data retrieval through hint injections. Users are supplied with a list of possible cases generated by Tuner, and can select and update a statement with the most efficient path to reduce load and improve efficiency, overall. The Analysis tab of the tuner provides a graphical diagram of an SQL query to show how the tables in the query would be joined to satisfy the query and also provides suggestions for the creation of indexes that may increase query performance.

SQL Editor: Simplifies SQL development by utilizing features that improve productivity and reduce errors. It provides a rich interface that offers code completion, real-time error checking, code formatting, and sophisticated object validation tools. In context with SQL Profiler and SQL Tuner, it provides an interface for viewing and editing SQL files and database packages, as accessed through Profiler and Tuner functionality.

LIMITATIONS OF THE TRIAL VERSION

The following limitations apply to the trial version of DB Optimizer. These limitations are removed once the product has been purchased.

- Command line profiling and tuning are not available.
- There are a limited number of profiling sessions you can run.
- Profiling is limited in that you will not be able to see the captured SQL text of the data source profiled.
- Only the Ad Hoc SQL tab is available for specifying tuning sources. You cannot tune database objects, SQL files or the Oracle SGA.
- There are a limited number of statement extractions from the profiler to the tuner.
- In the tuning editor the following case actions are disabled: Clone Case, Edit Case, and Create Custom Case.

LIMITATIONS OF DB OPTIMIZER XE DEVELOPER

The following features are not supported on DB Optimizer XE Developer:

- Visual SQL Tuning.
- Hint analysis via SQL hint injection for Oracle data sources.
- Saving profile sessions to an Oracle data source repository.

NEW FEATURES OF DB OPTIMIZER

This section comprises the following topics:

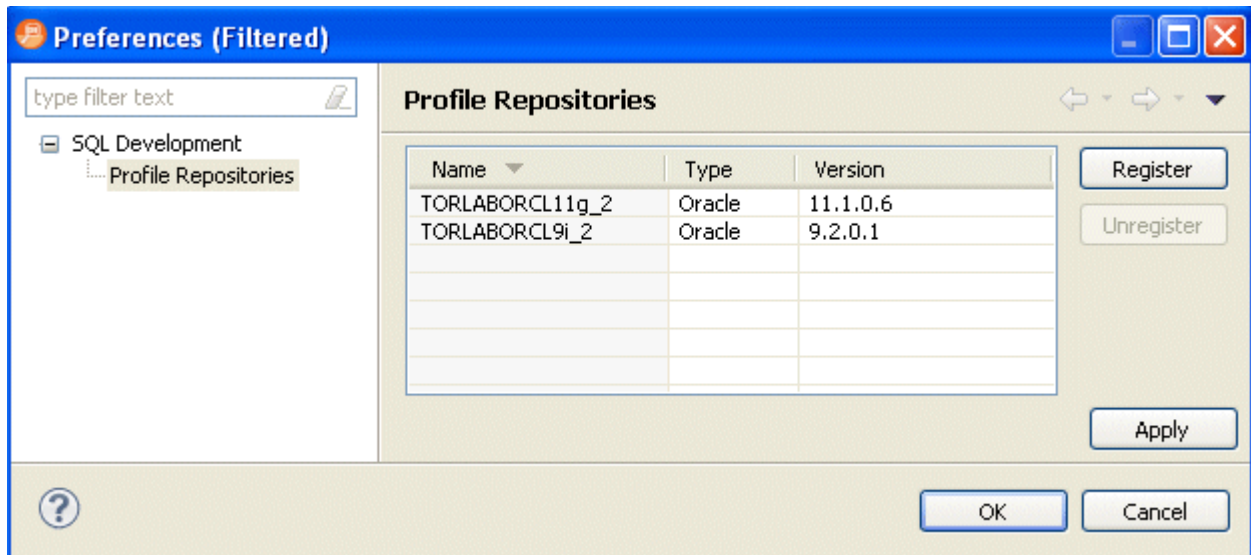
- [New Features of DB Optimizer 3.5](#) on page 10
- [New Features of DB Optimizer 3.0](#) on page 20

NEW FEATURES OF DB OPTIMIZER 3.5

PROFILING

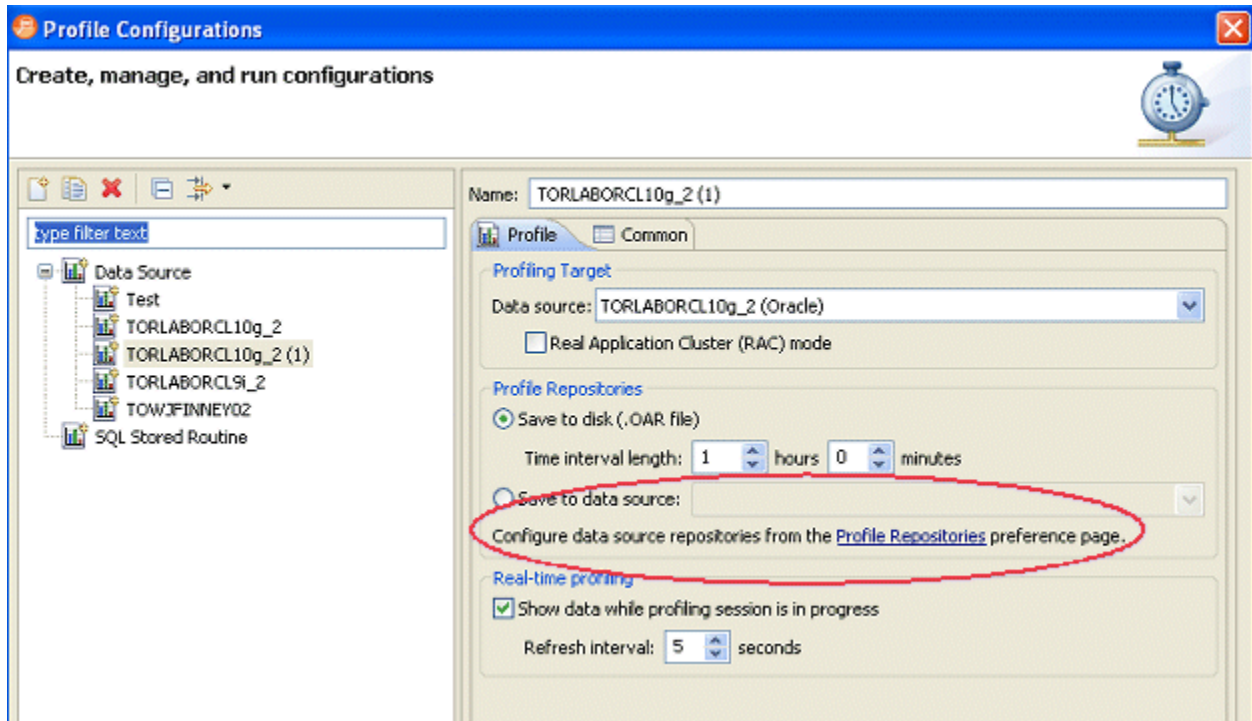
Repository management has been improved and simplified.

There is now a new preference page, entitled Profile Repositories:



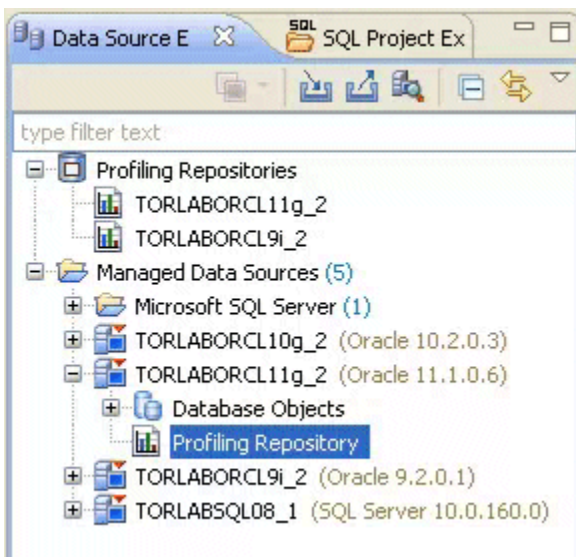
The Profile Repositories preferences are available from the Preferences Menu. Here you can register and unregister Oracle 9+ data sources as repositories. A data source must be available in the Data Source Explorer before it can be used as a repository.

You can also access the Profile Repositories Preferences dialog from the Profile Configurations dialog.



In the Profiling Target area of the Profile Configurations dialog, you can click the Data source list to choose from a list of registered Profiling Repository data sources.

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.




Unregistering a data source prompts you to optionally delete all sessions on the repository.

New filters for Profiling

You can now further limit the activity you see for a data source in a Profiling job, using the new filters that are available.

- Application/Program filter available on all four platforms
- User Name filter available on all four platforms.
- Schema Name filter available for Oracle and DB2.
- Database filter available for SQL Server and Sybase.




Variable/parameter capture for Oracle

Profiling an Oracle datasource now captures and sets all properties of variables/parameters when tuning a SQL statement captured in a profiling session.

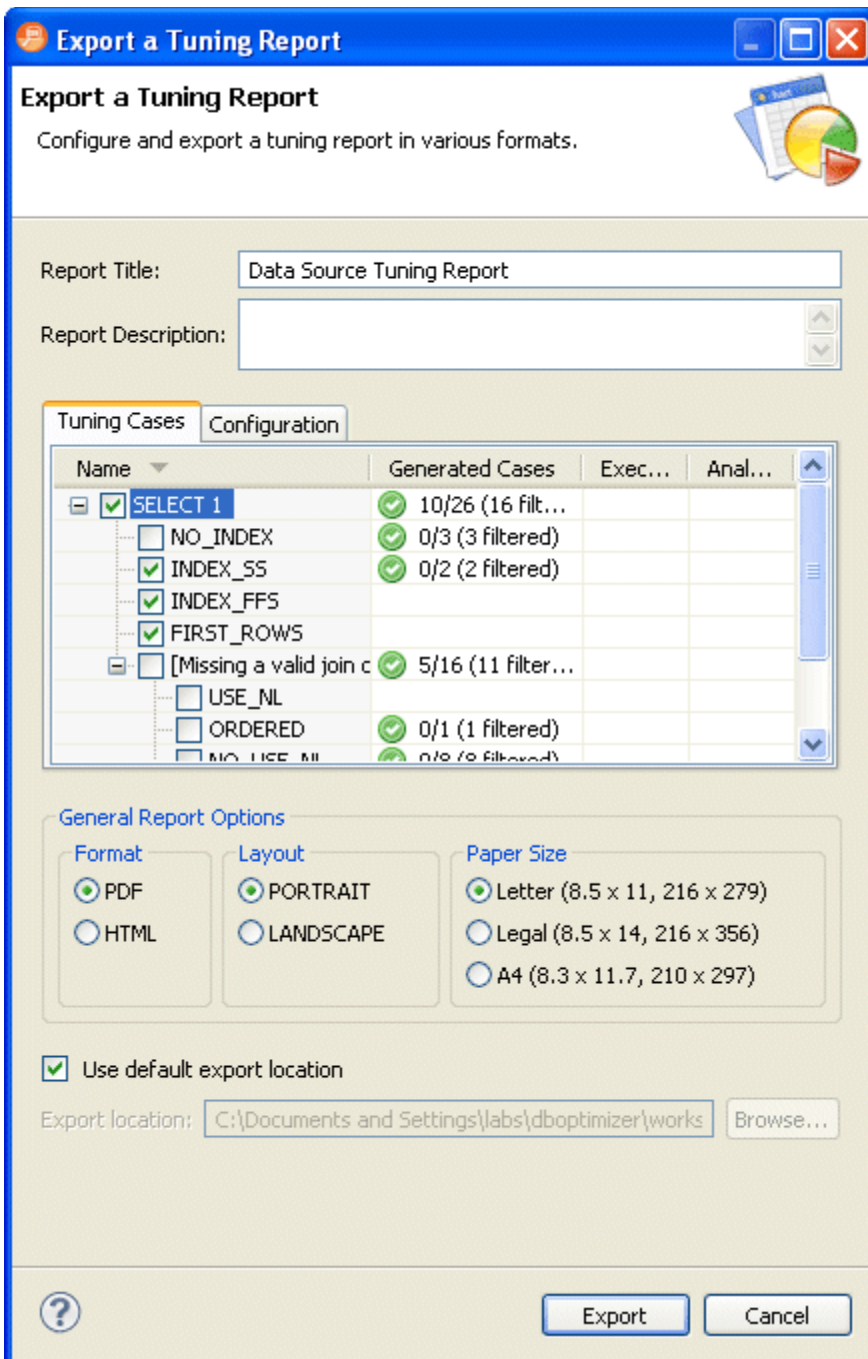
TUNING

In addition to offering a Profiling Report, DB Optimizer now offers a Tuning Report.

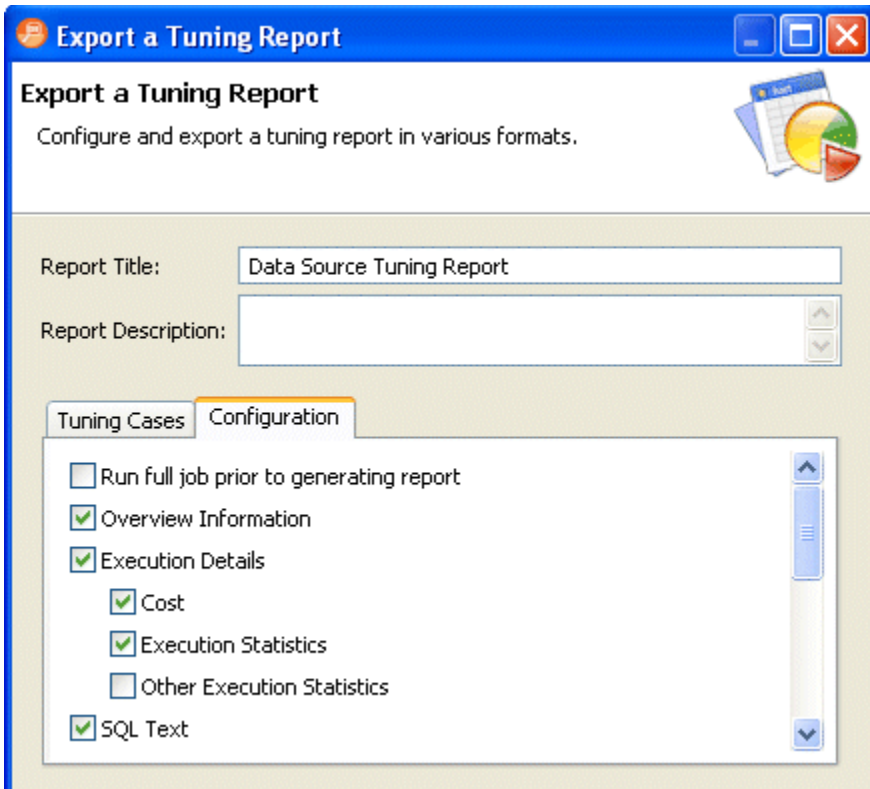
From a tuning session, click the **Export Report** button to export tuning session details.



On the **Tuning Cases** tab of the **Export a Tuning Report** dialog, you can choose to report on all the tuning cases or on selected cases only.



On the **Configuration** tab of the **Export a Tuning Report** dialog you can customize the content of the report by choosing from a variety of options.



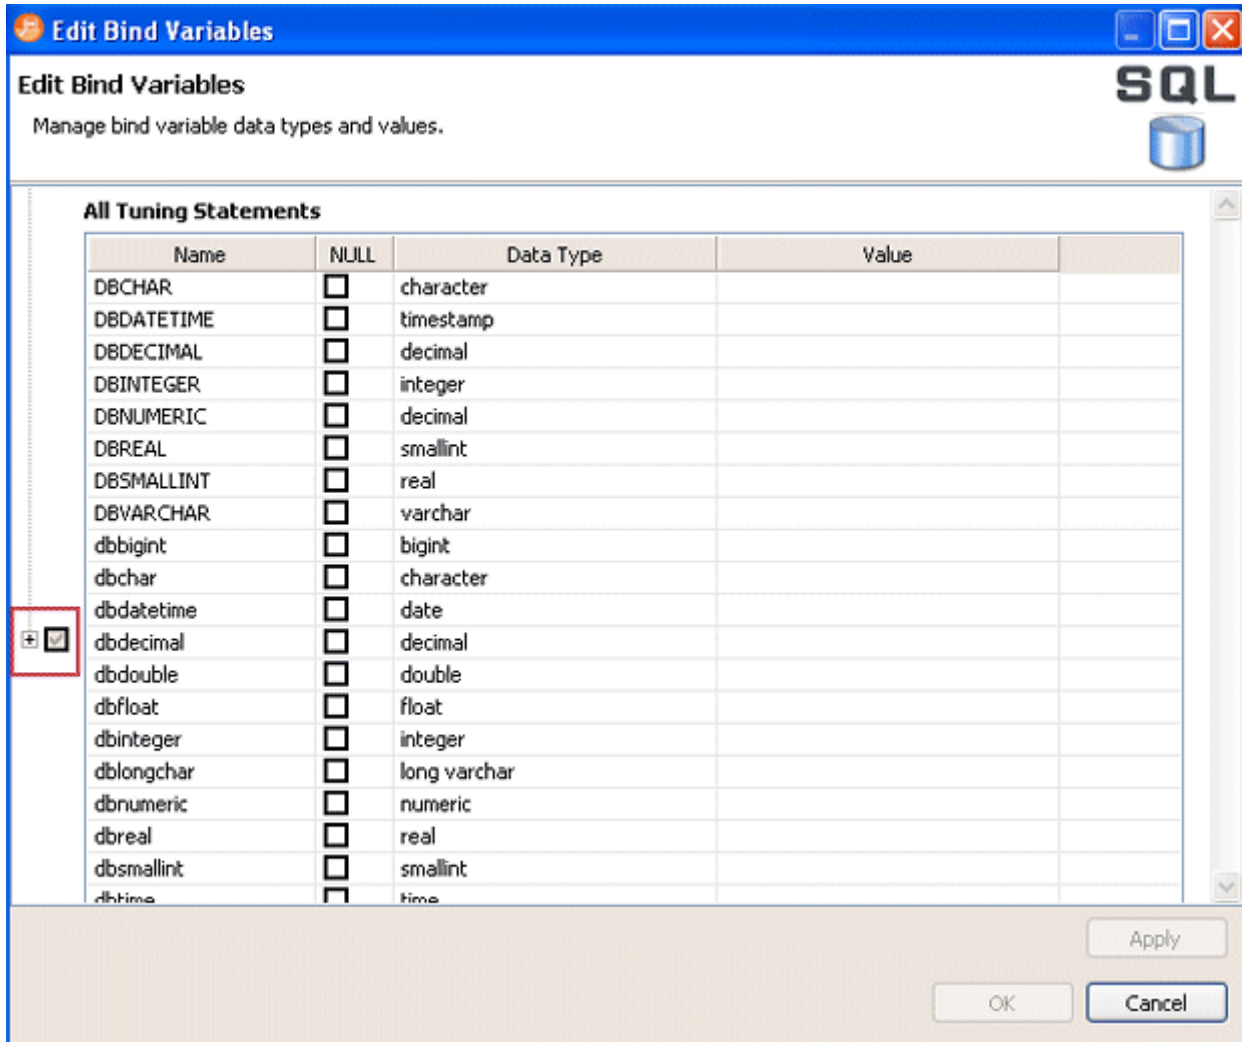
Vastly improved bind variable management

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

- 1 Click the **Edit Bind Variable** icon.



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 Statements** 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.

SELECT 11 (select from MOVIES.CUSTOMER, MOVIES.MOVIERENTAL)

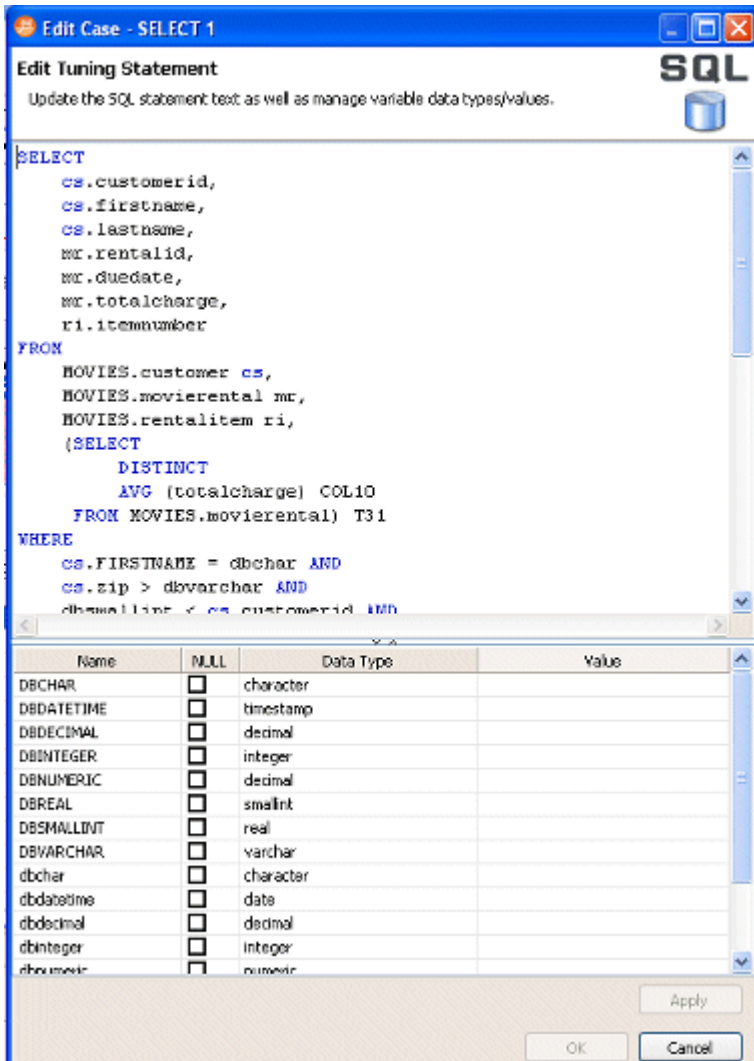
Name	NULL	Data Type	Value
<input checked="" type="checkbox"/> DBCHAR	<input type="checkbox"/>	character	
<input type="checkbox"/> dbchar	<input type="checkbox"/>	date	
<input type="checkbox"/> dulongchar	<input type="checkbox"/>	integer	
		bigint	
		blob	
		boolean	

SELECT 12 (select from employee)

In this manner you can set the bind variables for 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.

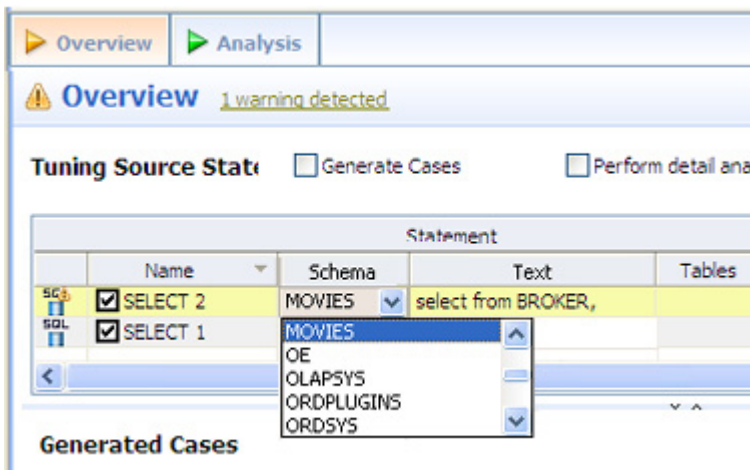


Automatic bind variable extraction for Oracle SGA

When you select to tune statements by scanning SGA on an Oracle data source, bind variables and their variable are automatically extracted.

More intuitive UI (catalog/schema overrides, simplified workflows)

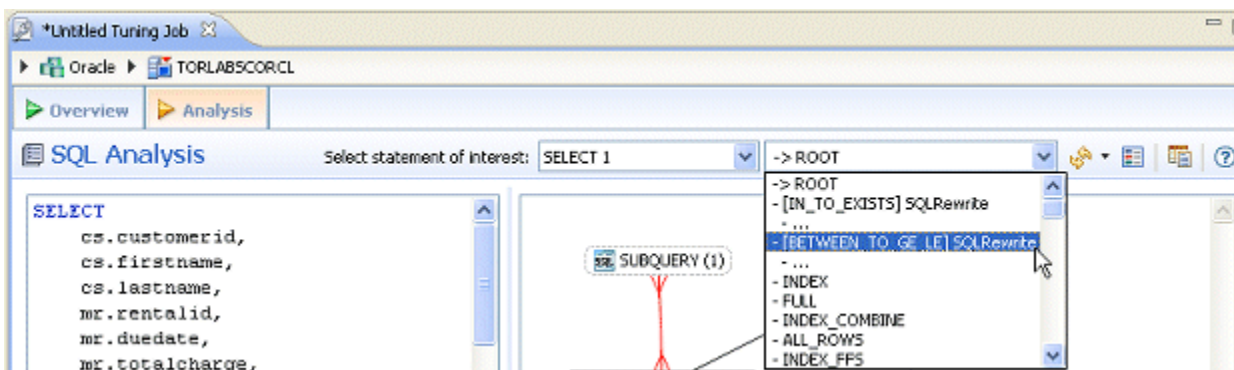
Use the **Schema** and **Catalog** selectors to select a schema and catalog for the tuning job. The catalog selector is available only for SQL Server and Sybase data sources. By specifying the schema and catalog, the tuner can use the paths of the schema and catalog selected to find the tables queried 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 select **Refresh Tuning Statements**.



VST diagrams

You can now select any generated case and see a VST diagram for it.

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 or case you chose. If warranted, a new diagram is displayed and the Index Analysis for all data sources is updated. In addition, for Oracle data sources, the Table Statistics, Column Statistics And Histograms, and Outlines are also updated.

Improved SQL coverage, across all platforms

Added full support for aliases and synonyms in the tuning editor for all platforms, previously only supported on the Oracle platform.

Added full support for Materialized Views on Oracle and Materialized Query Tables on DB2.

Query rewrites, previously only available on Oracle, are now available on all platforms: DB2, Oracle, SQL Server, and Sybase.

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
select * from t1 where EXISTS (select null from t2 where t2.key = t1.key);	select * from t1 where t1.key IN (select t2.key from t2);
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;
For DB2 and Oracle only 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)

Other

- DB Optimizer has been updated to take advantage of the improvements in Eclipse 3.6 and Java 1.6.
- The user interface and user documentation of DB Optimizer has been localized to Chinese.

NEW FEATURES OF DB OPTIMIZER 3.0

User Interface Changes

- **Removal of Eclipse-specific UI elements:** To provide a better focus on DB Optimizer-specific functionality and workflows, UI elements provided by the underlying Eclipse framework aimed primarily at Java development have been removed. These include the **Navigate**, **Project**, and **Run** menus, as well as toolbar buttons for debugging and code navigation.
- **Shortcuts to DB Optimizer Preferences:** A new **Preferences** menu has been added with shortcuts to the preference pages relevant to DB Optimizer configuration. These pages (and their associated menu entries) are **Data Sources**, **SQL Development**, **SQL Editor**, **SQL Execution**, **SQL Filters**, **Data Source Indexing**, **Profile Alerts**, and **Tuning Job Editor**. The set of all preferences is available under **Preferences > Other** option, and the previously-available **Window > Preferences** option.

Object Indexing Enabled by Default

Previous versions of DB Optimizer had the option to create an index of schema objects on a data source in order to provide code completion and semantic validation of object names in the SQL Editor, but this was disabled by default, because it could affect data source performance. In DB Optimizer 3.0, the speed and efficiency of Object Indexing has been significantly improved, so this functionality is now enabled by default. This option can still be disabled from the preferences window (available from the **Preferences > Data Source Indexing** option).

New Actions on Oracle Sessions During Profiling

While profiling an Oracle data source, the **Sessions** tab in the profiling editor has two new context menu actions on rows in the session grid:

- **Kill Session:** Instructs the database to kill the given session. The session will still appear within the profiling editor, based on its activity prior to being killed.
- **Trace Session:** Instructs the database to begin a full trace operation on the session. (This simply executes `set_sql_trace_in_session` from Oracle's `dbms_system` package for the given session.)

MS SQL Server and Sybase ASE Profiling Improved

Profiling on Microsoft SQL Server and Sybase Adaptive Server Enterprise are now more efficient, contributing less than 2% load on the profiled data source under most conditions. Furthermore, more accurate data is collected.

Improved User Interface for the Tuning Editor

The dedicated **Input** tab of the Tuning Editor has been removed, allowing the user to specify statements for tuning directly within the **Overview** tab.

In the top-right of the Tuning Editor, there is an input selection button, to add statements for tuning from the following sources:

- **New SQL Text:** Input/pasted by the user into a dialog box.
- **Extract from Database Object:** Use SQL DML statements from an appropriate schema object (such as views, procedures, triggers, functions, materialized views, and packages). Alternatively, objects can be dragged and dropped onto the Tuning Statements grid from the Data Source Explorer tree.
- **Import from File (Workspace):** Import statements from a file in an existing SQL Project.
- **Import from File (System):** Import statements from a file on a local disk or from a network share.

To view or modify a tuning statement, or to bind variable values within the statement, double-click on the statement in the grid or right-click and select **Edit...** from the context menu.

Oracle RAC Profiling

All instances of an Oracle Real Application Cluster (RAC) can be profiled using a single profiling session (and a single connection to the cluster). To enable this functionality for a RAC data source, you must access the **Profile Configurations** dialog, in one of the following ways:

- Right-click the data source in the Data Source Explorer and select **Profile As > Profile Configurations....**
- Click the down-arrow button next to the **New Profiling Session** button in the toolbar, and select **Profile Configurations....**

Within the **Profile Configurations** dialog, check the box for **Real Application Cluster (RAC) mode**.

Within a RAC profiling session, the filter bar at the top of the profiling editor includes an **Instances** filter, which allows you to view profiling results for individual instances in the cluster, or show results for sessions from all cluster instances.

Furthermore, several tabs from the **Profiling Details** view will show the ID of the instance on which the selected statement/event/session was observed.

Profiling Alerts

During profiling, alerts can be triggered based on thresholds set on the number of:

- Active sessions
- Sessions experiencing a wait condition
- Active sessions not on CPU

These alerts can produce a system notification (from the Windows status tray) and/or generate an email to be sent to one or more addresses. This allows you to begin a profiling session, continue with other work, and return to DB Optimizer when the alert has fired to investigate the underlying cause.

Alerts can be set up from the preferences page available from the menu item at **Preferences > Profile Alerts**. Alerts may be configured per data source and multiple alerts may be configured for a profiling session.

In order to generate email messages for alerts, the mail server options must be specified on the **Email Settings** page under the **Profile Alerts** preferences, and one or more recipients must be specified on the **Email Contacts** page.

Profiling Reports

The results of a profiling session can now be exported as an HTML or PDF document that can be saved for later, shared, and printed.

Within the Profiling Editor, the **Export Report** button is located in the top right, between the filter options and the **Help** button.

By default, the report is generated for the entire duration of the profiling session. To restrict the report to a particular time interval, click and drag on the profile session chart. To clear the selection, click once on the chart outside the selection.

Support for Common Table Expressions on Oracle and MS SQL Server

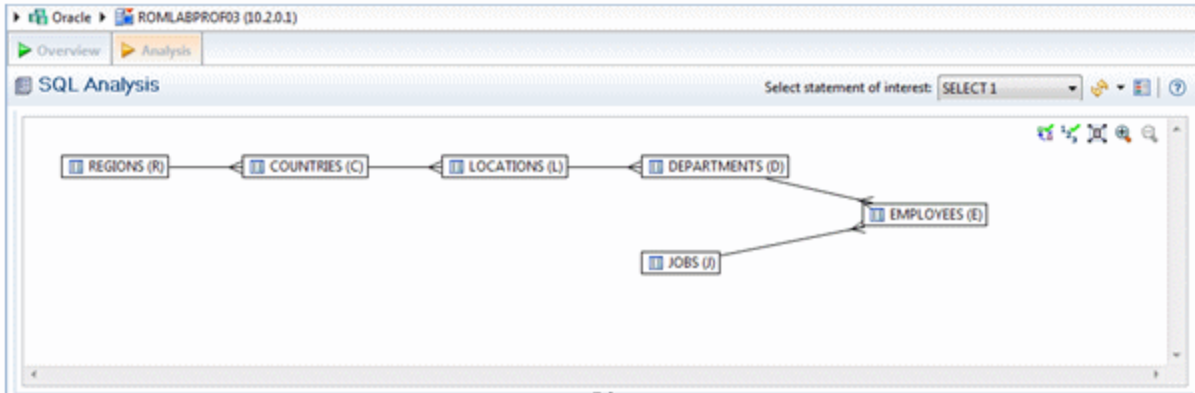
The **SQL Editor** and **Tuning Editor** now both support common table expressions (CTEs) for Oracle and SQL Server. This includes syntax highlighting and checking in the editor, as well as generation of explain plans for CTEs from the **SQL Editor** and the **Tuning Editor**.

Improved Visual SQL Tuning for Oracle

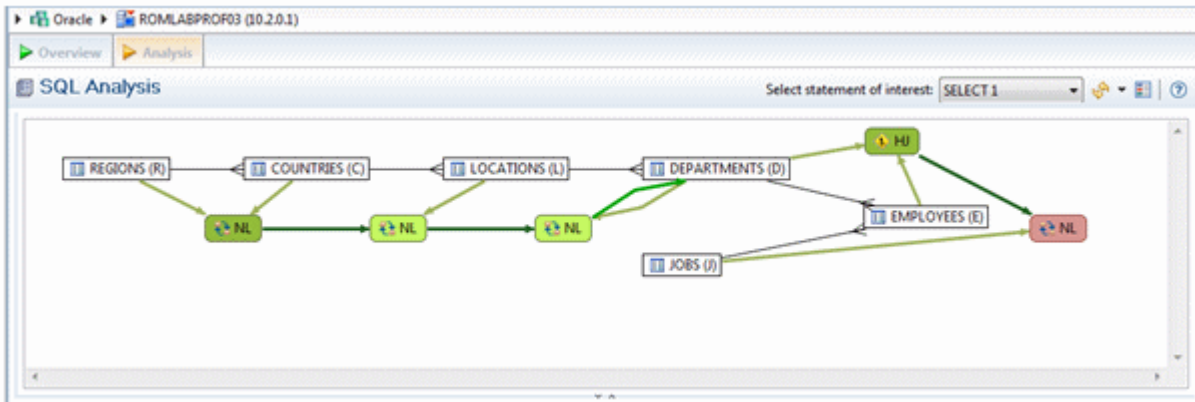
The **Visual SQL Tuning** (VST) diagram for Oracle 10g and 11g now has the option of showing Explain Plan details in addition to the column relationships. The **Explain Plan** display can be toggled on/off by clicking the leftmost button that appears when hovering over the **SQL Analysis** diagram.

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.

Compare the following VST diagram without **Explain Plan** overlay (largely unchanged from earlier versions):



Versus this VST diagram for the same query with the operations (four nested loops and a hash join) from the Explain Plan shown:



INSTALLATION

This section helps you install DB Optimizer. It is comprised of the following topics:

- [Installation Issues](#) on page 25
- [Installing DB Optimizer](#) on page 25

NOTE: For information on installing DB Optimizer XE, which includes Embarcadero AppWave, please see the *AppWave Quick Start Guide* and the *AppWave Administrator Guide* at docs.embarcadero.com/products/appwave.

TECHNICAL REQUIREMENTS

For information on the technical requirements of DB Optimizer, including supported operating systems and databases supported, see the Readme at docs.embarcadero.com/products/db_optimizer.

INSTALLATION ISSUES

DB Optimizer should be installed in a different directory than previous versions of DB Optimizer. If DB Optimizer is installed in the same directory, the following error will result in at application start up: "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 will result: "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

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

INSTALLATION ISSUES

DB Optimizer should be installed in a different directory than previous versions of DB Optimizer. If DB Optimizer is installed in the same directory, the following error will result in at application start up: "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 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 will result: "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

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

```
chmod 755 dbo_<version>.bin
```

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

```
./dbo_<version>.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 two options and 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 3.5 or higher or an Eclipse Based RCP application.
- 6 On the **Choose Install 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 click **Next**.

- 9 On the **Pre-Installation Summary**, review the installation details and 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 Embarcadero DB Optimizer** if you want to start the application for a first time use and click **Done**.

LICENSING DB OPTIMIZER

DB Optimizer supports the following types of licenses:

- **Trial licenses:** A license for a 14-day, full-featured trial version of the product. The trial license must be registered before you can use the product. For more information, see [Registering a Trial or Workstation License at Application Startup](#) on page 27.
- **Workstation or Network Named licenses:** A license or set of licenses that is tied to a particular workstation. The product can only be used on that workstation. For more information, see [Registering a Workstation License after Application Startup](#) on page 28.
- **Concurrent Licenses:** If you are using AppWave and DB Optimizer you can also have concurrent licenses which are floating licenses that allow a specific number of users to concurrently use a license, without tying the license to a specific workstation serial number. For more information, see the "Licensing" section of the *AppWave Administrator Guide*.

If you have problems registering, see [Registering by Phone](#) on page 28.

Registering a Trial or Workstation License at Application Startup

Shortly after downloading a trial version of an Embarcadero product, you should receive an email from Embarcadero with a serial number, which must be registered as follows. These instructions also apply if you have purchased a product and received a serial number prior to running the application for the first time, or if the trial period has expired.

- 1 Start the application.

The **Embarcadero License Registration** dialog appears.

- 2 Copy the serial number from the email, and then paste it into the **Serial Number** field.
- 3 Enter your Embarcadero Developer Network (EDN) account credentials in the **Login or Email** and **Password** fields.

If you have not previously created an EDN account, or have forgotten your password, click **I need to create ...** or **I've lost my password**.

- 4 Click **Register**.

The activation file should be downloaded and installed automatically. If this does not happen, click the **Trouble Connecting? Try Web Registration** link and follow the prompts. If you still have problems, see [Registering by Phone](#) on page 28.

Registering a Workstation License after Application Startup

The following instructions assume that you have received a workstation license by email and that you currently have a valid trial license. If you did not install a trial version or if the trial period has expired, follow the instructions in [Registering a Trial or Workstation License at Application Startup](#) on page 27 instead.

- 1 Click **Help > Launch License Manager**.
- 2 In the **License Manager** dialog, click **Serial > Add**.
- 3 Copy the serial number from the email, paste it into the **Add Serial Number** dialog, and then click **OK**.
- 4 Expand **Unregistered Serial Numbers**, right-click the serial number just added, and then click **Register**.

A registration dialog appears.

NOTE: The **Registration Code** box shows a machine-specific identifier required with other registration methods.

- 5 Ensure that **Register using Online Registration** is selected.
- 6 Provide Developer Network credentials in the **Login or Email** and **Password** fields.
- 7 Click **Register**.

The activation file should download and install automatically. If this does not happen, click the **Trouble Connecting? Try Web Registration** link and follow the prompts. If you still have problems, see [Registering by Phone](#) on page 28.

Registering by Phone

If you have problems with either of the above procedures, you can register licenses by phone. You must provide the following information:

- Embarcadero Developer Network credentials
- The registration code displayed in the **Embarcadero License Registration** dialog that appears when you start an unlicensed application
- The product base license serial number
- The license serial numbers for any additional features you have purchased.

For North America, Latin America, and Asia Pacific: call (415) 834-3131 option 2, and then follow the prompts. The hours are Monday to Friday, 6:00 A.M. to 6:00 P.M. Pacific time.

For Europe, Africa, and the Middle East: Call +44 (0)1628-684 494. The hours are Monday to Friday, 9 A.M. to 5:30 P.M. U.K. time.

Shortly after phoning in, you will receive an email containing an activation file. Then do the following:

- 1 Save the activation file to the desktop or to a scratch directory such as `c:\temp`.
- 2 From within the application, click **Help > Launch License Manager**.
- 3 Click **License > Import**.
- 4 Navigate to the directory where you saved the activation file, and then double-click the activation file.
- 5 Click the **Import** button to import the activation file, and then click **Finish**.

GETTING STARTED WITH DB OPTIMIZER

The following tutorial sessions enable you to get started using DB Optimizer. Their purpose is to provide you with the foundation needed to fully utilize the features and benefits of DB Optimizer, and apply them to your own enterprise. This section contains information on how to register data sources from your enterprise, and how to profile, analyze, and tune SQL statements, sessions, and events for the purpose of optimizing the efficiency of your data sources, as well as identify and prevent system bottlenecks and other wait-related issues.

This guide is comprised of the following sections:

- [Start DB Optimizer](#) on page 31
- [Overview of the User Interface](#) on page 32
- [Working with Data Source Explorer](#) on page 34
- [Profiling a Data Source](#) on page 39
- [Tuning SQL Statements](#) on page 48
- [SQL Code Assist and Execution](#) on page 68

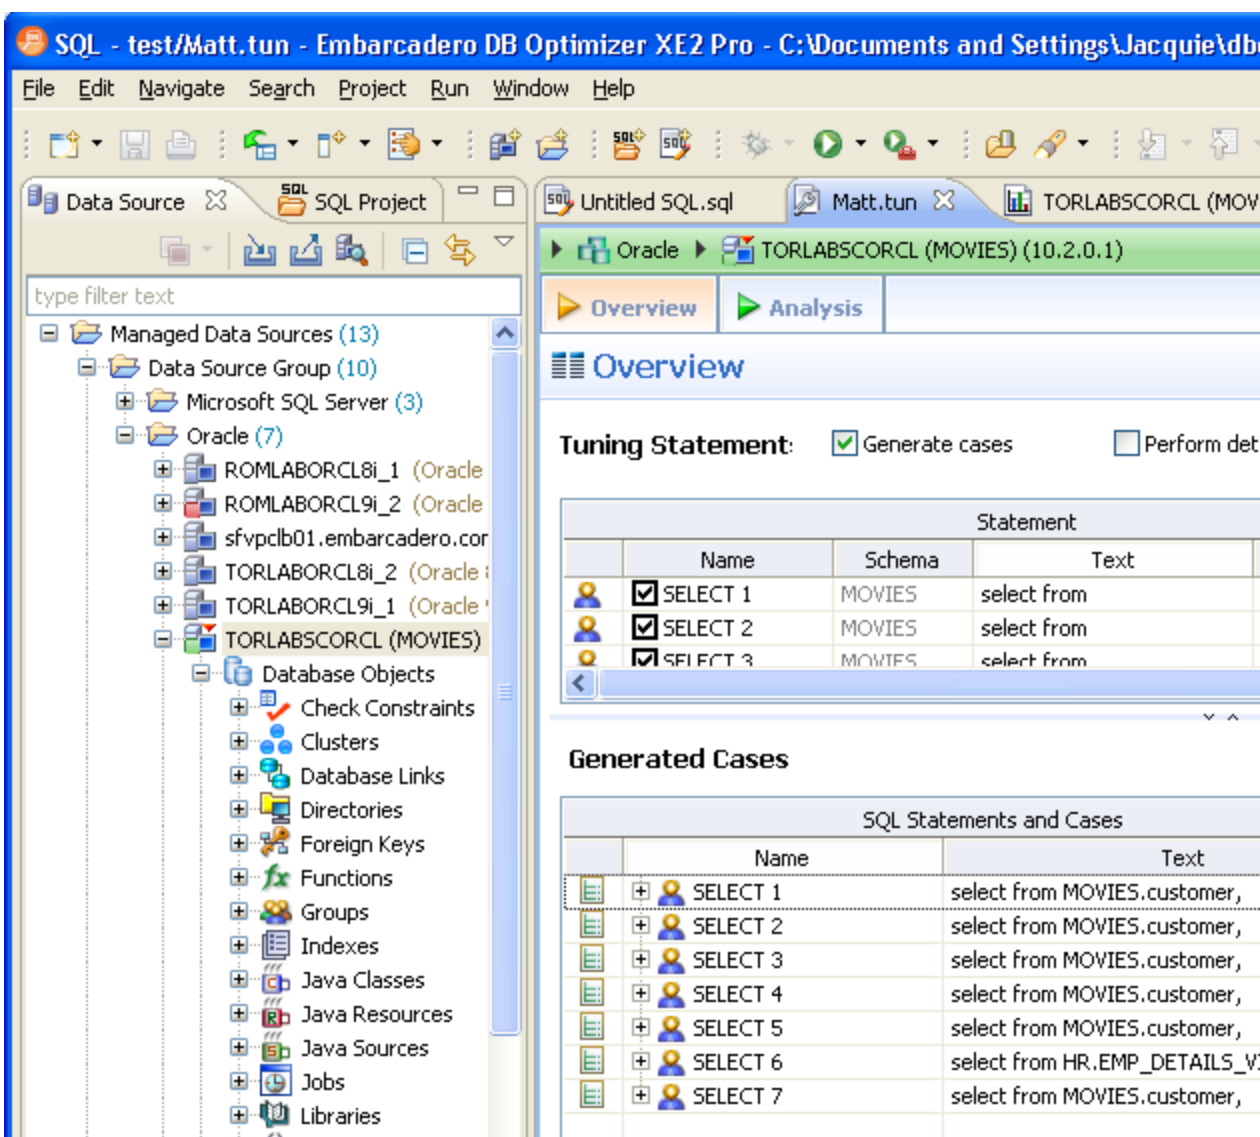
START DB OPTIMIZER

From the Desktop, choose the Windows Start menu and then select **Programs > Embarcadero > Embarcadero DB Optimizer X.X > Embarcadero DB Optimizer X.X**.

NOTE: Once you have started the application, you can select **Help** from the menu bar to find additional resources that complement and build upon the features and tasks presented in this guide.

OVERVIEW OF THE USER INTERFACE

The DB Optimizer application environment is known as the **Workbench** that provides an interface through which you manage data sources and analyze and tune statements. The Workbench is composed of common interface elements that provide tools to accomplish these tasks. These objects provide a uniform system for working with tuning jobs, profile sessions, and data sources.



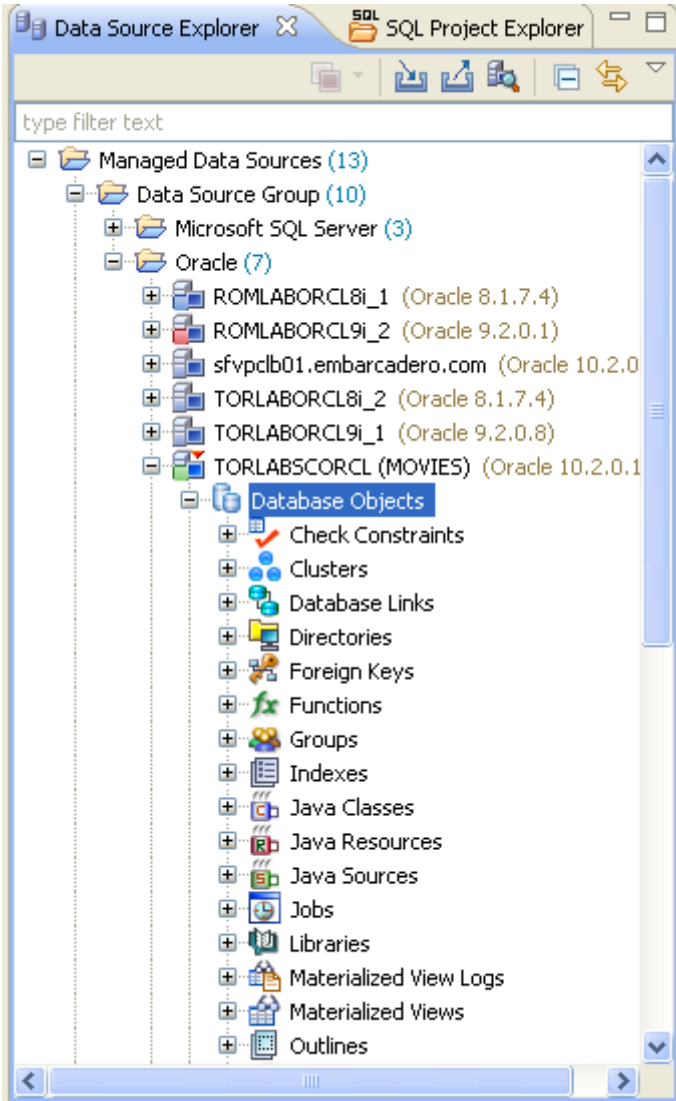
The Workbench provides an environment to build profiling sessions and tune query statements.

The **Workbench Space** is composed of views, editors, and the menu bar and command toolbar. Views and editors enable you to perform DB Optimizer functions and work tasks, as well as manage resources.

- **Views** are used to navigate a hierarchy of information, open Editors or display the properties of various application elements. For example, the Data Source Explorer view provides a tree of all data sources in the environment and the comparison jobs associated with each. You can launch these jobs directly from Data Source Explorer, modify the connection properties of data sources, or create and edit configuration archives from this View.
- **Editors** are used to access DB Optimizer functionality. For example, the **SQL Editor** provides a means to view, edit, and save SQL code. Editors are differentiated from Views in that they operate on an individual level rather than a supportive one. The **SQL Profiler** and **SQL Tuner** are also considered to be editors.
- The **Menu Bar** and **Command Toolbar** contain commands that execute various functional aspects of the application such as launching views and editors, navigation commands, and environment preference settings. The command toolbar contains icons that represent specific menu commands.

WORKING WITH DATA SOURCE EXPLORER

Data Source Explorer provides a tree view of all registered data sources. It breaks down the components of each data source and categorizes databases and corresponding database objects by object type and underlying SQL code. This feature provides you with a view of the contents of data sources in your enterprise in an easily navigated and cataloged interface.




Data Source Explorer sorts databases and database objects by category within DB Optimizer.

If data sources are particularly large or complex, or you are only developing specific objects, you can apply database object filters on the view in Data Source Explorer.

ADDING DATA SOURCES

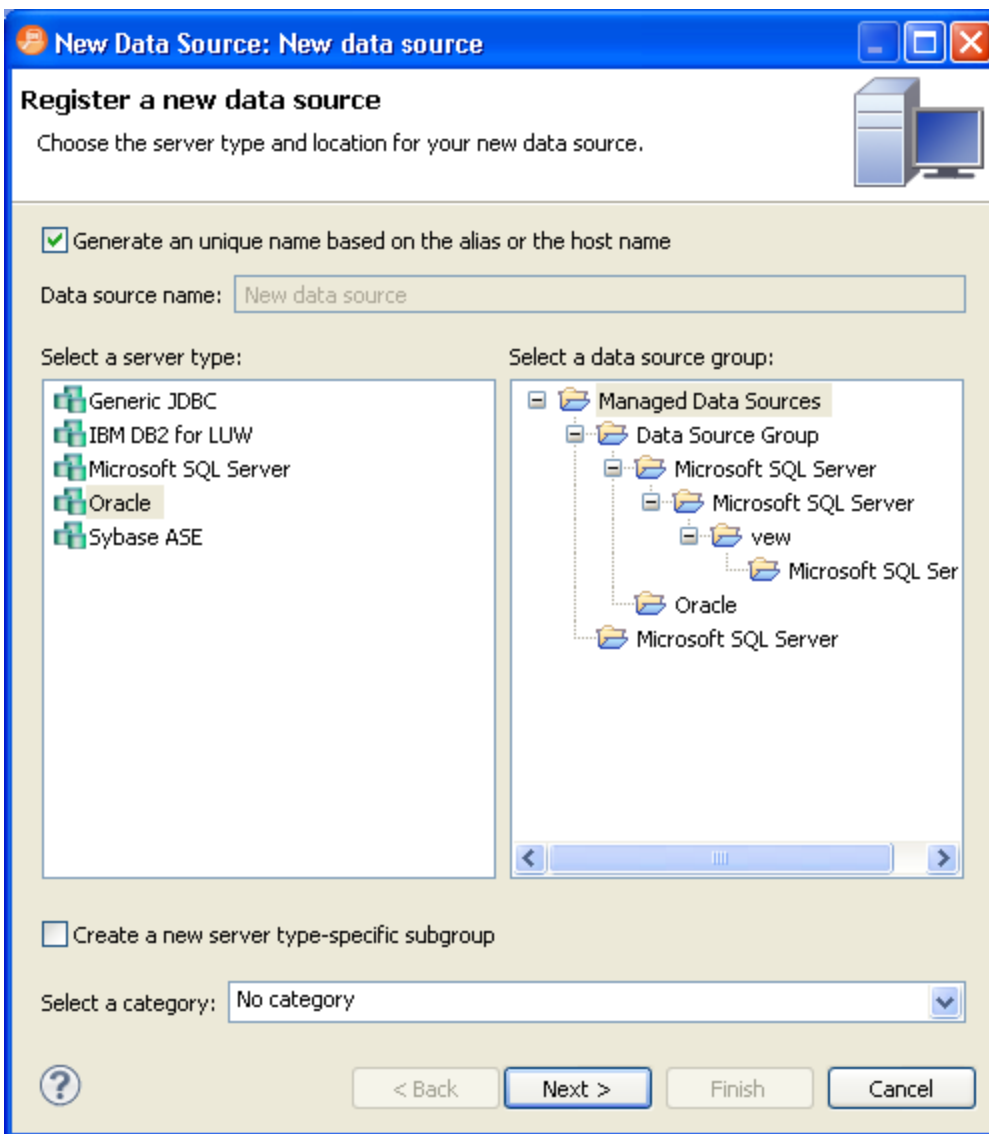
In order to profile and tune statements, you need to register the data sources to be analyzed in the environment by providing connection information and other details to DB Optimizer. Data sources are registered and managed in Data Source Explorer. Each time you register a new data source you need to specify its connection information and organize it in the View, as needed. Once a data source has been registered, it remains stored in a catalog and does not need to be registered again each time you open DB Optimizer. Furthermore, it can be used in multiple jobs, archived, or otherwise “shared” with regards to the general functionality of the application.



The Data Source Explorer view provides an organizational tree of registered data sources and the parameters associated with them.

ADD A DATA SOURCE

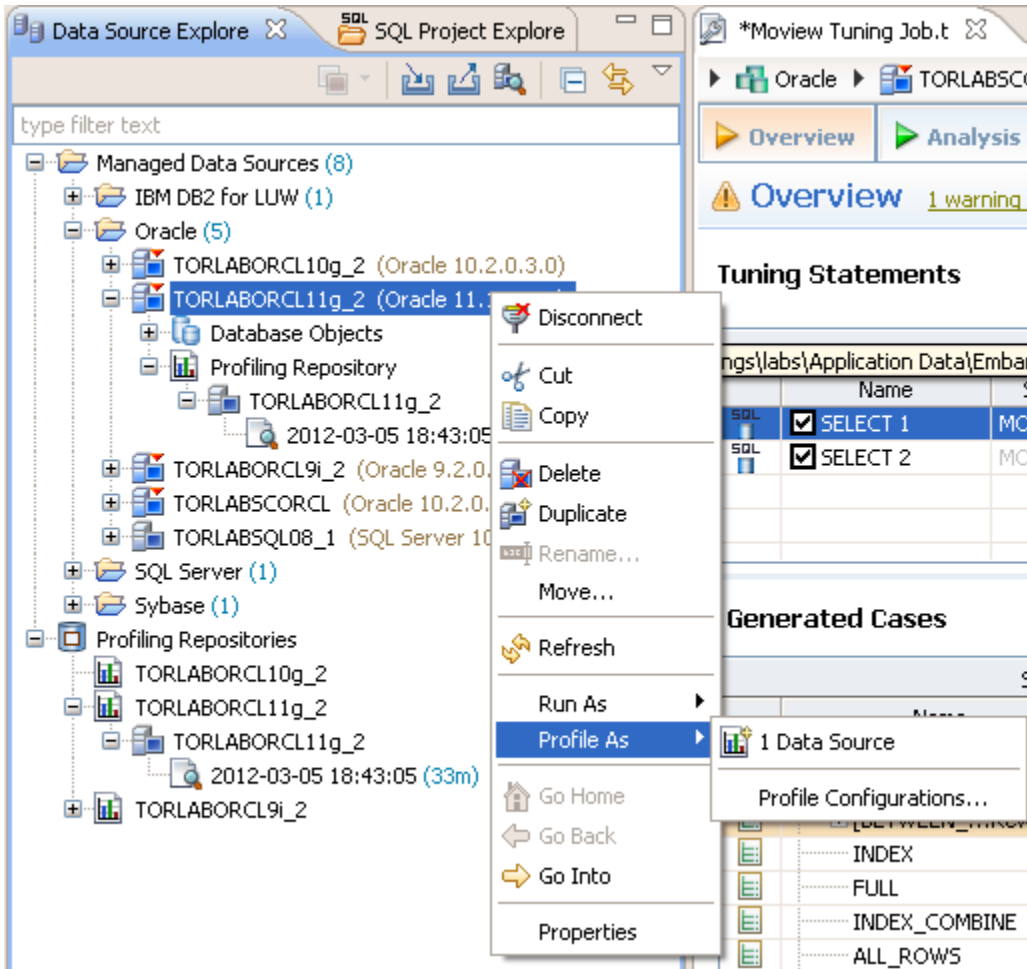
Select **Data Source Explorer** and choose **File > New > Data Source** from the Menu bar. The **New Data Source** Wizard appears. Follow the prompts to register the data source.



TIP: You can categorize data sources in order to color code them in the **Data Source Explorer** which makes the data source easier to find when you have a long list of data sources. The data source will be decorated with the color associated with the category selected for the data source. Categories include Development, Test, QA, and Production. Existing data sources can be categorized using the **Data Source Properties** page; right-click a disconnected data source and select **Properties**.

BROWSING DATA SOURCES

To analyze and tune statements on data sources within your enterprise, you must access the data source objects within the application. It is important to view databases and underlying code in an organized manner, especially when maintaining a large system. Data Source Explorer's tree structure can be used to view databases, tables, and other information about data sources. Expand the nodes of each registered data source to view details about each data source. Data Source Explorer is also used to launch profiling sessions, by enabling you to select the data source that you want to execute, and then launching a profiling session from your selection.



The Data Source Explorer tree provides a list of databases, and enables you to launch Optimizer features, such as SQL Profiler, via the context menu.

All objects listed in Data Source Explorer are sorted by data source name. Additionally, if you double-click on an object, SQL Editor will open and display the underlying SQL code.

To browse a data source

Expand the node of the data source that you want to examine. Double-click an object to view the code in SQL Editor.

Notice in the previous screenshot, there are Profiling Repositories listed in the Data Source Explorer. For Oracle data sources only, you can save your profiling sessions to the Profiling Repository for later examination. We'll discuss this more later in "[Saving a Profiling Session](#)" on page 46.

PROFILING A DATA SOURCE

NOTE: The SQL Profiler is not available in the “Basic” version of DB Optimizer.

SQL Profiler samples a data source and builds a statistical model of the load on the database. Profiling is used to locate and diagnose problematic SQL code and event-based bottlenecks. Profiler also enables you to investigate execution and wait time event details for individual stored routines. Results are presented in the profiling editor, where you can identify problem areas and view individual SQL statements, as needed.



SQL Profiler analyzes and provides a statistical model of database load that identifies problematic code and event-based bottlenecks.

SQL Profiler is composed of the following diagnostic parts:

- The **Load Graph** provides a display of the overall load on the system. The bars represent individual aspects of the data source.


- The **Top Activity** section displays where load originates. Specifically, Top Activity displays the top SQL statements, the top events the database spends time in, and the top activity sessions that may be causing issues in terms of query time and response.
- The **Profiling Details View** displays detailed information for any item selected from Top Activity. For example, examining a SQL statement from Top Activity displays the identification parameters and the execution statistics of that specific statement selection.

STARTING A PROFILING SESSION

In order to access SQL Profiler, you need to run a profiling session on a data source registered in Data Source Explorer.

To run a profiling session

In Data Source Explorer, right-click on a data source and select **Profile As > Data Source**. The profiling session begins.



The Progress view indicates that the profiling session has been initiated.

Once a profiling session launches, it runs until you stop it or it has run for its specified length of time. You can stop a session by clicking the **Stop** icon on the right-hand side of the Profiling progress bar.

When the profiling session is complete, the first two sections (**Load Chart** and **Top Activity**) will be populated with information about the database load. You can then begin analyzing the data and identifying problem areas.

ANALYZING SESSION DATA


The SQL Profiler is composed of three essential diagnostic views that provide information about load on a particular database in the system. These views enable you to identify bottlenecks and view details about specific queries that execute and wait over the course of a profiling session.

The following describe the three components of SQL Profiler, in descending order of granularity:

- [Load Chart](#) on page 41
- [Top Activity](#) on page 41
- [Profiling Details](#) on page 43

LOAD CHART

The Load Chart provides an overview of the load on the system, and is designed as a high level entry point to reading session results. The colored bars represent individual aspects of the database, and the graph can be used to discover bottlenecks.



The Load Chart displays the overall load of the database that you analyzed with Profiler.

Time is displayed on the X axis, and the Y axis shows the average number of sessions waiting or executing. Each support platform type has a specific set of wait event times. For example, Sybase platforms will display CPU, Lock, Memory, I/O, Network, and Other. Use the chart legend to understand the graph. It displays a color and code scheme for executing and waiting session categories in the upper right-hand corner of the chart.

TOP ACTIVITY

Below the Load Graph, the Top Activity section displays where the load originates and outlines the top SQL statements, top events, and the top activity sessions on the database. It is composed of a series of tabs that provide detailed statistics on individual SQL statements and sessions that are waiting or executing over the length of the profiling session.

U...e	Program	SID	S...#	Act...%)	Machine	Ses...pe	Cli...nfo
SYSTEM		259	1865	91.02	rolcgrigore01	USER	10....52
SYSTEM		285	2193	82.04	rolcgrigore01	USER	10....52
SYSTEM		251	1809	57.49	rolcgrigore01	USER	10....52
SYSTEM		286	566		rolcgrigore01	USER	10 52

The Top Activity section displays a more detailed view of the Load Graph. It identifies top statements, events, and sessions that are waiting or executing over the length of the profiling session.

- The **Overview** tab provides summary information about SQL statements and events and their activity levels, and sessions, their system process IDs and their activity level. You can reorder the rows in any of the three sections of this tab. For example, clicking the **Event** column in the **Events** section changes the alphabetical order to ascending or descending.

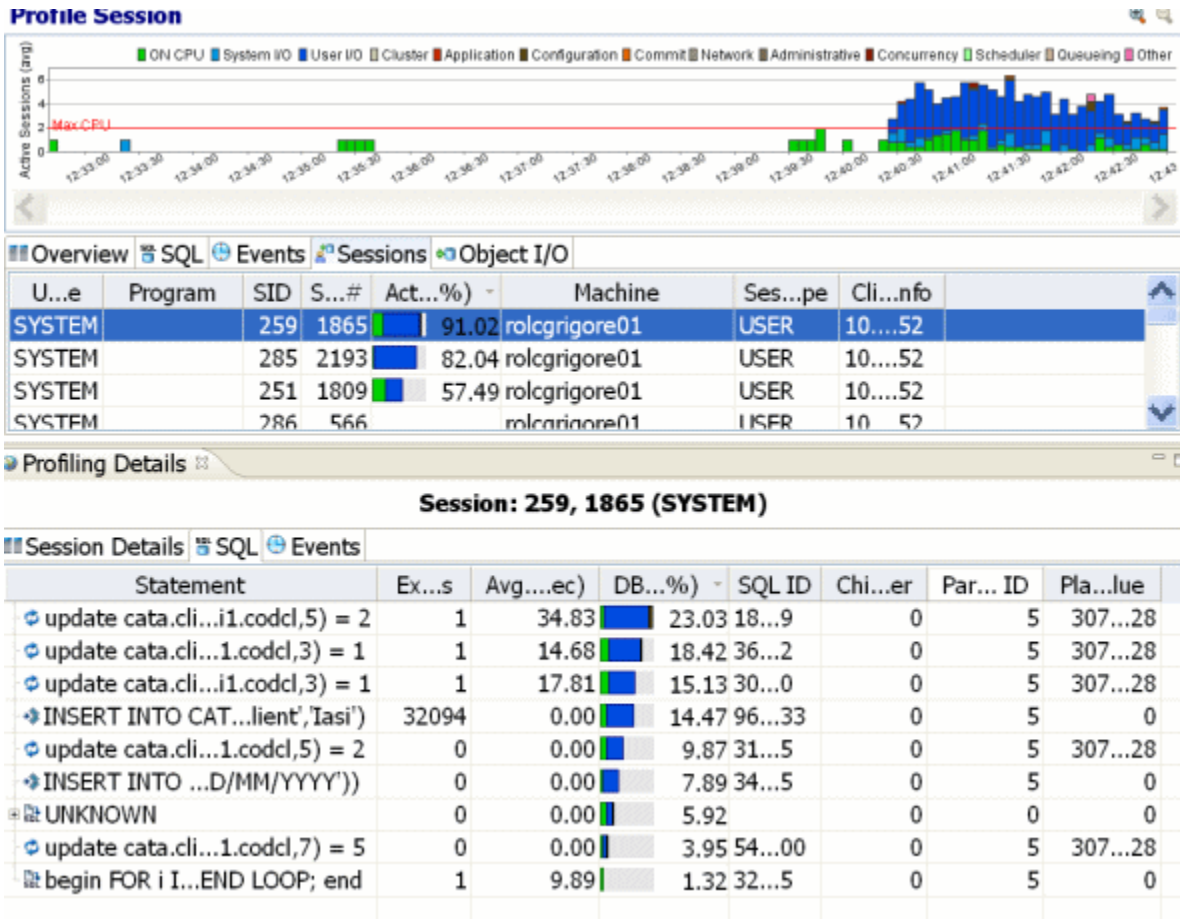
- The **SQL** tab provides information about SQL statements and procedures. This includes all INSERT, SELECT, DELETE, and UPDATE statements that are executing or waiting to execute over the length of the profiling session.
- The **Events** tab displays information about wait events, and should be used 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 the top events are related to database configuration, then the database setup should be investigated.
- The **Sessions** tab displays information about sessions, and can be used to discover sessions that are very active or bottlenecked.
- The **Object I/O** tab provides information about I/O when profiling an Oracle data source.
- The **Procedures** tab provides information about the procedures running in the session when profiling a SQL Server or Sybase data source.

PROFILING DETAILS

When you select any item from Top Activity, details are displayed on the **Profiling Details** view.

TIP: You may have to select **Windows > Show View > Profiling Details** to display the Profiling Details view.

The tabs that compose the **Profiling Details** view are dependent on the nature of the object selected in Top Activity, in order to reflect that item's specific information. The tabs are also dependent on the data source platform. For example, the **Object I/O** tab is available only for the Oracle platform whereas the Procedures tab is available only for the SQL Server and Sybase platforms.



Profiling Details displays detailed parameter information on a statement, event, or session that was selected in Top Activity. The data displayed also varies by database platform.

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.

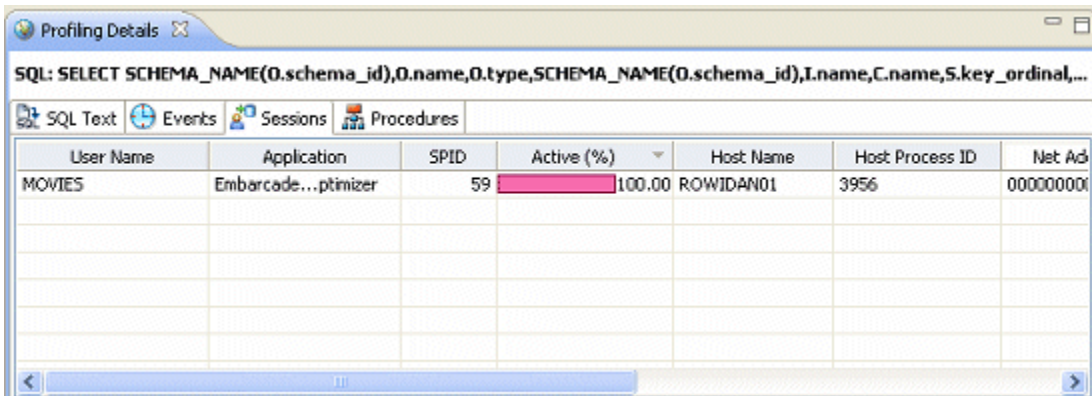
Depending on the top activity selected and the profiled platform types, some tabs may not be available.

NOTE: When right-clicking on a SQL statement in the Top Activity section in the SQL Profiler, if the SQL statement is run by a different user than the user who is running DB Optimizer, then 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 the tuning job under the original user.

VIEW SESSION DETAILS

- 1 In the **Profile Session** area of the **SQL Profiler**, in the **Sessions** column click anywhere in the row of an application that ran during the profiling session.
- 2 In the **Profiling Details** area, click the **Sessions** tab.

Details of the session are displayed.



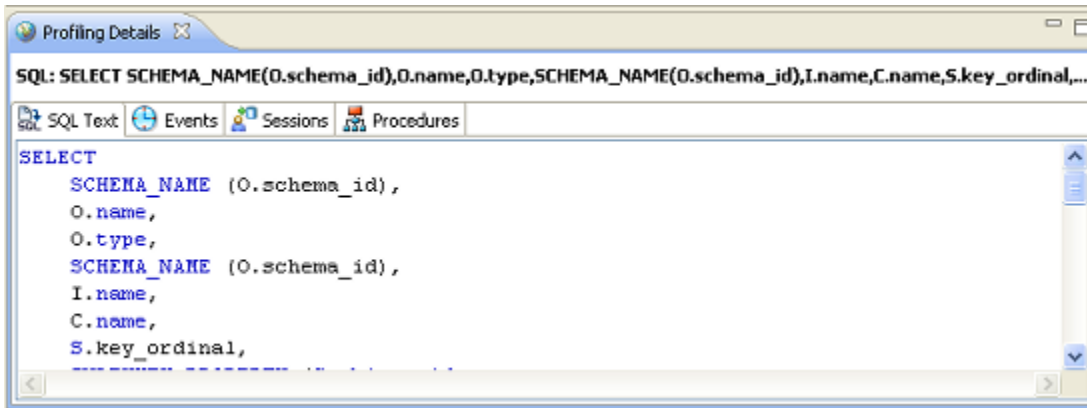
The screenshot shows a window titled "Profiling Details" with a SQL query at the top: `SQL: SELECT SCHEMA_NAME(O.schema_id),O.name,O.type,SCHEMA_NAME(O.schema_id),I.name,C.name,S.key_ordinal,...`. Below the query are four tabs: "SQL Text", "Events", "Sessions", and "Procedures". The "Sessions" tab is selected, displaying a table with the following data:

User Name	Application	SPID	Active (%)	Host Name	Host Process ID	Net Ad
MOVIES	Embarcade...ptimizer	59	100.00	ROWIDAN01	3956	00000000

VIEW SQL

- 1 In the **Profile Session** area of the **SQL Profiler**, click anywhere in the row of an application that ran during the profiling session.
- 2 In the **Profiling Details** area, click the **SQL Text** or **SQL** tab.

The SQL text for the application that selected object displays.



TIP: From the SQL tab you can easily tune a statement by right-clicking a statement on the SQL tab to initiate the tuner, which then opens with the selected statement in the Ad hoc SQL tab of the Tuner Input.

SAVING A PROFILING SESSION

A profiling session can be saved to a file with a `.oar` suffix that contains the name of the data source. When profiling Oracle data sources you can alternatively save the session data to the Profiling Repository. This enables you to open the file at a later time for analysis.

SAVE A PROFILING SESSION AS A .OAR FILE

Select the profiling 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 project is added to SQL Project Explorer.



Once you have saved a profiling session, 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.

SAVE A PROFILING SESSION IN THE PROFILING REPOSITORY

You can also have the session saved automatically as a Profiling Repository file in an Oracle data source. This allows you to profile your data source for a longer period of time, for days or even weeks.

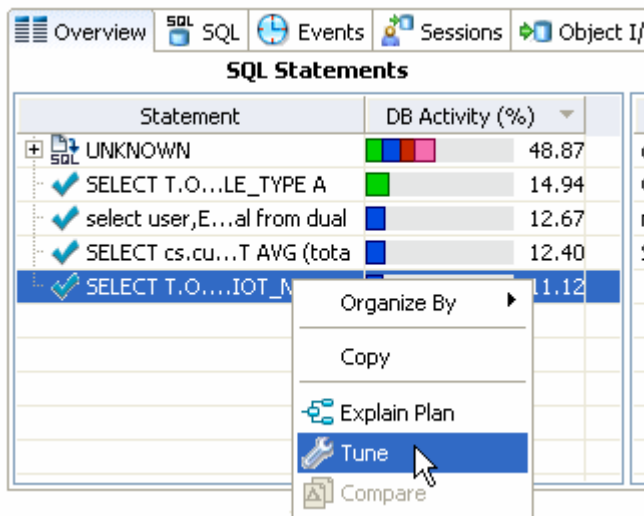
Right-click the data source you want to profile and select **Profile As > Profile Configurations**.

In the **Profile Configurations** dialog, select **Save to data source** and then select the data source where you want to save the profile sessions.

The saved session file, named using the date and time when the profiling session was stopped, can be viewed from the Data Source Explorer and thus available to share with other DB Optimizer users.

IMPORTING STATEMENTS TO SQL TUNER

SQL Profiler enables you to submit one or more statements into SQL Tuner. This enables you to take advantage of Tuner's hint-based and transformation-based suggestions if you want to tune a problem statement that you detected over the course of a profiling session.



Context menu commands in SQL Profiler enable you to import statements to a tuning job directly from the Profiler interface.

Import a statement from SQL Profiler into SQL Tuner

Select one or more statements in Profiler, right-click and select **Tune** from the context menu. SQL Tuner opens and contains the selected statements in a new tuning job. You can now proceed to tune the problematic statements.

TUNING SQL STATEMENTS

SQL Tuner 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. Tuner enables the optimization of poorly-performing SQL code through the detection and modification of execution paths used in data retrieval. This is primarily performed through hint injection, and index and statistics analysis.

The screenshot displays the Oracle SQL Tuner interface. The 'Overview' tab is active, showing the 'Tuning Status' section with three checked options: 'Generate cases', 'Perform detailed analysis', and 'Execute each generated case' (set to 3). Below this is a table of statements being analyzed:

Statement						Time	
Name	Schema	Text	Tables	Views	Elapsed (s)	Improved (s)	
SELECT 1	MOVIES	select from	5	0	0.63	0.63	
SELECT 2	MOVIES	select from	4	0	0.04	0.04	
SELECT 3	MOVIES	select from	8	0	0.09	0.09	
SELECT 4	MOVIES	select from	5	0	10.10	10.10	
SELECT 5	MOVIES	select from	0	0	0.00	0.00	

Below the statements table is the 'Generated Cases' section, which shows a table of SQL statements and their generated cases:

SQL Statements and Cases			Cost	Execution Statistics
Name	Text	Value	Elapsed Time (s)	Physical
SELECT 1	select from MOVIES.customer,	21.0	0.63	
SELECT 2	select from MOVIES.customer,	11.0	0.04	
SELECT 3	select from MOVIES.customer,	25.0	0.09	
SELECT 4	select from MOVIES.customer,	22.0	10.10	
SELECT 5	select from MOVIES.customer,	25.0	0.08	
SELECT 6	select from HR.EMP_DETAILS_VIEW			
SELECT 7	select from MOVIES.customer,	21.0	0.07	


Tuner analyzes specified SQL statements and then supplies execution path directives. This enables you to select alternate paths for queries, thus optimizing system performance based on analysis.

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 will select alternate paths, and enable you to change the original path to one of the alternates. Execution paths slower than the original path 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.

CREATING A NEW TUNING JOB

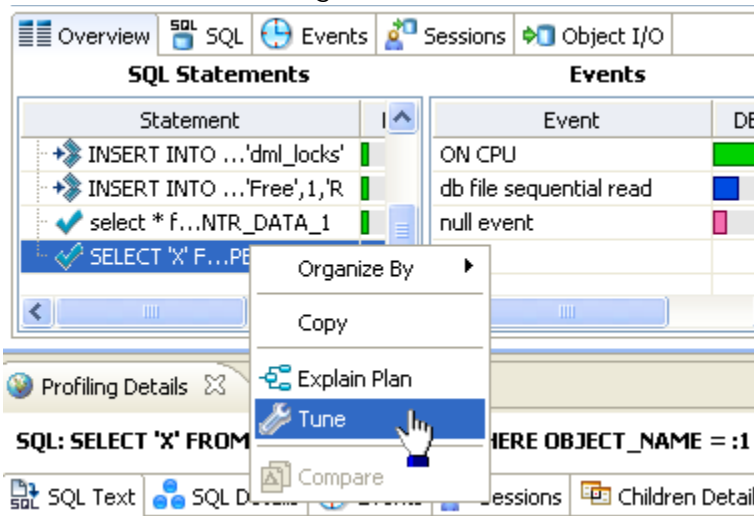
New tuning jobs are created from scratch where you can specify the statements to be tuned from a variety of sources, or statements can be directly imported from existing profiling sessions on data sources currently registered in the environment.



Tuning jobs are defined in SQL Tuner by specifying the data source and corresponding statements to be tuned and then executing the tuning job process.

CREATE A NEW TUNING JOB

If you determined from a profiling session that a specific SQL statement should be tuned, in the Profile Session, right-click the statement and select **Tune** as follows:



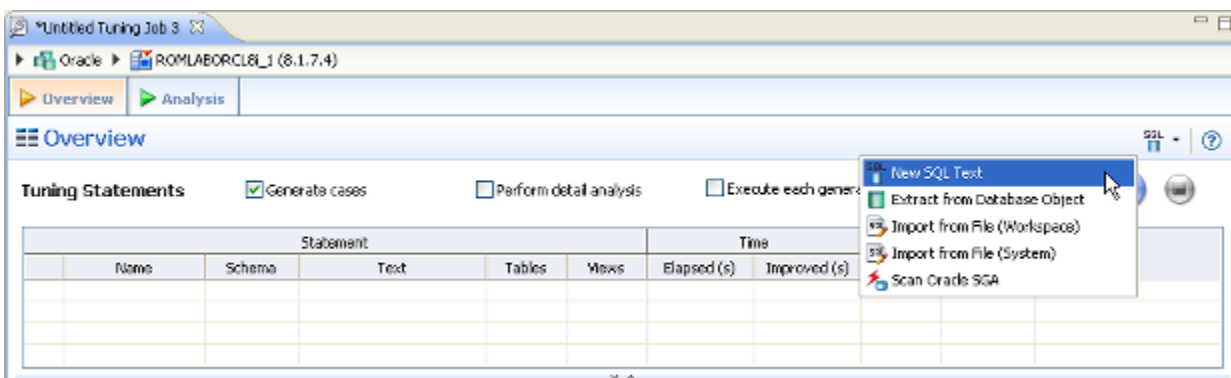
OR

Select **File > New > Tuning Job**, or click the New Tuning Job icon on the Toolbar. SQL Tuner opens and you can proceed to set up the parameters of the new job.

Once you have defined the input of the tuning job, you can save the file with a `.tun` suffix via the **Save As...** command. The job is added to the SQL Project Explorer and can be re-opened and re-run at any time once it is saved to the system.

ADDING SQL STATEMENTS TO A TUNING JOB

Once you have created a name for the tuning job and indicated its source, you can add SQL statements that you want the job to tune. Statements are added to a job via the options available in the SQL list on the Overview tab. All standard DML statements (SELECT, INSERT, DELETE, UPDATE) are viable for the tuning procedure.



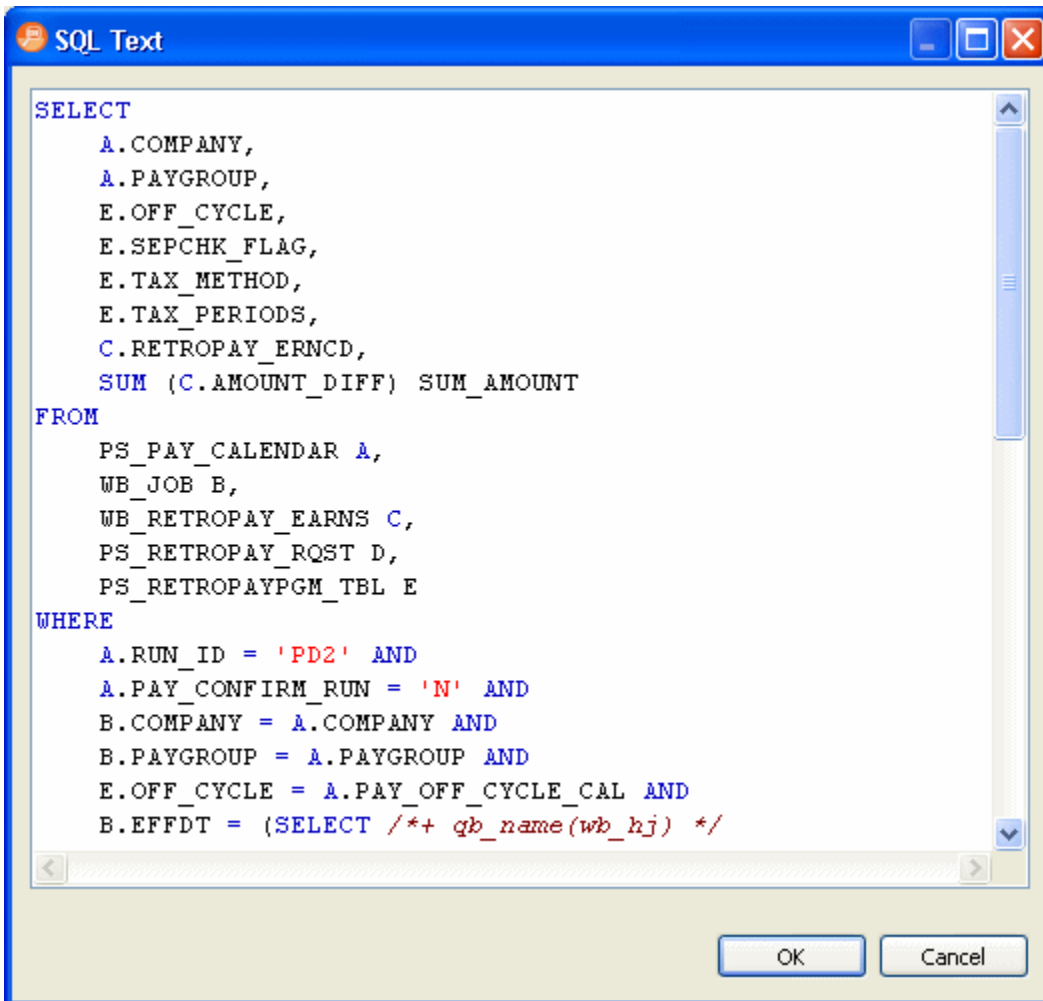
Statements are added to a tuning job via the Overview tab.

There are four or five different ways to add SQL statements to a job, as reflected by the options available in the SQL list as shown previously:

- **New SQL Text:** Copy/paste SQL statements to the New SQL Text window or write queries by hand. You can also paste a statement directly into the Overview statement grid by copying the statement and then right-clicking anywhere on that grid and selecting **Paste**.
- **Extract from Database Object:** Drag and drop database objects from the Data Source Explorer to the **Statements** grid on the **Overview** tab.
- **Import from File (Workspace)** and **Import from File (System):** Browse the workspace or file system and select SQL files.
- **Scan Oracle SGA:** For the Oracle platform only, you can also scan the System Global Area (SGA) for statements to tune.

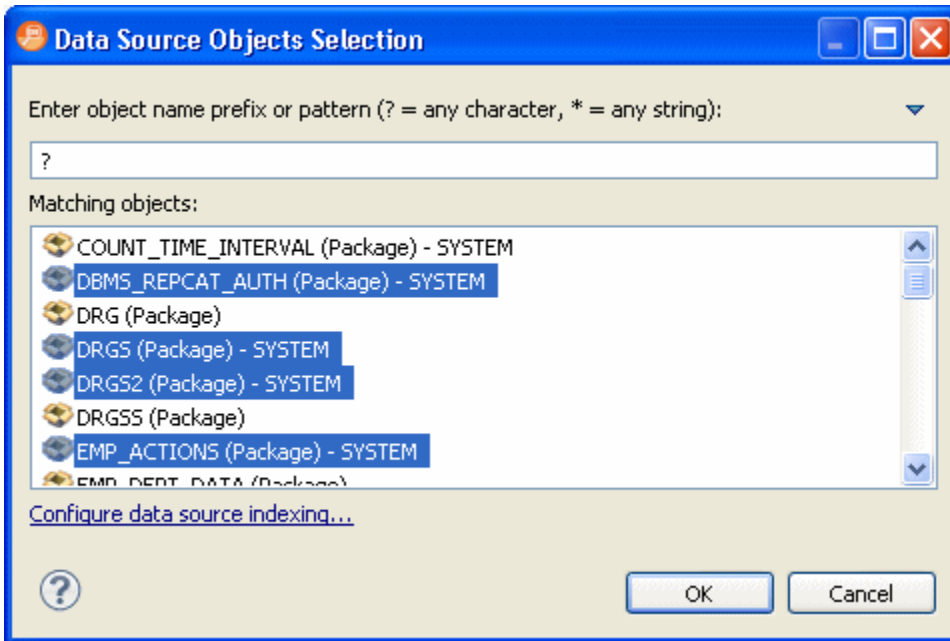
ADD NEW SQL TEXT

From the **SQL** list, select **New SQL Text** and manually type an SQL statement in the window. Alternatively, copy/paste the statement from another source.



ADD A DATABASE OBJECT

From the **SQL** list, select **Extract from Database Object**. The **Data Source Objects Selection** dialog appears. Type an object name prefix or pattern in the field provided and choose a statement from the window as it populates to match what you typed.



TIP: You can also drag an object, such as a Packages, Package Bodies, Views, and Functions, from the Data Source Explorer to the grid on the **Overview** tab.

NOTE: In order to drag a database object into the statement grid on the Overview tab, the database object in the Data Source Explorer must be in the same database that displays in the bread crumbs at the top of the Tuning window.

ADD A SAVED SQL FILE

From the **SQL** list, select **Import from File (Workspace)** or **Import from File (File System)**, depending on where the file you want to add is stored. Select a file from the dialog that appears and it will be added to the tuning job.

ADD ACTIVE SQL IN SGA

From the **SQL** list, select **Scan Oracle SGA**. Specify any filters as required and then click **Next**. From the list of SQL statements retrieved from the SGA, select those you want to optimize, and then click **Finish**. The selected statements are copied to the tuning job.

RUNNING A TUNING JOB

After you add SQL statements to the job, click the **Overview** tab. Once you choose your tuning options and click the **Run Job** icon, the DML is parsed from the statements and added to the **Generated Cases** area. The Generated Cases are alternative execution paths or explain paths that could be better or worse than the default path the database uses. When these cases are executed you can use the execution statistics to determine which case would optimize performance.

Each extracted statement is listed by Name and Text. Additionally, each statement has a Cost, Elapsed Time, and Other Execution Statistics value that provide information on how effectively each case executes on the specified data source. These parameters let you compare the efficiency of the original statements to the cases generated by the tuning process when it is executed.

TIP: You can double-click a generated case to view or edit the SQL source of the statement.

The screenshot displays the DB Optimizer interface with several key components and annotations:

- Tuning Source Status:** Includes checkboxes for 'Generate Cases', 'Perform detail analysis', and 'Execute each generated case'. A 'Run/Cancel Job Controls' button is visible on the right.
- Statement Table:** A table with columns: Name, Schema, Text, Tables, Views, Elapsed (s), Improved (s), Cases, and Index. It lists 'SELECT 2' and 'SELECT 1' with their respective execution times and case counts.
- Generated Cases Table:** A table with columns: Name, Text, Cost, Elapsed Time (s), Physical Reads, and Logical Reads. It lists various cases like 'SELECT 2', 'Missing a ...sformation', 'SELECT 1', and several 'LEADING' and 'INDEX' cases.
- Annotations:**
 - Tuning Status Indicator:** Points to the checkmarks in the 'Statement' table.
 - Enable Execution Check Box:** Points to the 'Execute each generated case' checkbox.
 - Schema and Catalog Selectors:** Points to the 'Schema' dropdown menu.
 - Pane Size Control Increase/Decrease:** Points to the vertical scrollbar.
 - Column Set Expand/Collapse Control:** Points to the '+' and '-' icons in the 'Generated Cases' table.
 - Filter Control:** Points to the filter icon in the 'Generated Cases' table.
 - Generated Case Expand/Collapse Control:** Points to the '+' and '-' icons in the 'Statement' table.
 - Transformation Case:** Points to the 'Missing a ...sformation' case.
 - Hint-Based Cases:** Points to the 'LEADING' and 'INDEX' cases.
 - Extracted SQL Statements:** Points to the 'Text' column in the 'Generated Cases' table.

The Generated Cases table enables you to measure the various load costs of the original tuning statements and generated cases for each, which suggest alternative query paths to optimizing your data source.

The **Tuning Status Indicator** provides the status of each statement or case, and indicates if they are ready for execution. In some cases, SQL code may need to be corrected or bind variables may need to be set prior to executing statements. When you try to tune a statement containing a bind variable you are now warned that either the type is not set or the value is not set.

Use the check boxes to select which statements and cases you want to run and then click the **Run** icon in the upper right-hand corner of the screen. The **Execute each generated case** field enables you to execute each selected statement or case.

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 select **Refresh Tuning Statements**.

Once you have executed a tuning job, the **Generated Cases** tab will reflect SQL Tuner's analysis of the specified statements. Once these have been analyzed, you can proceed to modifying the Tuner results and applying specified cases on the data source to optimize its performance.

RUN A TUNING JOB

- 1 Once you have a SQL statement that is a tuning candidate, navigate to the **Overview** tab.
- 2 In the **Tuning Statements** area, select the checkbox next to the Statement name that you want to analyze and then:

*To analyze the SQL statements, click **Generate cases**.*

*To perform the analysis that populates the Analysis tab now, click **Perform detail analysis**. Otherwise, this analysis is performed when you click the Analysis tab.*

*To have the system generate execution statistics, click **Execute each generated case** and then select the number of times 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.*


- 3 Then click the **Run Job** icon at the top right-hand side of the window. The tuning job runs, analyzing each statement and case, and providing values in the appropriate columns.

ANALYZING TUNER RESULTS ON THE OVERVIEW TAB

When you have executed a tuning job, the **Generated Cases** area of the **Overview** tab reflects Tuner's analysis of the specified statements and cases. The **Generated cases** create alternative execution or explain paths that could be more or less efficient than the default path the databases uses. Executing these cases provides the statistics necessary to optimize performance.

Once a tuning job has executed, use the **Cost** and **Execution Statistics** value columns to determine the fastest execution path for each statement. The **Cost** column shows the performance cost of an execution path as determined by the database. The **Execution Statistics** are the actual results of running the SQL statement using the generated case. This is where DB Optimizer can help you find where the database default path is actually not the optimal path. The **Elapsed Time(s)** and **Results** columns can more accurately show the most efficient execution path.

In the **Cost** and **Execution Statistics** columns, the values of the original statement are considered to be the baseline values. A **Cost** column can be expanded to provide a graphical representation of the values for statements and cases. Similarly, the **Execution Statistics** column can also be expanded to display a graphical representation of values as well. The bar length and colors are intended as an aid in comparing values, particularly among cases.



Case query times based on the original statement can be represented as colored bars on the **Generated Cases** area of the **Overview** tab to help you determine the fastest execution path for the given selections.

The baseline value of the original statement spans half the width of the column, in terms of bar length. For cases of the original statement, if one or more cases show a degradation value, the largest value will span the width of the column. Bar lengths for all other cases will then be displayed in comparison length to the highest degradation value.

The cost and execution results are color-coded as follows:

- **Light blue:** These cases are within the degradation and improvement threshold. Applying these changes may marginally improve or degrade the efficiency of the SQL statement.
- **Green:** These cases have values less than the improvement threshold. There is a high probability that changing the SQL statement with this alternative execution path will improve efficiency.
- **Red:** These cases have values higher than the improvement threshold. Implementing these changes will degrade the efficiency of the SQL statement.

DETERMINE THE BEST CASES FOR STATEMENT EXECUTION PATH TIME

Once the tuning job has executed, view the **Generated Cases** area of the **Overview** tab and determine the best possible case in terms of the **Execution Statistics** column values. This will indicate the most optimized query path for a given statement. Once you have determined the best case, you can execute that case on the specified data source and alter the database code to run the statement as that case on the native environment.

If you don't find an acceptably fast path, go to the **Analysis** tab. The **Analysis** tab can identify missing indexes and by examining the diagram you may be able to determine if there is something wrong with the SQL or schema.

FINDING MISSING INDEXES AND SQL PROBLEMS

The tuner performs the index and SQL analysis as part of the tuner run job performed on the **Overview** tab if **Perform Detail Analysis** is selected. Otherwise, the analysis is performed when you click the **Analysis** tab.

The screenshot shows the Oracle SQL Analysis tool interface. The top navigation bar includes "Overview" and "Analysis" tabs. The main window is titled "SQL Analysis" and displays the following:

- SQL Statement: `UPDATE MOVIERENTAL XXX SET TOTALCHARGE = (SELECT FROM VHEF)`
- Visual SQL Tuning Diagram (VST): A graphical representation of the query execution plan. It shows a root node "RENTALITEM (ITM)" connected to "MOVIERENTAL (XXX)" and "MOVIECOPY (COP)". "MOVIECOPY (COP)" is further connected to "MOVJETITLE (TIT)".
- Index Analysis Tab: A table listing indexes for the tables involved in the query.

A red arrow points to the "Create Index Button" in the Index Analysis tab.

Index Name	Table Owner	Table Name	Index Type
MOVIECATEGORY_PK	MOVIES	MOVIECATEGORY	CATEGOR
MOVIECOPY_PK	MOVIES	MOVIECOPY	MOVIECC
MOVJETITLE_PK	MOVIES	MOVJETITLE	MOVIEID
RENTALITEM_PK	MOVIES	RENTALITEM	RENTALII

DB Optimizer can 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. The VST diagram can be displayed in either Summary Mode or Detail Mode. The VST diagram 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 troubleshooting and analysis.

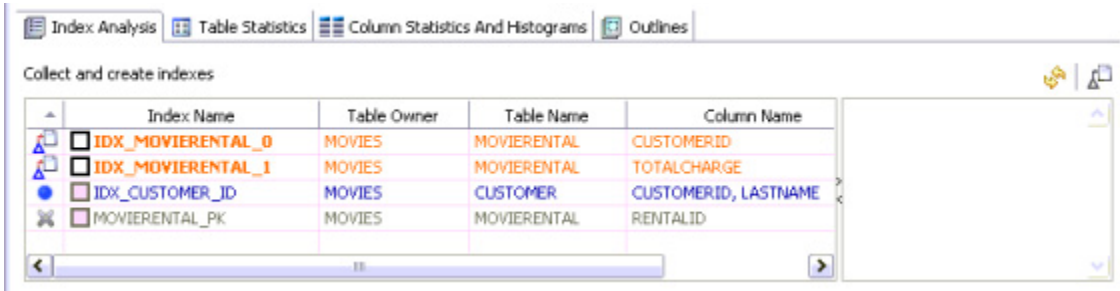
This section is comprised of the following topics:

- [Finding Missing Indexes](#) on page 58
- [Exploring the VST Diagram](#) on page 59
- [Finding Problematic SQL or Schema](#) on page 66

FINDING MISSING INDEXES


Missing indexes are color-coded **orange** in the **Collect and create indexes** area of the **Overview** tab. Creating a missing index can improve the execution path of the SQL statement being analyzed.

TIP: Indexes that are used are color-coded **green**. Indexes that exist in the table but are not used in this execution path are color-coded **grey**. Indexes that are usable but not used by the current execution path are color-coded **blue**.



Index Name	Table Owner	Table Name	Column Name
IDX_MOVIEWERENTAL_0	MOVIES	MOVIERENTAL	CUSTOMERID
IDX_MOVIEWERENTAL_1	MOVIES	MOVIERENTAL	TOTALCHARGE
Idx_CUSTOMER_ID	MOVIES	CUSTOMER	CUSTOMERID, LASTNAME
MOVIERENTAL_PK	MOVIES	MOVIERENTAL	RENTALID

EXPLORING THE VST DIAGRAM

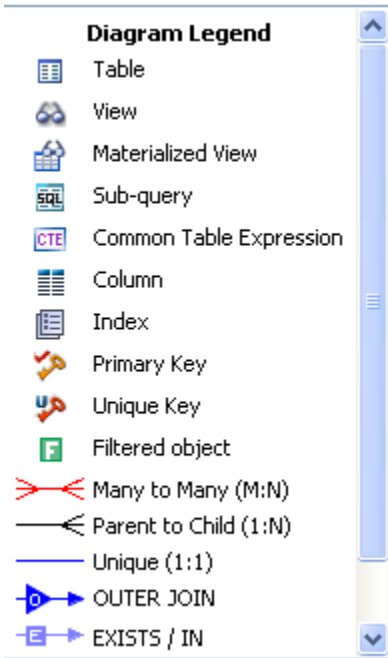
**Legend**

- 1 Display/Hide Diagram Legend
- 2 Display/Hide Explain Plan
- 3 Display/Hide Ratio Counts
- 4 Display Detail Mode or Summary Mode
- 5 Zoom In
- 6 Zoom Out

DISPLAYING OR HIDING THE 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.

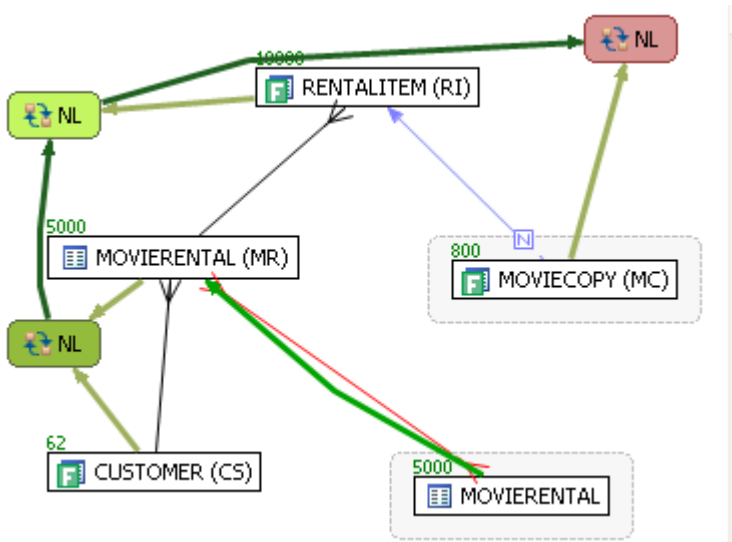


DISPLAYING OR HIDING THE EXPLAIN PLAN

NOTE: The Explain Plan is available only for the Oracle 10g and 11g 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.

DISPLAYING OR HIDING 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 tooltip contents is larger than the size of the tooltip rectangle, you can hover the mouse on top of the tooltip for a second, then it will turn into a dialog you can re-size, scroll in, and select text from to copy into the Clipboard.



DISPLAYING OR HIDING TABLE COLUMNS AND INDEXES

Click the **Detail/Summary** switch to display or hide details of the table display, including table columns and indexes.

Only fields that are used in the WHERE clause are displayed in detail mode. For information on interpreting the icons used in Detail mode, activate the Diagram Legend.

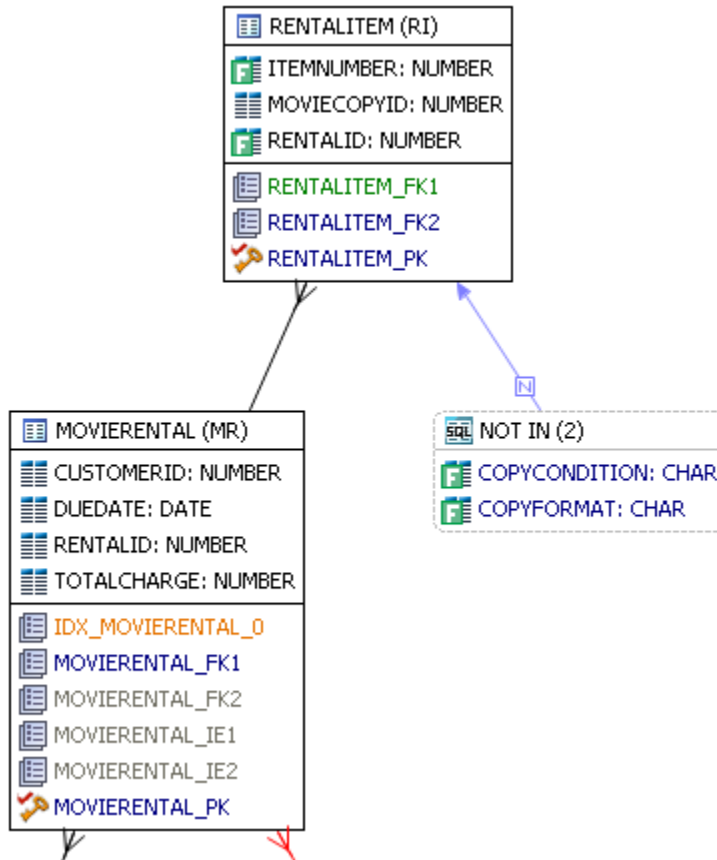
TIP: You can also switch between Summary Mode and Detail Mode for a specific table or view by double-clicking the object name.

NOTE: 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.

VIEWING ALL FIELDS IN A TABLE

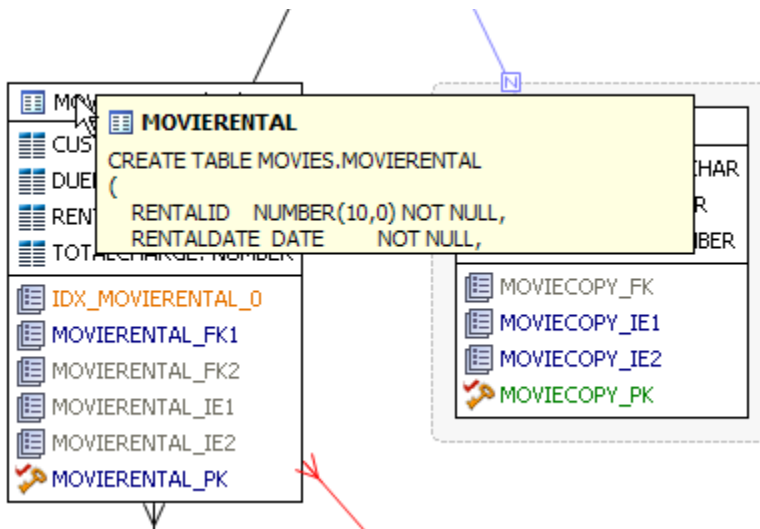
Hover the mouse over the table name to view all fields in the table in a pop-up window.

The illustration below shows the pop-up window that appears when hovering over the table.



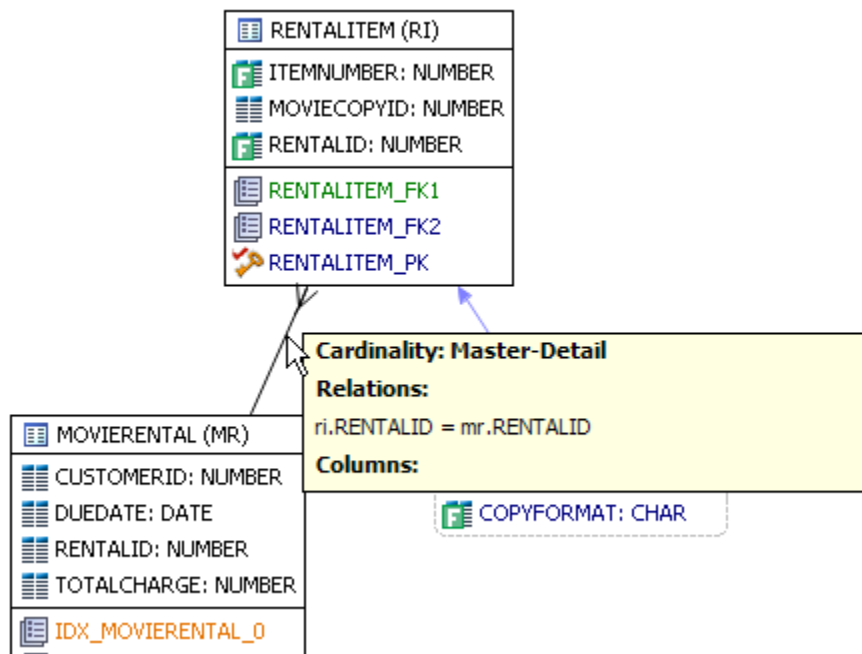
VIEWING OBJECT SQL

Hover the mouse over the table name, field, or index displays the SQL required to create that object.



VIEWING RELATIONSHIPS

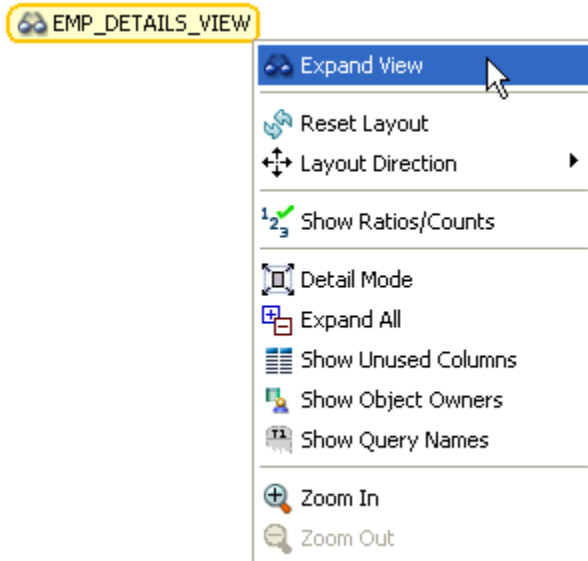
Hover over the join between two tables displays the relationship between the two tables.



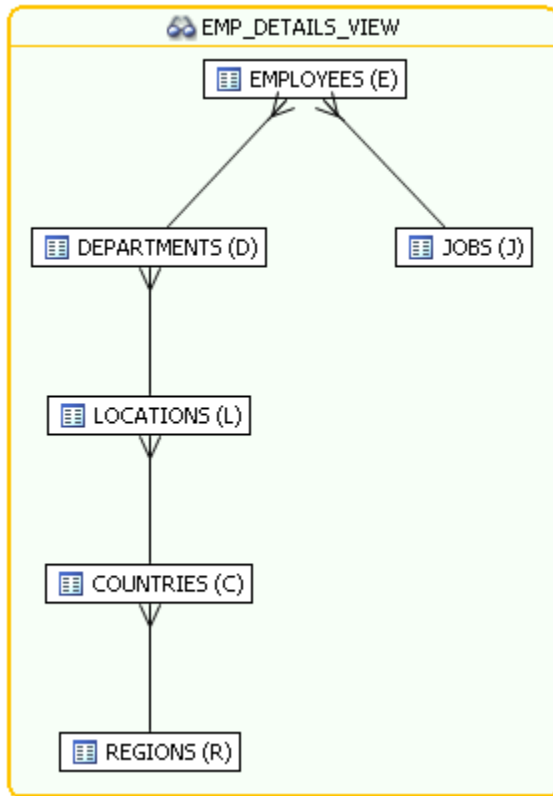
EXPANDING VIEWS IN THE VST DIAGRAM

For example, the following is the default layout from query join table CLIENT (c) to view TRANSACTIONS (t).

Right click on the view, and choose **Expand View**

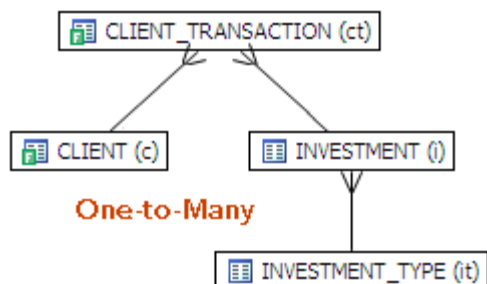


Now we can see the objects in the view.

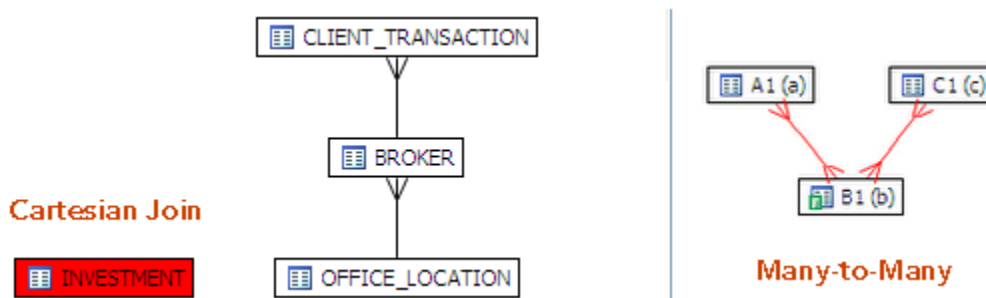


FINDING PROBLEMATIC SQL OR SCHEMA

In the Visual SQL Tuning (VST) diagram, a well-formed query would resemble the following:



Problems such as Cartesian joins, implied Cartesian joins, and many-to-many relationships are clearly represented in the VST.



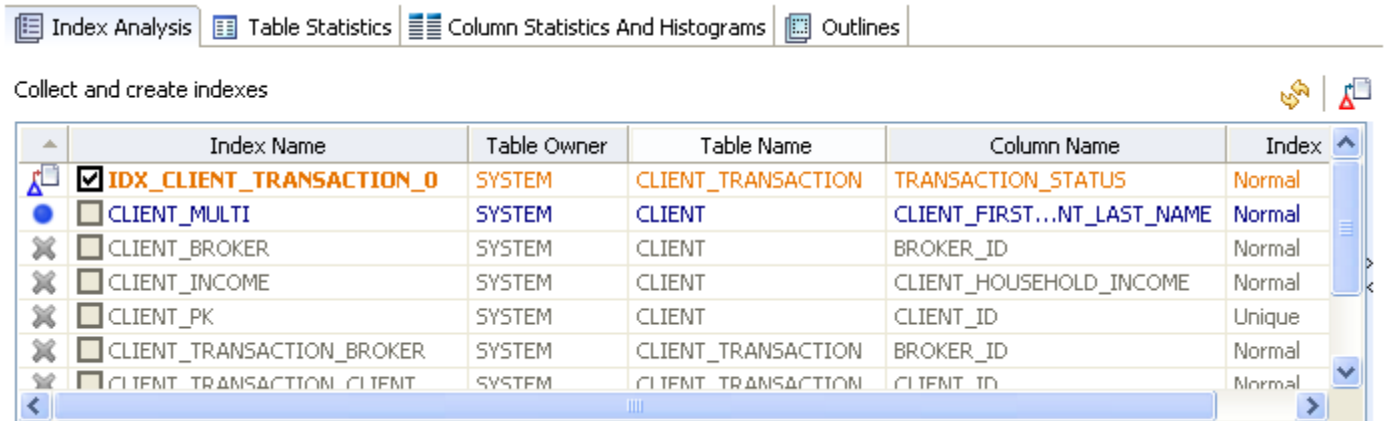
Joins are represented with connecting lines between nodes. The following describes when a particular type of connecting line is used and the default positioning of the line.

Connecting Lines	When used
	One-to-One Join relationships are graphed horizontally using blue lines. One-to-one joins exist when two tables are joined on their primary key.
	One-to-Many Join relationships are graphed with the many table above the one table.
	Cartesian Join shows the table highlighted in red with no connectors to indicate that it is joined in via a Cartesian join. The query is missing join criteria. A Cartesian join problem could be resolved by implementing a rewrite suggestion shown in the Generated Cases area of the Overview tab.
	Many-to-Many Join relationships are connected by a red line and the relative location is not restricted. If master detail information is missing then the VST diagram will have many-to-many connectors. The optimizer can more consistently optimize a well-formed query, so the query will run faster.

IMPLEMENTING RECOMMENDATIONS ON THE INDEX ANALYSIS TAB

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, an index that DB Optimizer recommends that you create is marked in orange and has a create index icon next to its checkbox.

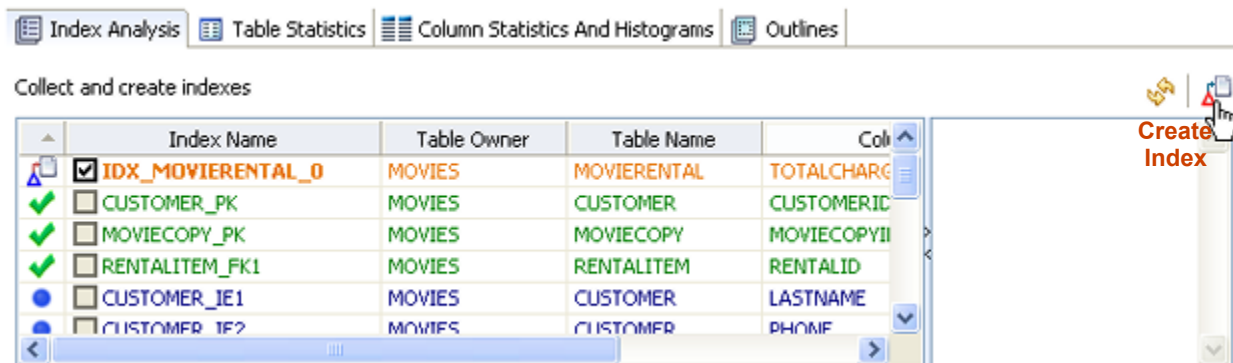


Index Name	Table Owner	Table Name	Column Name	Index
<input checked="" type="checkbox"/> IDX_CLIENT_TRANSACTION_0	SYSTEM	CLIENT_TRANSACTION	TRANSACTION_STATUS	Normal
<input type="checkbox"/> CLIENT_MULTI	SYSTEM	CLIENT	CLIENT_FIRST...NT_LAST_NAME	Normal
<input type="checkbox"/> CLIENT_BROKER	SYSTEM	CLIENT	BROKER_ID	Normal
<input type="checkbox"/> CLIENT_INCOME	SYSTEM	CLIENT	CLIENT_HOUSEHOLD_INCOME	Normal
<input type="checkbox"/> CLIENT_PK	SYSTEM	CLIENT	CLIENT_ID	Unique
<input type="checkbox"/> CLIENT_TRANSACTION_BROKER	SYSTEM	CLIENT_TRANSACTION	BROKER_ID	Normal
<input type="checkbox"/> CLIENT_TRANSACTION_CLIENT	SYSTEM	CLIENT_TRANSACTION	CLIENT_ID	Normal

ACCEPT THE SUGGESTION AND AUTOMATICALLY GENERATE AN INDEX

- 1 For any recommended index, click the checkbox to the left of the index you want to create.

For a selected index you can modify the Index type by clicking in the **Index Type** column and then selecting a type from the list: **Normal**, **Bitmap**, **Reverse Key**, **Reverse Key Unique**, or **Unique**.



Index Name	Table Owner	Table Name	Column Name	Index Type
<input checked="" type="checkbox"/> IDX_MOVIERENTAL_0	MOVIES	MOVIERENTAL	TOTALCHARGE	Normal
<input checked="" type="checkbox"/> CUSTOMER_PK	MOVIES	CUSTOMER	CUSTOMERID	Unique
<input checked="" type="checkbox"/> MOVIECOPY_PK	MOVIES	MOVIECOPY	MOVIECOPYID	Unique
<input checked="" type="checkbox"/> RENTALITEM_FK1	MOVIES	RENTALITEM	RENTALID	Normal
<input type="checkbox"/> CUSTOMER_I1	MOVIES	CUSTOMER	LASTNAME	Normal
<input type="checkbox"/> CUSTOMER_I2	MOVIES	CUSTOMER	PHONE	Normal

- 2 Click the **Create Index** button.

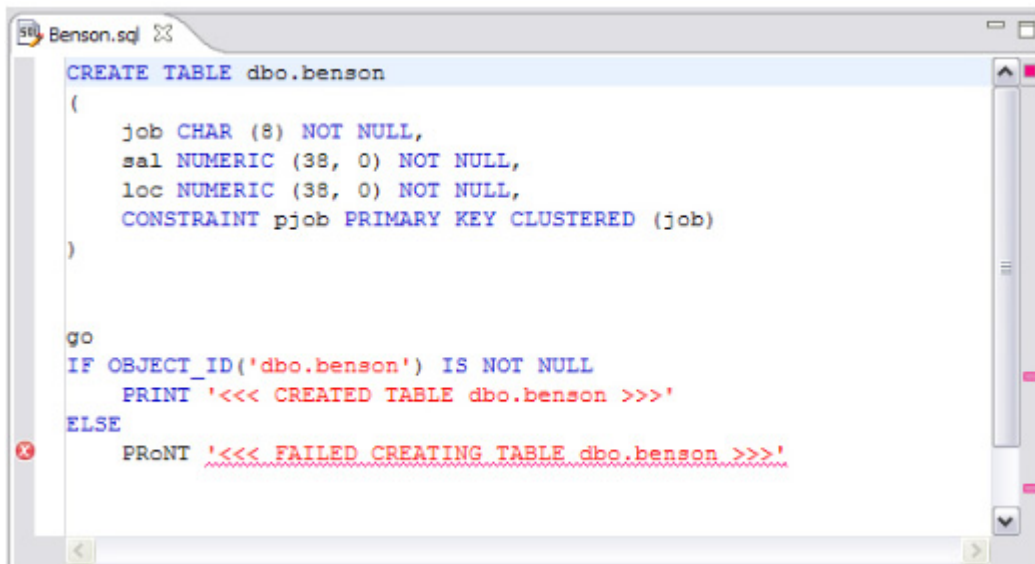
The **Index Analysis** dialog appears.

- 3 To view the index SQL in an editor for later implementation, click the statement and then click **Open in a SQL editor**.
- 4 To run the index SQL and create the index on the selected database, click **Execute**.

SQL CODE ASSIST AND EXECUTION

SQL Code Assist is an interface component that enables the development and formatting of SQL code for the purposes of creating and modifying database objects. It provides a front-end application for the delivery of code through its object code extraction capabilities.

Code Assist provides a number of key features that assist in the coding process, ensuring greater accuracy in coding, faster development cycles, and a general increase in efficiency overall.



```
Benson.sql
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 >>>'
```

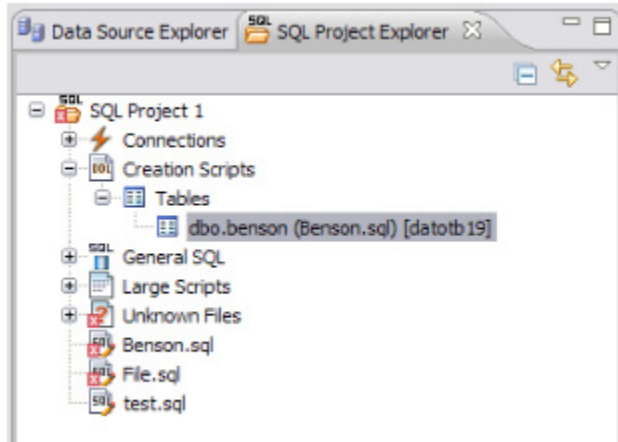
SQL Editor provides key features to ensure an increase in code accuracy and overall development efficiency.

The following key features are provided with SQL Code Assist:

- [Code Extraction](#) on page 69
- [Code Highlighting](#) on page 70
- [Automatic Error Detection](#) on page 70
- [Code Complete](#) on page 72
- [Hyperlinks](#) on page 72
- [Code Formatting](#) on page 73
- [Code Folding](#) on page 74
- [Code Quality Checks](#) on page 74

In order to access SQL Editor, you need to create a new file or edit existing code from Data Source Explorer.

Additionally, SQL Project Explorer provides a tree structure for all files created in DB Optimizer. You can access files from here by double-clicking the file name you want to open as well.



SQL Project Explorer provides a tree of all files developed and saved in DB Optimizer.

To create a new file

Choose **File > New > SQL File**.

A blank instance of SQL Editor appears in the Workbench. If you save this file, it is automatically added to the SQL Project Explorer.

To edit an existing file

Use Data Source Explorer or Project Explorer to navigate to the code you want to modify and double-click it.

An instance of SQL Editor appears in the Workbench, populated with the extracted code of the specified object or SQL project file.

CODE EXTRACTION

SQL Editor provides the ability to extract the underlying SQL code of database objects registered in DB Optimizer to provide a front-end application for the development and modification of data sources in your enterprise.

To extract underlying SQL code

Navigate to a database object in Data Source Explorer and select **Extract** via the right-click menu.

The object's underlying SQL code appears in SQL Editor and is ready for editing and further modification.

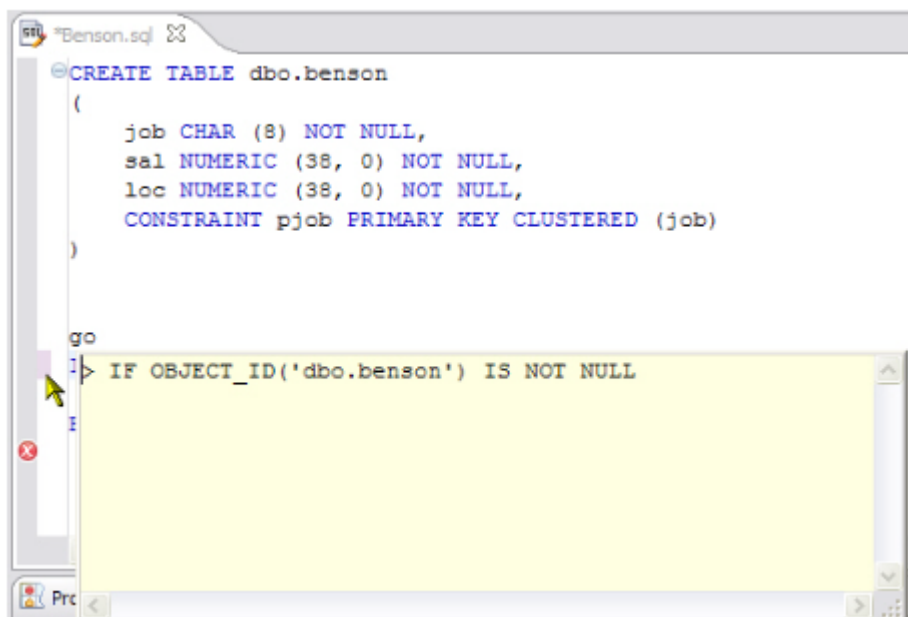
CODE HIGHLIGHTING

SQL Editor identifies commands and provides syntax highlighting changes that are automatically added to the code as you add lines, which enables you to clearly and quickly understand code when you read it in the interface.

The following syntax highlighting is automatically added to lines of code in SQL Editor:

Code	Formatting
Comments	<i>green italics</i>
SQL Commands	dark blue
Syntax Errors	<u>red underline</u>
Coding Errors	red

Additionally, SQL Editor provides a purple change bar in the left-hand column that indicates if a line of code has been modified from the original text. You can hover over this change bar to view the original code line. A red square icon in the right-hand column indicates that there are errors in the code line. You can hover over the icon to view the error count.

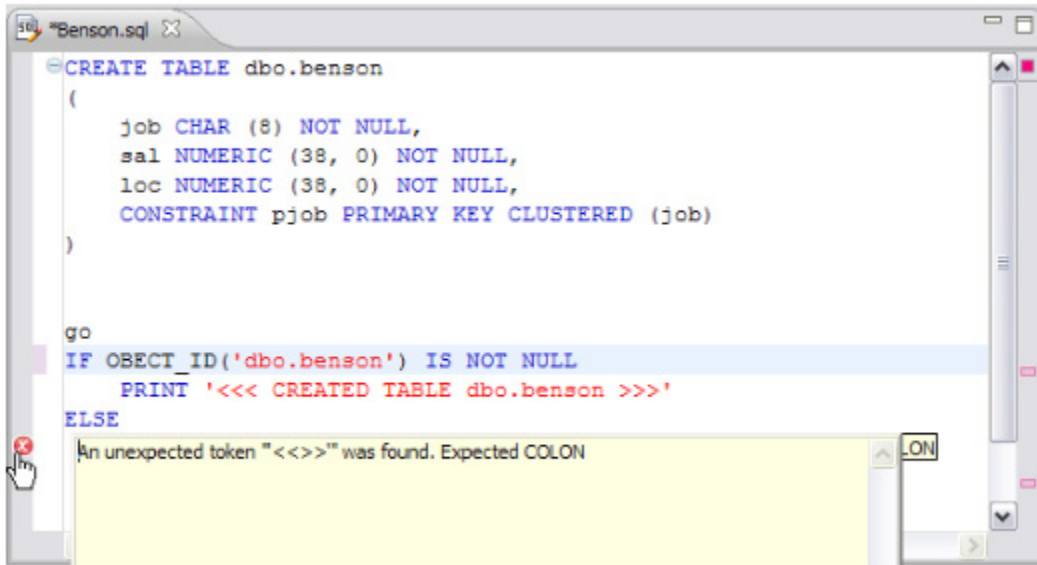


The purple change bar indicates if a line of code has been changed from its original text. Hover your mouse over the change bar to view the original text.

AUTOMATIC ERROR DETECTION

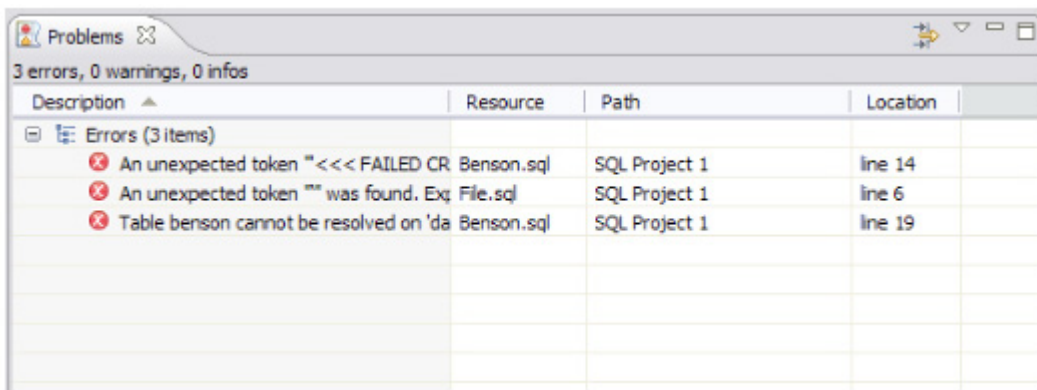
The automatic error detection functionality of the editor highlights errors and typos in the code as you work in "real time".

Automatic error detection automatically identifies and analyzes SELECT, FROM, WHERE, GROUP BY, HAVING, and ORDER BY statements. If it detects any syntax errors while you type these statements, the line is automatically flagged by the error icon in the left-hand column of SQL Editor. You can hover your mouse over the icon to view any errors.



Syntax errors are automatically flagged by line as you work with code in SQL Editor. Hover the mouse over the error icon to view the specific error message.

Additionally, all semantic errors are recorded in the **Problems** view, an interface component that automatically logs errors and warnings as you work with files.



The Problems view logs errors and warnings as you work with files in SQL Editor.

You can double-click on a line in the **Problems** view and DB Optimizer will automatically navigate you to that issue in SQL Editor.

CODE COMPLETE

SQL Editor provides suggestion capabilities for both DML statements and objects. It has the ability to look up object names in order to avoid errors when defining table names, columns, etc. in the development process. This feature provides lists of object and code suggestions that will make the development process more efficient as you will not be forced to manually look up object names and other statement values. SQL Editor provides code assist for SELECT, UPDATE, INSERT, and DELETE statements and object suggestion support for tables, alias tables, columns, alias columns, schemas, and catalogs.

To activate code assist

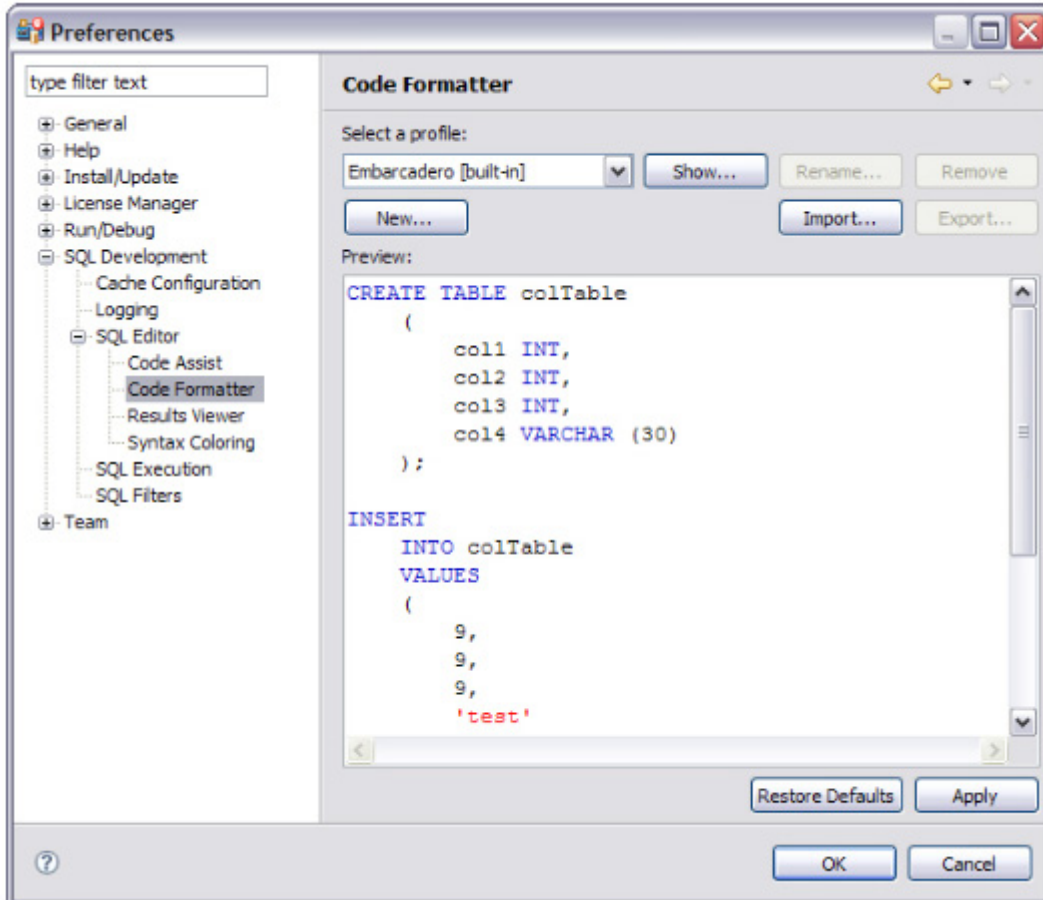
- 1 Click the line on which you want to activate the code assist feature.
- 2 Press **CTRL + Spacebar** on your keyboard. Code assist analyzes the line and presents a list of suggestions as appropriate based on the elements of the statement.

HYPERLINKS

Hyperlinks are used in SQL Editor to provide links to tables, columns, packages and other reference objects. When you select a hyperlinked object from a piece of code, a new editor opens and displays the source. Additionally, hyperlinks can be used to link procedures or the function of a call statement, as well as function calls in DML statements. To enable a hyperlink, hover your mouse over the object name and hold the CTRL key. It becomes underlined and changes color.

CODE FORMATTING

Code formatting is automatically applied to a file as you develop the file in SQL code. This enables you to set global formatting preferences one time, and then apply them to all code development, saving time and allowing for a more efficient code development process. The code formatter can be accessed by selecting **CTRL + Shift + F** in the editor. All code is automatically formatted based on parameters specified on the **Code Formatter** node of the **Preferences** panel. (Select **Window > Preferences** in the main menu to access this panel.)



Code formatting parameters can be globally set and then applied to all development work in SQL Editor.

In addition to formatting code per individual file, you can also format an entire group of files from **Project Explorer**. Select the directory of files that you want to apply formatting to and execute the **Format** command from the right-click menu. The files will be automatically formatted based on the global preferences.

CODE FOLDING

The code folding feature automatically sorts code into a tree-like outline structure in SQL Editor. This increases navigation and clarification capacities during the code development process. It ensures that the file will be easily understood should any future work be required on the code. As you work in SQL Editor, collapsible nodes are automatically inserted into the appropriate lines of code. Statements can then be expanded or collapsed as needed, and this feature is especially useful when working on parts of particularly large or complicated files.

CODE QUALITY CHECKS

On Oracle-based code, the code quality checking feature provides suggestions regarding improved code on a statement-by-statement basis. As you work in SQL Editor, markers provide annotations that prevent and fix common mistakes in the code.

Notes regarding code quality suggestions appear in a window on any line of code where the editor detects an error, or otherwise detects that the code may not be as efficient as it might be. Code quality check annotations are activated by clicking the light bulb icon in the margin, or by selecting **Ctrl + I** on your keyboard.

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

- Statement is missing valid JOIN criteria
- Invalid or missing outer join operator
- Transitivity issues
- Nested query in WHERE clause
- Wrong place for conditions in a HAVING clause
- Index suppressed by a function or an arithmetic operator
- Mismatched or incompatible column types
- Null column comparison

To activate code quality checks:


- Click the light bulb icon in the margin of the editor or select **Ctrl + I** on your keyboard.

The editor suggestions appear in a window beneath the selected statement. When you click a suggested amendment, the affected code is automatically updated.

SQL EXECUTION

When you have finished developing or modifying code, you can then execute the file from within the DB Optimizer environment, on the database of your choosing. This enables you to immediately execute code upon completion of its development. Alternatively, you can save files for execution at a later point in time.

In order to execute a file, you must first associate it with a target database. This is performed by using the drop-down menus located in the Toolbar. When a SQL file is open in the Workbench, the menus are enabled. Select a data source and a corresponding database to associate the file with and then click the green arrow icon to execute the file.




The pair of drop-down menus indicate that the SQL file is associated with the data source and the database. When the green arrow icon on the right-hand side of the menus is selected, the file is executed on the specified data source and database.

Additionally, if you have turned off auto-committal in the **Preferences** panel (**Window > Preferences**) you can commit and execute transactions via the **Commit Transaction** and **Start Transaction** icons located beside the **Execute** icon.


EXECUTE A FILE

Open the file you want to run and ensure it is associated with the correct database, then click the **Execute** icon [▶]. DB Optimizer executes the code on the database you specified.

EXECUTE A TRANSACTION

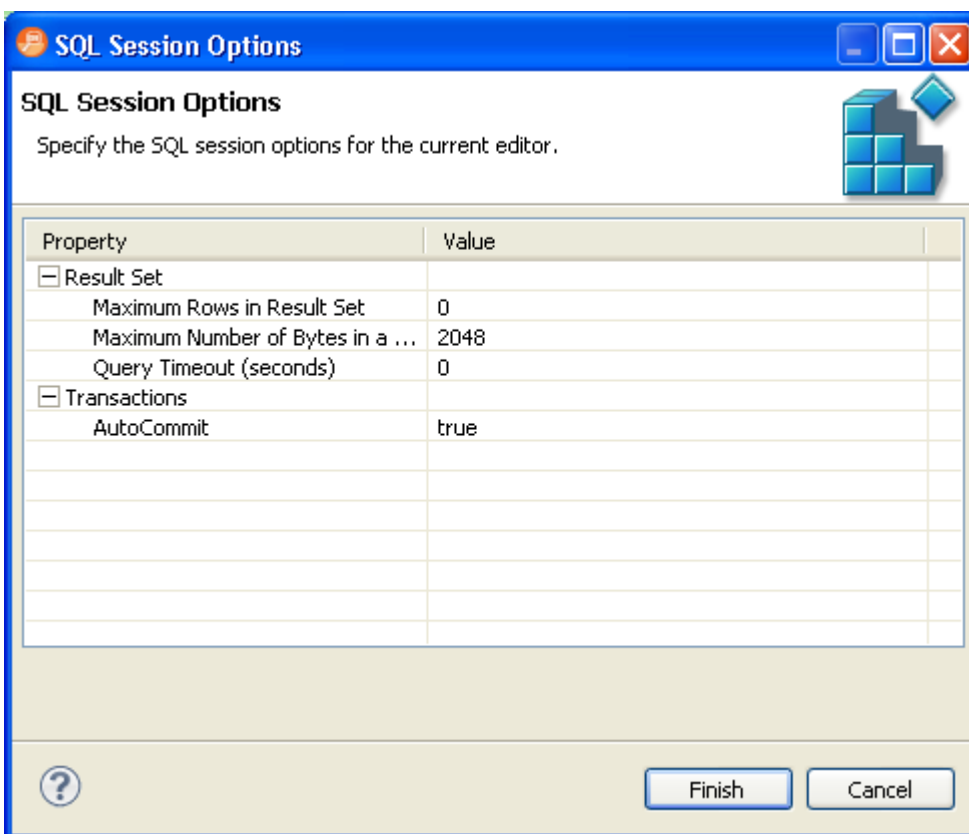
Open the transaction file you want to run and ensure it is associated with the correct database, then click the **Start Transaction** icon []. DB Optimizer executes the transaction on the database you specified.

COMMIT A TRANSACTION

Open the transaction file you want to commit, ensure it is associated with the correct database, and click the **Commit Transaction** icon []. DB Optimizer commits the transaction on the database you specified.

CONFIGURING SQL EXECUTION PARAMETERS

Use the **SQL Session Options** dialog to modify the configuration parameters that determine how DB Optimizer executes code. These options ensure that code is executed the way you want on an execution basis, ensuring accuracy and flexibility when running new or modified code.



SQL Sessions Options provide you with the flexibility to adjust execution parameters on a session-by-session basis.

MODIFY SQL SESSION OPTIONS

- 1 Click the SQL Session Options icon on the toolbar. The **SQL Sessions Options** dialog appears.
- 2 Click on the individual parameters in the **Value** column to change the configuration of each property, as specified.
- 3 Click **Finish**. The session options are changed and DB Optimizer executes the code as specified by your options.

NOTE: SQL Session Options are only applied to the currently-selected code, and are not retained across different files with regards to execution.