



Embarcadero® DB Optimizer™ 1.0 New Features Guide

Published: June 26, 2008

Copyright © 1994-2008 Embarcadero Technologies, Inc.

Embarcadero Technologies, Inc.
100 California Street, 12th Floor
San Francisco, CA 94111 U.S.A.
All rights reserved.

All brands and product names are trademarks or registered trademarks of their respective owners.
This software/documentation contains proprietary information of Embarcadero Technologies, Inc.; it is provided under a license agreement containing restrictions on use and disclosure and is also protected by copyright law. Reverse engineering of the software is prohibited.

If this software/documentation is delivered to a U.S. Government Agency of the Department of Defense, then it is delivered with Restricted Rights and the following legend is applicable:

Restricted Rights Legend Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of DFARS 252.227-7013, Rights in Technical Data and Computer Software (October 1988).

If this software/documentation is delivered to a U.S. Government Agency not within the Department of Defense, then it is delivered with Restricted Rights, as defined in FAR 552.227-14, Rights in Data-General, including Alternate III (June 1987).

Information in this document is subject to change without notice. Revisions may be issued to advise of such changes and additions. Embarcadero Technologies, Inc. does not warrant that this documentation is error-free.

Contents

NEW FEATURES SUMMARY	4
TOP NEW FEATURES	5
SQL Tuner.....	5
SQL Profiler.....	5
Data Source Migration and Enterprise Management	6
SQL Code Development	9
SQL Execution	10
SQL Project Insight	11
ADDITIONAL RESOURCES.....	13
Licensing Your Embarcadero Technologies Product.....	13
Embarcadero Technologies Product Support.....	13
Embarcadero Technologies Technical Support.....	13
Embarcadero Technologies on the Web	13

New Features Summary

Embarcadero DB Optimizer maximizes the performance of your database by enabling developers and DBAs to:

- Quickly discover, diagnose, and optimize poor-performing SQL
- Eliminate performance bottlenecks by discovering data intensive or frequently executed queries
- Focus in on specific SQL statements through query statistics (CPU, I/O, wait times)
- Optimize any problematic statements.

The DB Optimizer user interface offers easy to use analysis tools and improves overall productivity. DB Optimizer offers native support for IBM® DB2® for LUW, Oracle®, Microsoft® SQL Server and Sybase® as well as JDBC support for other DBMS.

SQL Tuner and SQL Profiler

These features let you locate and diagnose problematic SQL. The profiling and tuning components offer wait event-based analysis, explain plan cost details, SQL Optimizer hint-based suggestions, and a variety of execution statistics. Both features offer high level, graphical starting points you can use to spot problem areas. Drill down capabilities let you identify specific, problematic SQL statements.

Data Source Migration and Enterprise Management

Data Source Explorer provides a tree-structured interface that organizes data sources objects within the development environment. Once a data source has been registered in DB Optimizer, commands can be issued to extract code from objects, compile stored routines, and drop objects. Data Source Explorer also features search capabilities and filters, as well as declaration and reference search commands. Migration functionality imports data sources from Eclipse DTP (Data Tools Project) or Quest TOAD, and automatically detects pre-existing catalogs on local machines.

SQL Code Development

SQL Editor provides the ability to create and edit SQL code files, and associate those files with registered data sources. It provides a number of features including SQL Code Assist and Formatting Profiles. Additionally, error detection and code extraction functionality is provided, and real-time SQL validation ensures that there are no parser violations or use of object names that are not associated with the database.

SQL Execution

DB Optimizer executes code directly on registered data sources. Code execution however, is initiated through the interface. When a file is opened in SQL editor, you can specify a database and associated catalog on which you want to execute the file via drop down lists. You can then execute the code on the specified database, or commit and start a transaction. Additionally, each execution session can be customized through a series of session options you can customize prior to execution.

SQL Project Insight

SQL Project Insight provides project-level SQL file cataloging and search features to streamline project organization and maintenance.

Top New Features

SQL Tuner

SQL Tuner lets you view cost details on collections of SQL statements you add to a tuning session. As a means of identifying bottlenecks or problem queries, SQL Tuner lets you:

- View high-level explain plan costs for statements as well as offering the option to view operation-by-operation explain plan details
- Generate execution statistics such as CPU time as well as Read and Sort details
- Work with automatically-generated SQL Optimizer hint-based variations on all statements in a tuning session

Used as a development aid, SQL Tuner lets you fine tune queries before putting them into production. As a maintenance tool, SQL Tuner can help you isolate and optimize queries in production systems.

SQL Statements and_Cases				Cost	Elapsed Time		Other Execution Statistics	
Name	Text	Source	Index Analysis	Value	Value (s)	Result	Physical Reads	Logical Reads
<input checked="" type="checkbox"/>	FIRST_ROWS							
<input checked="" type="checkbox"/>	ALL_ROWS							
<input checked="" type="checkbox"/>	Statement 50	select from TEMP_CL GI...WN	OK					
<input checked="" type="checkbox"/>	Statement 53	select from user_tab	Oracle SGA OK	10.0	0.001		0	6
<input checked="" type="checkbox"/>	NO_EXPAND			10.0	0.031		0	8
<input checked="" type="checkbox"/>	USE_CONCAT			10.0	0.188		0	8
<input checked="" type="checkbox"/>	Statement 52	select from < Nestex	Oracle SGA OK	2.0	0.439		0	6
<input checked="" type="checkbox"/>	Statement 51	select from	Oracle SGA Optimize	2.0	0.016		2	5
<input checked="" type="checkbox"/>	ALL_ROWS			2.0	0.031		0	5
<input checked="" type="checkbox"/>	FIRST_ROWS			2.0	0.062		0	5
<input checked="" type="checkbox"/>	NO_PARALLEL			2.0	0.032		0	5
<input checked="" type="checkbox"/>	STAR_...ATION			2.0	0		0	5
<input checked="" type="checkbox"/>	NO_ST...TION			2.0				
<input checked="" type="checkbox"/>	CACHE			2.0	0.015		0	5
<input checked="" type="checkbox"/>	NOCACHE			2.0				

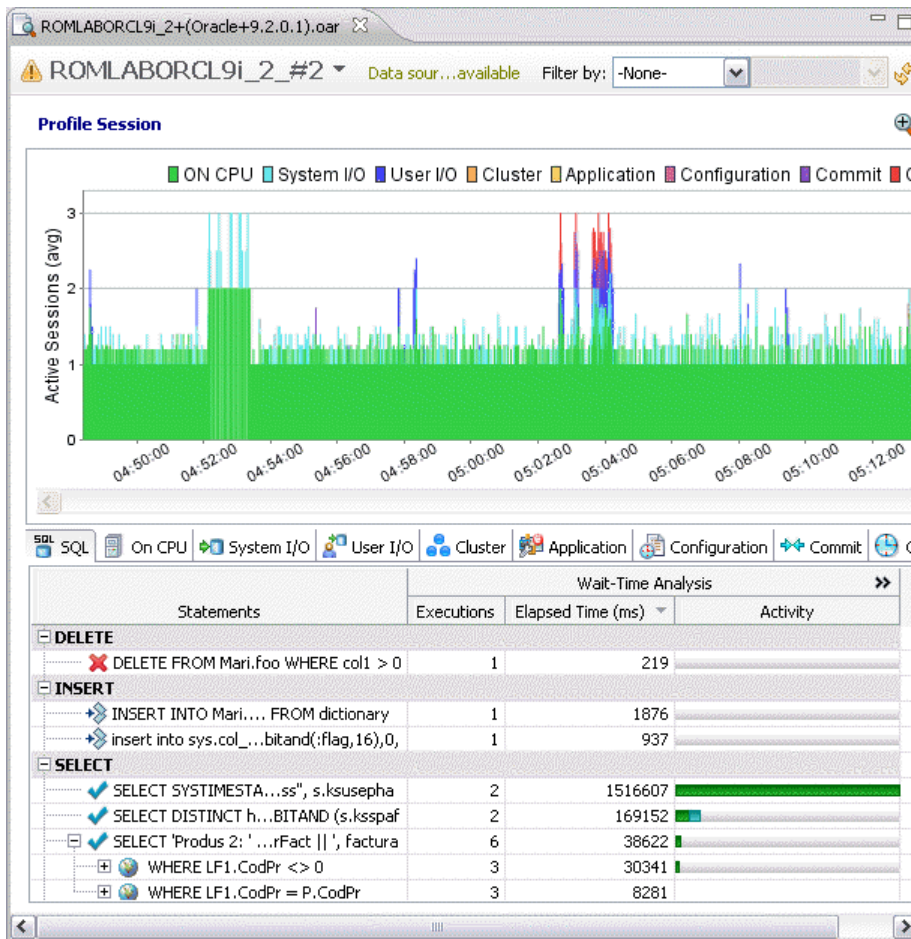
Your tuning sessions can include SQL from the following sources:

- Typed or pasted statements
- DML statements extracted from SQL-containing database objects: functions, materialized views, package bodies, procedures, stored outlines, triggers, and views
- DML statements extracted from SQL Files
- Statements currently active in the Oracle SGA

Support for simultaneous tuning of multiple statements lets you perform high-level, “at a glance” comparisons as well as detailed, low-level analysis. In combination with the ability to save your tuning session in an archive, this lets you group logically-related or functionally-related code, revisit your analysis, add or remove SQL, and routinely generate new statistics for critical systems.

SQL Profiler

SQL Profiler lets you locate and diagnose problematic SQL and wait event-based bottlenecks. At a high level, SQL Profiler lets you take fixed-duration snapshots of query activity on a data source. At a lower level, SQL Profiler lets you investigate execution and wait event details for individual stored routines. SQL Profiler results, presented in a profiling editor, let you quickly identify problem areas and subsequently drill down to individual, problematic SQL statements.



At a high level, you can view the distribution of sessions that were executing over the duration of the profile and those that were waiting on DBMS-specific event categories. This is your starting point in identifying problem areas. Tabs at the bottom of the profiling editor provide:

- Total number of executions, execution time, and wait time for SQL statements waiting or executing over the profile duration
- Elapsed time details for specific wait event types and for executing sessions

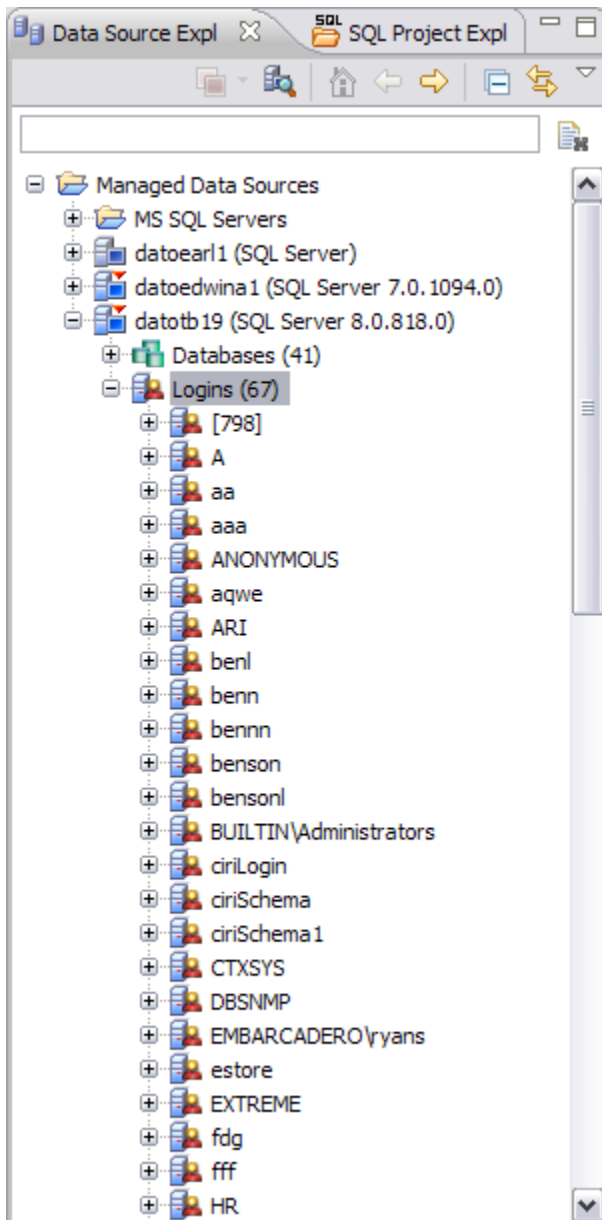
Zooming in and selecting one or more intervals of the graph updates tab details to display only statistics for those intervals, letting you find particular problem areas. You can then drill down to investigate SQL statements, statement components, and associated wait events active during those intervals. In short, SQL Profiler provides the tools to not only indicate the general area of concern, but also to isolate the specific causes of wait-based inefficiency. Integration with other components speeds up correction of problems found.

The profiling editor can run in real-time monitor mode, displaying statistics as the session runs, or simply open when the session completes. As well, the ability to save your profiling sessions lets you return to your analysis at your convenience.

Data Source Migration and Enterprise Management

Data Source Explorer provides a hierarchical tree view used to organize and manage data sources within your enterprise.

It displays databases, database objects, logins, and other information as nodes in the tree. Data sources are, by default, organized by DBMS. The SQL code for each object residing on a data source is accessible under the specified node, and can be extracted for editing, modification, or reuse via **SQL Editor**.



The Data Source Explorer organizes your data sources and provides a graphical display of the database objects that reside on them.

When you first install DB Optimizer, it automatically detects data sources that may already be registered on other Embarcadero products such as *DBArtisan* or *Rapid SQL*, or via DBMS clients installed on your machine. A dialog prompts you to add any data sources that it finds.

If DB Optimizer does not detect a data source automatically, the **Auto-Discover Data Sources** command provides a means of manually registering new data sources.

Once registered, Data Source Explorer displays all databases and corresponding objects residing on the data source. These objects are displayed by category beneath the appropriate data source node.

Commands are issued against data sources and database objects from Data Source Explorer. Some common commands enable you to extract the DDL of database objects, drop objects from the database, and compile stored routines.

To access these commands, right click on a node to access the shortcut menu.

To register a data source:

1. Select **File > New > Data Source** from the Main Menu. The **New Data Source Wizard** appears.
2. Enter the appropriate connection parameters associated with the data source and click **OK**.
The data source appears in the tree view.
3. View the database objects that reside on the new data source by expanding the node. You can use the shortcut menu to execute available commands as needed.

To extract database objects:

Right-click on the database object from which you want to extract the DDL and choose **Extract to Project**.

The code appears in SQL Editor, where it can be modified, copied, or edited, executed, and saved.

To drop database objects:

Right-click on the database object you want to remove and choose **Drop** from the menu.

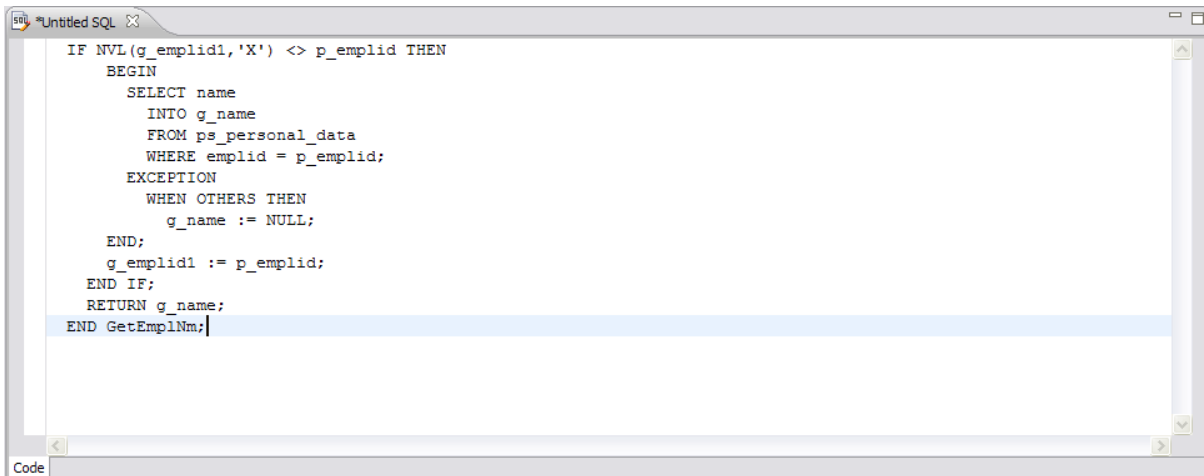
The object is removed from the specified data source.

To compile a stored routine:

1. Right-click on a database object and choose **Compile**.
The **Compile** wizard appears.
2. Specify if you want to compile with debug options and/or dependent objects and view the script in the preview window.

SQL Code Development

SQL Editor enables you to create SQL files and associate them with data sources that are registered in DB Optimizer as well as edit existing files. It provides features such as code assistance that makes suggestions for partially-entered code strings, code folding and syntax coloring to provide enhancement to the organization of code structures, and formatting templates that enable code format customization in the editor window. These features aid developers by increasing the development process and increasing overall productivity.



The SQL editor provides an interface for creating and editing SQL code.

While you are developing files, DB Optimizer checks code for syntactic and semantic errors. All parser violations, as well as references to tables, views and columns that are absent from the target database are flagged.

Additionally, in SQL Editor, syntax errors that are detected are color highlighted, and validation checks are performed every time you create, edit, or save a SQL file.

Once code has been developed, commands enable script and stored routine execution commands to be run externally on specified data sources. Various logs and views return execution problems and code statements of SQL commands run by you or the system.

To create a new file:

1. Select **File>New>SQL File** from the Menu Bar.

SQL Editor opens.

2. Enter the new code in the editor window. When you are finished, select **Save**.

You can store the new file in a project directory, or associate it with a registered data source. Once a file has been associated with a data source, it can then be executed on the target database.

To edit an existing file:

1. In **Data Source Explorer**, select the file you want to edit and choose **Edit**.

SQL Editor opens pre-populated with the corresponding code.

2. Edit the code, as needed. When you are finished, select **Save**.

You can run the modified file on the associated data source or associate it with a new one.

To associate a file with an existing data source:

Ensure that the Editor with the file open is selected and then use the drop down menu that appears at the top of the workspace and select a data source to associate the open file with.

SQL Execution

All SQL scripts and stored routines in the system can be executed from within the development environment, customized to any of the supported database platforms. Additionally, DB Optimizer provides the ability to start, commit, and rollback transactions. The results of executed scripts can be viewed in the Results Viewer and SQL Log view.



Icons in the Menu Bar provide easy access to script execution and transaction commands.

Once you have executed a script or routine, the Results Viewer provides the results of the execution process. Results can be sorted by the column headings of the table.



The Results Viewer provides messages that indicate the results of the SQL execution process.

Additionally, DB Optimizer logs all statements that were executed by the system or the user during the process in the SQL Log View. Statements can be sorted by statement name, date executed, or the name of the data source on which the file was executed.

SQL Statement	Date	Host/Server	DBMS	Service	User
<input checked="" type="checkbox"/> IF OBJECT_ID('dbo.benson') IS NOT NULL PRINT '<<< CREATED TABLE dbo.benson >>>'	2008-02-04 11:06:53.171	datotb19	SQLServer	--	sa
<input checked="" type="checkbox"/> CREATE TABLE dbo.benson (job char(8) NOT NULL, sal numeric(38,0) NOT NULL,	2008-02-04 11:06:53.00	datotb19	SQLServer	--	sa
<input checked="" type="checkbox"/> ALTER TABLE dbo.testapps ADD CONSTRAINT CK_123 CHECK ([a3] > 40)	2008-02-04 11:06:12.656	datotb19	SQLServer	--	sa

The SQL Log view logs statements and provides information regarding the execution process.

To execute SQL scripts:

Open the file you want to execute and click the SQL execution icon in the Menu Bar. The script executes and results are displayed in the Results Viewer. Statements are captured to the SQL Log.

To execute stored routines:

Open the file you want to execute and select **Run > Run As > SQL Stored Routine**. The routine executes and results are displayed in the Results Viewer. Statements are captured to the SQL Log.

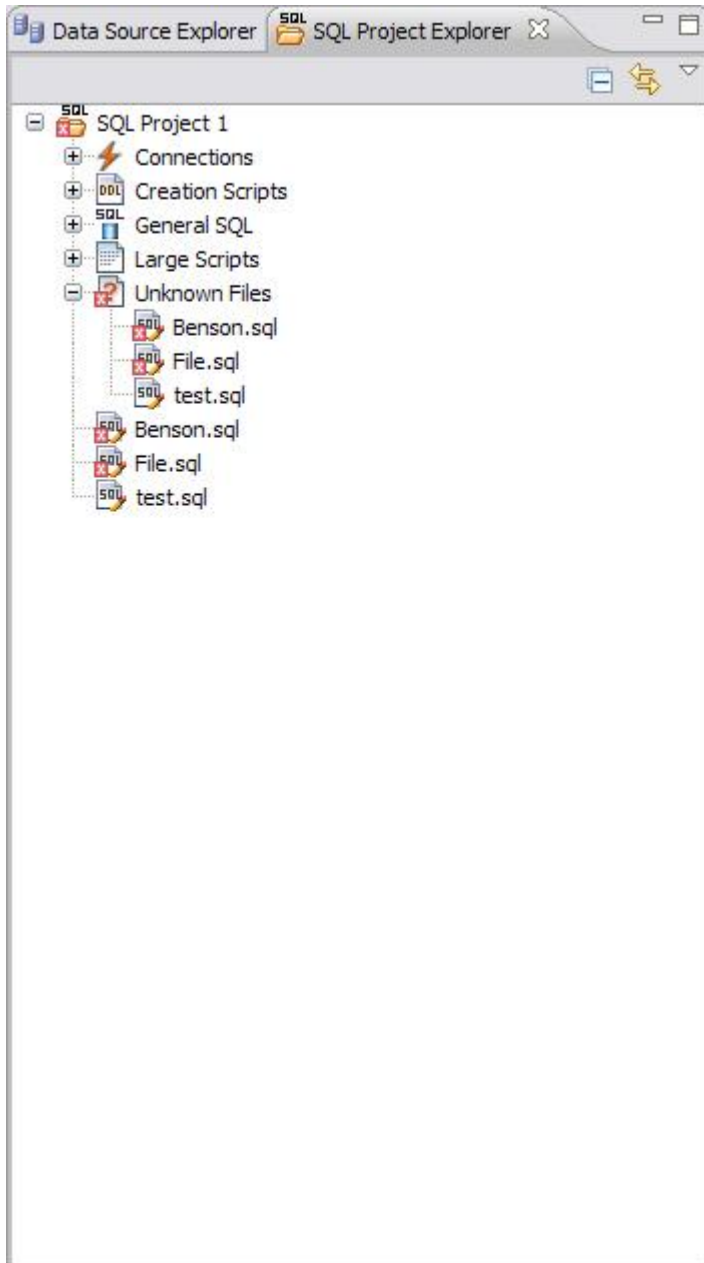
To execute, commit, or rollback transactions:

Open the file you want to execute, commit, or rollback, and click the respective icon in the Menu Bar. The process runs and results are displayed in the Results Viewer. You need to ensure that auto-

commit is turned off prior to committing transactions by selecting **Window > Preferences** and then modifying auto-commit parameters on the **SQL Editor** pane.

SQL Project Insight

SQL Project Explorer provides a tree view of all projects created in DB Optimizer for the purpose of organizing and storing SQL development files. This provides a place to keep work in progress organized, as well as maintaining a common directory structure from which the development and execution of files can be implemented.



SQL Project Explorer provides a directory structure to organize and maintain SQL files.

When you save a file in SQL Editor, SQL Project Insight is able to determine the type of SQL file you created and automatically organizes it in the proper directory folder in SQL Project Explorer.

Additionally, SQL Project Explorer contains a folder for all data source connections you have indicated when creating a file:

- **Connections** lists all data source connections of the files associated with the project.
- **Creation Scripts** contains DDL statements and statements that define database objects.
- **General SQL** is a category for all other SQL files not used in database object creation. This includes DML files, etc.
- **Large Scripts** contains all files larger than the currently indicated file size preference. This limit can be modified on the **Preferences** panel. (Select **Windows>Preferences** in the Menu Bar.)

Other features include:

- Flagging of files that contain semantic or syntactic errors, letting you quickly scan a file list for erroneous files.
- A code indexer that provides a preview look at code in each file.
- Quick access to the data source associated with a file
- Hierarchical display and access to SQL code.

Additional Resources

Licensing Your Embarcadero Technologies Product

All Embarcadero Technologies products include a 14-day trial period. To continue using the product without interruption, we recommend that you license it as soon as possible. To license your product, use the License Request Wizard found in the Help menu of your respective product. If you have not yet purchased your Embarcadero Technologies product, contact sales@embarcadero.com, or uk.sales@embarcadero.com for sales in the EMEA region.

Embarcadero Technologies Product Support

The Embarcadero Technologies Web site is an excellent source for additional product information, including white papers, articles, FAQs, discussion groups, and the Embarcadero Knowledge Base. Go to www.embarcadero.com/resources, or click any of the links below, to find:

- [Documentation](#)
- [Online Demos](#)
- [Technical Papers](#)
- [Discussion Groups](#)
- [Knowledge Base](#)

Embarcadero Technologies Technical Support

If you have a valid maintenance contract with Embarcadero Technologies, the Embarcadero Technical Support team is available to assist you with any problems you have with our applications. Our maintenance contract also entitles registered users of Embarcadero Technologies products to download free software upgrades during the active contract period. Evaluators receive free technical support for the term of their evaluation (14 days).

We encourage you to open technical support cases via the [Technical Support request form](#) at the [Embarcadero Technologies Web site](#). For additional information about Embarcadero Technologies Technical Support, go to the Support page on our Web site.

Embarcadero Technologies on the Web

To download evaluations of other Embarcadero Technologies products or to learn more about our company and our products visit us at www.embarcadero.com.