



Product Documentation

---

# Rapid SQL® XE and Rapid SQL® 7.7.1

## User Guide

1st Edition August, 2010

---

© 2010 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](http://www.embarcadero.com).

August 16, 2010

**CORPORATE HEADQUARTERS**

100 CALIFORNIA STREET  
12TH FLOOR  
SAN FRANCISCO, CALIFORNIA  
94111 USA

**EMEA HEADQUARTERS**

YORK HOUSE  
18 YORK ROAD  
MAIDENHEAD, BERKSHIRE  
SL6 1SF, UNITED KINGDOM

**ASIA-PACIFIC HEADQUARTERS**

L7. 313 LA TROBE STREET  
MELBOURNE VIC 3000  
AUSTRALIA

# Contents

Welcome to Rapid SQL .....	13
Technical Requirements .....	13
Hardware Requirements .....	13
Operating System Requirements .....	13
Database Support .....	16
IBM DB2 for Z/OS Stored Procedure Requirements .....	16
Licensing .....	17
Licensing Overview .....	17
Licensing Your Application .....	19
Online/Offline Mode and Concurrent License Checkout .....	21
More Information on Using this Document .....	22
Accessing Third Party Documentation .....	22
Shorthand for Third Party Product References .....	22
.....	23
Rapid SQL Tutorial .....	25
Session 1: Getting Started with Rapid SQL .....	25
Overview .....	26
Session 2: Productivity Enhancers .....	27
Creating an Object Using the Object Creation Wizard .....	28
Working With an Existing Object Using the Object Editor .....	28
Working With Code, Files and Data .....	29
Working with Scripts and Files .....	31
Session 3: Scripting .....	33
Generating Code .....	33
Code Assistance .....	35
Session 4: Building a Database Project .....	36
Creating a New Rapid SQL Project .....	36
Adding a Project to Version Control (Sample - Microsoft Visual Source Safe) .....	37

Section 5: Visual Query Builder and Data Editor .....	37
Section 6: SQL Debugging and Profiling .....	39
Additional Evaluation Resources .....	40
Application Basics .....	43
Product Design .....	43
Database Explorer .....	43
Rapid SQL Windows .....	52
Menus .....	65
Toolbars .....	67
Customizing General User Interface Appearance .....	69
Keyboard Shortcuts .....	71
Full Screen Mode .....	74
Specifying Rapid SQL application and feature options .....	75
Browsers Options .....	76
Code Analyst Options .....	77
Connection Options .....	78
Data Editor Options .....	78
Data Transfer Options .....	79
Datasource Options .....	80
DDL Extract Options .....	80
Debug Options .....	82
Directories Options .....	82
Explorer Options .....	83
General Options .....	84
Grid Properties (Results window) Options .....	85
ISQL Options .....	85
Java Options .....	94
JDBC Options .....	94
Logging Options .....	95
MySQL Utilities Options .....	96
Oracle Utilities Options .....	96
Perf Center Options .....	96
Query Builder Options .....	97
Reports Options .....	98
Results (ISQL) Options .....	98
SMTP Mail Options .....	99
Version Control Options .....	100

Warnings Options.....	100
Using.....	103
Datasource Management .....	105
Datasources .....	105
Understanding the Datasource Catalog.....	106
Automatically Discovering Datasources .....	108
Changing Datasource Groups .....	108
Connect .....	109
Disconnect .....	110
Discover Datasource.....	111
Managing Datasources .....	112
Managing Datasource Properties .....	115
Registering or editing datasources.....	115
Selecting Datasources .....	121
Unregistering Datasource .....	122
Datasource Groups .....	123
Removing a Datasource Group.....	123
New Datasource Group .....	123
Rename Datasource Group.....	124
Categorizing datasources using color-coding and labelling .....	127
Supported Objects.....	129
Aliases .....	130
Blob Filters .....	131
Check Constraints .....	131
Clusters .....	132
Database Links .....	133
Databases .....	133
DBRM.....	135
Defaults .....	135
Domains.....	136
Encryption Keys .....	136
Exceptions.....	137
Extended Procedures.....	137
External Functions .....	138
Foreign Keys .....	138
Functions .....	139
Generators.....	140

Indexes .....	140
Instance .....	144
Java Classes .....	144
Java Resources .....	145
Java Sources .....	145
Job Queues .....	145
Libraries .....	146
Logins .....	146
Materialized Query Tables .....	147
Materialized Views .....	148
Materialized View Logs .....	149
Oracle Job Queues .....	149
Outlines .....	150
Packages .....	150
Package Bodies .....	152
Plans .....	152
Primary Keys .....	153
Procedures .....	153
Profiles .....	154
Rollback Segments .....	156
Rules .....	157
Segments .....	158
Sequences .....	159
Shadows .....	159
Snapshots .....	160
Snapshot Logs .....	160
Stogroups .....	161
Structured Types .....	161
Summary Tables .....	162
Synonyms .....	163
Tables .....	163
Tablespaces .....	165
Triggers .....	166
Types .....	167
Type Bodies .....	168
Unique Keys .....	168
User Datatypes .....	169
Users .....	170

Views .....	172
Modifying objects using editors .....	173
Overview and common usage of object editors .....	173
IBM DB2 for Linux, Unix, and Windows Object Editors .....	181
IBM DB2 for z/OS Object Editors .....	199
InterBase/Firebird Object Editors .....	216
Microsoft SQL Server Object Editors .....	228
Oracle Object Editors .....	243
Sybase ASE Object Editors .....	272
Creating objects .....	289
Overview and Common Usage of Object Wizards .....	289
IBM DB2 for Linux, Unix, and Windows Object Wizards .....	293
IBM DB2 for z/OS Object Wizards .....	318
InterBase/Firebird Object Wizards .....	344
Microsoft SQL Server Object Wizards .....	362
Oracle Object Wizards .....	382
Sybase ASE Object Wizards .....	426
Functionality .....	445
Activate Logging .....	447
Add/Modify Database File .....	448
Add Database Fragment .....	449
Add, Insert, or Modify Column .....	450
Add, Insert, or Modify Parameter .....	457
Add, Insert, or Modify Partition .....	458
Add or Modify Cluster Column .....	461
Add or Modify Datafile .....	462
Allocate Extent .....	463
Analyze .....	465
Bind Data Cache .....	469
Bind Package .....	470
Bind Plan .....	473
Build Query .....	478
Change Category .....	478
Change Container Size .....	479
Change Database Comment .....	479
Change Password .....	480
Change Status (Tablespaces) .....	481
Change Status (Triggers) .....	482

Checkpoint .....	482
Clone Partition .....	483
Clone Table .....	483
Coalesce .....	484
Compile .....	485
Connection Editor .....	488
Convert to Partitioned .....	489
Copy Object Names .....	491
Create Alias .....	491
Create Clone .....	493
Create Insert Statements .....	494
Create or Edit Java Source .....	494
Create Like .....	495
Create or Edit User Accounts .....	502
Create or Modify Constraint .....	502
Create Synonym .....	503
Create View .....	505
DataLink Options .....	505
DBCC .....	505
Deallocate Unused Space .....	514
Dependencies .....	516
Detach/Attach .....	517
Disable Index .....	519
Disable/Enable Job Queues .....	520
Disable/Enable Triggers .....	521
Disk Resize .....	522
Drop .....	523
Drop Clone .....	524
Edit Data .....	525
Error .....	525
Estimate Size .....	526
Exchange Data With Clone .....	527
Execute .....	528
Extract .....	531
Extract Data as XML .....	532
Flush Cache .....	532
Free Packages .....	533
Free Plans .....	533



Generate Anonymous Block .....	534
Generating Packages, Procedures, and Statements from Tables and Views .....	534
Index Constraint .....	536
Load Java .....	536
Lock .....	538
Modify Constraint .....	539
Move Log .....	539
Open .....	540
Package Editor .....	541
Partitioned Columns .....	541
Partition Editor .....	542
Partition Upper Bound .....	543
Paste Columns for Check Constraint .....	544
Place .....	545
Preview: Create .....	546
Preview .....	546
Quiesce Tablespaces .....	547
Reassign by Category .....	549
Rebind Packages .....	550
Rebind Plans .....	551
Rebuild Index .....	551
Rebuild Outlines .....	555
Recompile .....	555
Redistribute .....	555
Refresh Summary Table .....	556
Rename .....	557
Reorganize .....	559
Report .....	565
Restart .....	566
Schema .....	566
Select * From .....	567
Set Default .....	568
Set Online/Offline .....	568
Set Statistics .....	569
Set Tablespaces Quota .....	570
Set UNDO Tablespace .....	571
Shrink .....	572
Status .....	573

Summary Definition .....	576
Switch Online .....	576
Transfer Ownership .....	577
Truncate .....	578
Update Statistics .....	582
SQL Scripting .....	589
ISQL Editor .....	589
Results Editor .....	643
Permissions Management .....	661
Explicit Permissions .....	661
Cascading Permissions .....	661
Using Roles to Grant Permissions and Privileges .....	662
Using Roles to Revoke Permissions and Privileges .....	662
Grant Privilege(s) To .....	662
Revoke Privilege(s) From .....	663
Deny Privileges From .....	664
Project Management .....	665
Create a New Project .....	665
Working with Projects .....	673
Version Control .....	689
Version Control Integration .....	689
Version Control Configuration .....	691
Using Version Control .....	694
Tools .....	711
Find in Files .....	712
Database Search .....	713
Database Search Wizard - Panel 1 .....	713
Database Search Wizard - Panel 2 .....	714
Database Search Results .....	714
Script Execution Facility .....	715
File Execution Facility .....	715
Completing the Script/File Execution Facility .....	716
Script Execution Facility - Script Tab .....	716
File Execution Facility - Files Tab .....	717
File/Script Execution Facility - Target Tab .....	717
File/Script Execution Facility - Output Tab .....	718
File/Script Execution Facility - Notify Tab .....	718

Scheduling .....	719
Microsoft Task Scheduler .....	719
ETSQLX Command Line Utility .....	719
Visual Difference .....	720
Comparing Files .....	720
Comparing Database Objects .....	721
Navigating in the Visual Difference Dialog Box .....	722
Printing a Pane of the Visual Difference Dialog Box .....	722
Searching in the Visual Difference Dialog Box .....	722
Setting Options in the Visual Difference Dialog Box .....	723
Query Builder .....	724
Query Builder Design .....	725
Using Query Builder .....	734
Data Editor .....	765
Data Editor Design .....	765
Code Generation Facility .....	774
Import Data .....	775
Import Data Wizard - Panel 1 .....	776
Import Data Wizard - Panel 2 for Text Files .....	777
Import Data Wizard - Panel 3 for Text Files .....	777
Import Data Wizard - Panel 2 for Excel Files .....	777
Embarcadero Products .....	778
Code Workbench .....	778
Enabling the Code Templates and Auto Replace Features .....	778
Maintaining Code Template Definitions .....	779
Maintaining Auto Replace Expressions .....	780
Loading and Saving Custom Code Workbench Settings .....	781
Code Analyst .....	782
Code Analyst DBMS Notes .....	783
Code Analyst Requirements .....	787
Installing Code Analyst .....	787
Uninstalling Code Analyst .....	788
Code Analyst Product Design .....	789
Code Analyst Tutorial .....	791
Using the Code Analyst .....	796

Rapid SQL Add-On Tools .....	805
Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows .....	805
Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Features .....	806
Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Interface .....	809
Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Functionality ...	813
Using the Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows .....	822
Embarcadero SQL Debugger for Microsoft SQL Server .....	824
Embarcadero SQL Debugger for Microsoft Features .....	825
Embarcadero SQL Debugger for Microsoft Interface .....	832
Embarcadero SQL Debugger for Microsoft Functionality .....	835
Using the Embarcadero SQL Debugger for Microsoft .....	843
Debugging a Sample Script .....	845
Embarcadero SQL Debugger for Oracle .....	855
Debugging Features .....	856
Embarcadero SQL Debugger for Oracle Interface .....	858
Embarcadero SQL Debugger for Oracle Functionality .....	861
Using the Embarcadero SQL Debugger for Oracle .....	869
Debugging a Sample Script .....	872
Embarcadero SQL Debugger for Sybase ASE .....	882
Embarcadero SQL Debugger for Sybase Features .....	883
Embarcadero SQL Debugger for Sybase Interface .....	885
Embarcadero SQL Debugger for Sybase Functionality .....	888
Using Embarcadero SQL Debugger for Sybase .....	896
Rapid SQL PL/SQL Profiler .....	898
Setting Up Rapid SQL PL/SQL Profiler .....	899
Rapid SQL PL/SQL Profiler Explorer .....	899
Rapid SQL PL/SQL Profiler Functionality .....	900
Using Rapid SQL PL/SQL Profiler .....	907
Index .....	915

# WELCOME TO RAPID SQL

Rapid SQL is an integrated cross-platform database development environment that provides a highly-intuitive and well-integrated interface. Its graphical facilities simplify SQL scripting, object management, reverse-engineering, database project management, version control, and schema deployment. Additionally, it provides comprehensive tools for tuning and editing code to ensure high-performance and quality.

Rapid SQL offers support for all major databases including IBM DB2, InterBase/Firebird, Microsoft SQL Server, MySQL, Oracle, and Sybase.

See the following topics:

- [Rapid SQL Tutorial](#)
- [Application Basics](#)
- [Using](#)
- [Tools](#)
- [Rapid SQL Add-On Tools](#)

## TECHNICAL REQUIREMENTS

Rapid SQL is a 32-bit application that runs in a Microsoft Windows environment. Before using Rapid SQL, please verify that your environment meets the following requirements:

- [Hardware requirements](#)
- [Operating system requirements](#)

**NOTE:** Users need full registry privileges during the installation and access to the keys under HKEY\_CURRENT\_USER in the registry after installation.

## HARDWARE REQUIREMENTS

Embarcadero Technologies recommends the following minimum hardware requirements:

- 500 MB of memory
- 500 MB of disk space

## OPERATING SYSTEM REQUIREMENTS

Rapid SQL supports the following operating systems.

- Windows 7 (32-bit and 64-bit)
- Windows XP (32-bit and 64-bit)
- Windows Vista (32-bit and 64-bit)

For more detailed information, see the following topics:

- [XP Support](#)
- [Vista and Windows 7 Support](#)
- [32-bit Versus 64-bit Operating System Support and Restrictions](#)

## XP SUPPORT

Windows XP has two user security groups:

- Users
- Power Users

Microsoft intentionally does not grant members of the Users Group the authority to install applications or make global changes to the system. Restricted Users are members of the Users Group. Standard users belong to the Power Users Group. Microsoft grants members of the Power Users Group the authority to install programs. You must be a member of the Administrators Group in order to install and use Embarcadero Technologies applications.

Because Restricted Users are members of the Users Group, they cannot install and run Embarcadero Technologies applications.

**CAUTION:** You must be a member of the Administrators Group in order to install and use Embarcadero Technologies applications.

To open the Group Membership tab where you can determine your group and review the Microsoft security guidelines:

- 1 On the Control Panel, open User Accounts.
- 2 On the Users tab, select a user and then click the Properties button.
- 3 Click the Group Membership tab.

## VISTA AND WINDOWS 7 SUPPORT

Windows Vista UAC and Windows 7 provide two user types:

- Standard user
- Administrator

Rapid SQL can be installed or uninstalled by an administrator or by a standard user using an administrator token. Standard users can run Rapid SQL.

**CAUTION:** For the purposes of running Rapid SQL, default standard user token privileges should not be modified. Modifying standard user token privileges can result in licensing issues which will prevent Rapid SQL from operating properly.

### 32-BIT VERSUS 64-BIT OPERATING SYSTEM SUPPORT AND RESTRICTIONS

Rapid SQL is a 32-bit application that runs emulated on 64-bit versions of Windows. The only requirement is that Rapid SQL supports only 32-bit versions of the client software that must be installed for connectivity to each DBMS that Rapid SQL will work against.

## DATABASE SUPPORT

The table below describes the database platforms Rapid SQL supports and the server and client requirements:

Platform	Server	Client
IBM DB2 for OS/390 (z/OS)	IBM DB2 for OS/390 (z/OS) v7, v8, and v9.	When going through a DB2 Connect gateway, DB2 Client for Windows 6.0 or later. When going directly to the mainframe from the client, DB2 Connect Personal Edition v6 or higher on the client machine.
IBM DB2 for Linux, Unix, and Windows	IBM DB2 Universal Database 8 and 9	IBM DB2 Client for Windows 7.2 or later. NOTE: When using a v8 client, Rapid SQL only supports connecting to a v8 Database on Windows, Linux and Unix.
InterBase/Firebird	InterBase 2007 and InterBase 2009	DataDirect ODBC driver for InterBase and the InterBase Client
	Firebird 2.0	Firebird ODBC Driver
Microsoft SQL Server	Microsoft SQL Server 2000, 2005, and 2008 NOTE: For SQL Server 2005 and 2008, only object types like users, logins, roles, config parameters, etc., are supported. We do NOT support Microsoft .NET extensions at this time.	Microsoft SQL Server Client Library
MySQL	MySQL 4.x. (Rapid SQL is tolerant of 5.x but does not support Stored Procedures, Triggers, Views.)	MySQL ODBC driver 3.5.1 and above
Oracle	Oracle 8i, 9i, 10g, and 11g	Oracle SQL*Net Client
Sybase ASE	Sybase System 12.5, Sybase 15	Sybase Open Client

**NOTE:** Since Rapid SQL is a 32-bit application that runs emulated on 64-bit versions of Windows, only 32-bit versions of the DBMS client software are supported.

## IBM DB2 FOR Z/OS STORED PROCEDURE REQUIREMENTS

When working against an IBM DB2 for z/OS data source, Rapid SQL relies on the following stored procedures, provided as an optional installation step in setting up the DB2 subsystem:

- DSNWZP
- DSNUTILS
- ADMIN\_COMMAND\_DSN
- ADMIN\_COMMAND\_DB2



Prior to using Rapid SQL against an IBM DB2 for z/OS data source, ensure that these components are installed on the server. See <http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2.doc.inst/stpadd.htm#stpadd> for more information.

## LICENSING

Each Embarcadero client application requires one or more licenses in order to run. An Embarcadero product, such as DBArtisan or ER/Studio, has a baseline license which provides basic feature support for that product. In addition, incremental licenses may be required to support specific DBMS platforms, product add-ons, or other functions.

For more information, see the following topics:

- [Licensing Overview](#)
- [Licensing your Application](#)
- [Online/Offline Mode and Concurrent License Checkout](#)

## LICENSING OVERVIEW

The following discussion provides a high-level discussion of key licensing topics and directs you to sources of more detailed information.

[Viewing your License Type and Modules](#)

[Understanding Trial, Workstation, and Networked Licenses](#)

[Rapid SQL License Modules, DBMS Support, and Feature Availability](#)

[Rapid SQL XE License Modules, DBMS Support, and Feature Availability](#)

[Directing Queries Regarding Licenses](#)

## VIEWING YOUR LICENSE TYPE AND MODULES

The About... dialog, available from the Help menu, displays your license type and each individual license module currently registered.

## UNDERSTANDING TRIAL, WORKSTATION, AND NETWORKED LICENSES

Three kinds of licenses are available: Trial, Workstation, and Networked.

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.
----------------	--

Workstation licensing	A license or set of licenses is tied to a particular workstation. The product can only be used on that workstation.
Networked licensing	<p>Networked licenses are administered and distributed by a central License Server (Embarcadero License Center or Acresto FLEXnet Publisher). There are two types of networked licenses: <i>Concurrent</i> and <i>Networked Named User</i>. With Concurrent licensing, users on different machines take turns using licenses from a shared pool. With Networked Named User licensing, licenses are pre-assigned to specific users setup on the license server's user list. Those users are guaranteed to have licenses available any time.</p> <p>NOTE: Concurrent licenses can be borrowed for use without a network connection. For details, see <a href="#">Online/Offline Mode and Concurrent License Checkout</a>.</p>

For a detailed description of licensing options, see <http://www.embarcadero.com/software-licensing-solutions>.

## RAPID SQL LICENSE MODULES, DBMS SUPPORT, AND FEATURE AVAILABILITY

Individual license modules correspond to the DBMS platforms you are licensed to use. In addition, each DBMS license module corresponds to a product edition. The following table indicates feature availability for each license module edition:

Standard	Provides baseline support including datasource management, object management, SQL editing and execution, and standard tools and utilities.
Professional	Standard edition features plus SQL Debugger, SQL Profiler, and Code Analyst.

## RAPID SQL XE PRO LICENSE MODULES, DBMS SUPPORT, AND FEATURE AVAILABILITY

An XE license provides access to all supported DBMS platforms. XE modules and feature availability are as follows:

XE Pro	Provides baseline support including datasource management, object management, SQL editing and execution, and standard tools and utilities, SQL Debugger, SQL Profiler, and Code Analyst.
--------	--

## DIRECTING QUERIES REGARDING LICENSES

Questions regarding license availability, feature availability, and client or server licensing, should be directed as follows:

- If you work in an organization that uses networked licensing, direct any questions to your site's Rapid SQL administrator
- If you are using workstation licensing, direct licensing questions to your Embarcadero Technologies representative

## LICENSING YOUR APPLICATION

See the following topics for details on registering your product:

[Registering a trial or workstation license during installation](#)

[Registering a workstation license after application startup](#)

[Registering by phone](#)

[Registering a networked license](#)

### REGISTERING A TRIAL OR WORKSTATION LICENSE DURING INSTALLATION

Shortly after initiating download of a trial version of an Embarcadero product, you should receive an email with a serial number you must register during installation. Similarly, if you purchase an Embarcadero product while no trial version is active, you will receive a serial number that must be registered during installation.

- 1 Start the installation. An Embarcadero License Registration dialog appears.
- 2 Copy the serial number from the email and paste it in the Serial Number field.
- 3 Enter your Embarcadero Developer Network 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.

Your 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](#).

### 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 the trial period has expired, follow the instructions in [Registering a Trial or Workstation License During Installation](#) instead.

- 1 On the Help menu select About and then on the dialog that opens, click Manage to open a license manager dialog.
- 2 On the Serial menu, select Add Serial Number.
- 3 Copy the serial number from the email and paste it in the Add Serial Number dialog, then click OK.

- 4 Right-click on the serial number you added, and then select Register from the context menu. A registration dialog opens.

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

- 5 Ensure that the Register using Online Registration radio box is selected.
- 6 Provide Developer Network credentials in the Login name or Email and Password boxes. 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.
- 7 Click Register.
- 8 If prompted to restart the application, click Yes.

Your 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](#).

## REGISTERING BY PHONE

If you have problems with either of the above procedures, you can register licenses by phone. You will have to provide 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; and 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 through 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 file to the desktop or a scratch directory such as c:\temp.
- 1 On the Help menu select About and then on the dialog that opens, click Register. A registration dialog opens.
- 2 Select the I have received an activation file (\*.slip or reg\*.txt) radio box.
- 3 Click the Browse button and use the Select License Activation File dialog to locate and select the activation file you installed.
- 4 Click the Import button to import the activation file and when complete, click the Finish button.
- 5 If prompted to restart the application, click Yes.

## REGISTERING A NETWORKED LICENSE

If you work in an organization using Networked licensing, an administrator, department head, or someone providing a similar function will provide you with an activation file.

Once you receive the file, save it to the `license` subfolder of your product's main installation folder (typically `C:\Program Files\Embarcadero\<product><version>\license\`), then restart the application.

No additional steps are necessary.

## ONLINE/OFFLINE MODE AND CONCURRENT LICENSE CHECKOUT

Concurrent licenses can be used in both online and offline modes. In online mode, you must have a continuous network connection to your License Center. Licenses are checked out when you start Rapid SQL and checked back in when you shut Rapid SQL down.

You can also use a license in offline mode. When you explicitly check out a license for offline use, you can use Rapid SQL without a connection to your License Center for a specified duration. This lets you work while travelling or commuting, work away from your primary work area, or use Rapid SQL when a network connection is unavailable or not required.

**NOTE:** Contact your site administrator for information on offline license availability, the maximum duration, offline license policy at your site, or any other issues arising from online license usage.

### To check out a license for offline use

- 1 On the **Help** menu, select **Checkout License**. The **Check Out License For Offline Use** dialog opens.
- 2 Select the check box associated with each individual license you want to check out.
- 3 In the **Checkout Duration** box, type the number of hours that you can use the offline license without a network connection to the License Center.
- 4 Click **OK**.

You can work offline for the specified duration. The duration period begins immediately.

If you subsequently establish a network connection to the License Center before the license duration expires, you can indicate to the License Center that the offline license is no longer required.

### To indicate that an offline license is no longer required

- 1 On the **Help** menu, select **Checkin License**.

There is no interruption in Rapid SQL usage. The license is not actually checked in until you shut down Rapid SQL.

## MORE INFORMATION ON USING THIS DOCUMENT

The following topics provide additional information that is useful when working with product documentation:

- [Accessing third party documentation](#)
- [Shorthand for third party product references](#)

## ACCESSING THIRD PARTY DOCUMENTATION

Many Rapid SQL features provide support for functionality available in the supported third party DBMS platforms. In object management for example, properties available when creating or editing objects and actions available against object types, are direct equivalents of clauses, options, or keywords available for use with the third party equivalent.

In such cases, no attempt is made to duplicate detailed documentation by third parties. Casual descriptions are provided, noting the third party equivalent. You can consult the third party documentation for details or clarifications. The following links provide access to online documentation for DBMS platforms supported by Rapid SQL:

- [IBM DB2 for Linux, Unix, and Windows documentation](#)
- [IBM DB2 for z/OS documentation](#)
- [InterBase/Firebird documentation](#)
- [Microsoft SQL Server documentation](#)
- [MySQL documentation](#)
- [Oracle documentation](#)
- [Sybase documentation](#)

## SHORTHAND FOR THIRD PARTY PRODUCT REFERENCES

To save space in headings and table columns, the following shorthand is used to represent versions and variations of the DBMS platforms supported by Rapid SQL.

Shorthand	DBMS
DB2 LUW	IBM DB2 for Linux, Unix, and Windows
DB2 z/OS	IBM DB2 for z/OS
ITB/FBD	InterBase/Firebird
SQL SVR	Microsoft SQL Server
MySQL	MySQL
Oracle	Oracle

Shorthand	DBMS
Sybase	Sybase





# RAPID SQL TUTORIAL

Embarcadero Rapid SQL® is an integrated development environment that enables developers to create, edit, version, tune, and deploy server-side objects residing on Microsoft® SQL Server, InterBase/Firebird, Oracle®, Sybase Adaptive Server®, IBM® LUW DB2®, and IBM® DB2® for OS/390® databases. Its unified database development environment provides extensive graphical facilities that simplify SQL scripting, object management, reverse engineering, database project management, version control, and schema deployment. With Rapid SQL, programmers can develop and maintain high-quality, high-performance client/server and Web-based applications in less time, and with greater accuracy.

This tutorial is intended to help you get started using Embarcadero Rapid SQL. After completing this tutorial, you will have the foundation you need to explore the many features of Rapid SQL. You will have learned how to register and connect cross-platform datasources; navigate the database Explorer; work with the individual object browsers, editors, and wizards; build and manage projects; and leverage many of the productivity-focused features offered throughout Rapid SQL. Rapid SQL allows you to concentrate *more on what needs to be done and less on how it should be done*.

This tutorial is divided into six sessions. Do them all at once, or complete them individually as you have time.

- [Session 1: Getting Started with Rapid SQL](#)
- [Session 2: Productivity Enhancers](#)
- [Session 3: Database Object Management Made Easy](#)
- [Session 4: Building a Database Project](#)
- [Session 5: Visual Query Builder and Data Editor](#)
- [Session 6: SQL Debugging and Profiling](#)

You can use this tutorial as a road map of product highlights, and to help familiarize yourself with key product features.

Once you begin, select Help from the Toolbar to browse additional online help resources, including documentation that complements and builds upon the activities outlined in this tutorial.

## SESSION 1: GETTING STARTED WITH RAPID SQL

### Download and Install

You can obtain the latest version of the Rapid SQL software from the Embarcadero Web site at <http://www.embarcadero.com/products/rapid-sql>.

Provide the requested information and follow the steps indicated to download and install the software. When you first install an evaluation copy of Rapid SQL, you can use the tool for 14 days once you register your trial license. After that time, a permanent license is needed.

Proceed to [Overview](#).

## OVERVIEW

Rapid SQL allows you to view and manage datasources via the Database Explorer. The Database Explorer enables navigation between different datasource platforms from the same window. Additionally, the Rapid SQL environment enables multiple workspaces to be opened at the same time, allowing you to work while application processes run in the background.

### Start Rapid SQL

- 1 On the Start menu, navigate to Programs, Embarcadero RapidSQL 7.7.0, and select Rapid SQL 7.7.0.

The first time Rapid SQL starts, a dialog appears prompting a search for pre-existing datasources that have been registered in other Embarcadero applications. Selecting Yes prompts Rapid SQL to search and locate any datasources defined in the datasource catalog (defined on the machine identified in the Options field). For the purposes of this tutorial, select No.

- 2 Click OK to continue.

Rapid SQL provides a feature that automatically searches the DBMS configurations on the network or installation machine and finds datasources residing on the system that are not currently registered. The Discover Datasources dialog contains a list of these datasources that includes the name of the datasource server or instance, and the type of DBMS. Use the Discover Datasources dialog to register datasources in Rapid SQL.

### Registering Cross-Platform Datasources

Datasource registration is performed via the Datasource Registration Wizard. It opens automatically via the Discover Datasources dialog (when you click OK), and can also be opened manually from Rapid SQL's Main Menu. (Datasource>Register Datasource.)

### Registering a Datasource

Rapid SQL provides the same interface for registering all supported DBMS platforms. Each datasource's connection information and log on details only need to be set up once per datasource instance. This connection information can be saved locally, or in a common datasource catalog for use in other Embarcadero products requiring datasource registration.

You can configure Embarcadero products to use a datasource catalog stored in the system registry of your machine (locally), or use a datasource catalog located in the registry of a remote machine. This capability provides the ability to easily share datasource catalogs among multiple users and workstations, and creates a highly-maintainable development environment.

All Embarcadero database management products share the datasource catalog. You can set up your datasource catalog using one product (such as Rapid SQL), and the same list of datasource will be made available in other Embarcadero Technologies products. Any changes you make to the datasource catalog are reflected in all Embarcadero database management products.

To register a Microsoft SQL Server datasource

- 1 From the **Datasource** menu select **Register Datasource**.
- 2 On the first panel of the Datasource Registration Wizard, select Microsoft SQL Server as the database type and then click **Next**.
- 3 Use the controls on the **Connection Information** panel to select a Microsoft SQL Server datasource on your network and then click **Next**.
- 4 On the **Security Parameters** panel, provide valid credentials in the **User ID** and **Password** boxes and select the Auto-Connect? to avoid having to provide credentials each time you connect to this datasource. Click Next to continue.
- 5 Since the **Datasource Properties** panel references an advanced feature, click **Next**.
- 6 In the Datasource Group tree, indicate where the datasource you are currently registering will sit. Click Finish to complete the Wizard and when prompted, connect to the datasource. The datasource is registered with the instance of Rapid SQL, and it can now be manipulated via the Rapid SQL interface.

Proceed to [Session 2: Productivity Enhancers](#).

## SESSION 2: PRODUCTIVITY ENHANCERS

### Database Object Management Made Easy

Rapid SQL provides the Database Explorer Tree to provide an easy and intuitive way to drill down into atomic database objects. The Database Explorer Tree displays all registered datasources and serves as the entry-point to much of Rapid SQL's advanced functionality.

### Viewing Datasources and Datasource Objects in the Explorer Tree

- 1 On the Explorer, select By Object Type.  
The Explorer displays objects organized by type.
- 2 Under the datasource you registered, expand the Databases node and drill down further by expanding the Pubs node. The datasource object types for that database are displayed.
- 3 Expand the Tables node to view a list of table objects contained in the SAMPLE\_DATASOURCE database.

Proceed to [Creating an Object Using the Object Creation Wizard](#).

## CREATING AN OBJECT USING THE OBJECT CREATION WIZARD

You can create database objects via Object Creation Wizards accessed from the Database Explorer. The following steps provide an example of how to use the Table Object Creation Wizard to create a new Oracle table.

### Using the Table Object Creation Wizard

- 1 Right-click the Tables node and select New from the menu. The Create Table Wizard appears.
- 2 Enter the appropriate information on the panels of the wizard:
  - On the **Properties** panel, select a Schema, provide a table name of SAMPLE\_TABLE, and optionally modify other properties.
  - On the **Columns** panel, type a column name of sample\_column1, and optionally modify other column attributes. Experiment with the **Add Column** and **Delete** buttons.
  - Proceed through the **Indexes**, **Constraints**, and **Permissions** panels, optionally making changes.
  - On the **DDL View** panel, inspect the DDL generated by your choices and when ready click Execute. Your new table is created.

**NOTE:** Rapid SQL builds syntactically-correct code that is ready to run the first time it is generated. Manual coding is not required in any of the Object Creation Wizards.

In the case of tables, Rapid SQL automatically opens an editor on the newly created object.

- 3 For now, close the object editor window.

Proceed to [Working With an Existing Object Using the Object Editor](#).

## WORKING WITH AN EXISTING OBJECT USING THE OBJECT EDITOR

Once an object is created, it can be modified via the same Object Editor. Object Editors perform modification operations that would normally require painstaking and error-prone scripting, such as deleting or inserting columns in a table while preserving data, dependencies, and permissions. Rapid SQL analyzes the database catalog to determine its structure, and then automatically generates the appropriate SQL script required for the extended alteration.

For example, when a full table alteration is required, Rapid SQL automatically unloads and reloads the data, eliminating most of the tedious work a manual alteration would require.

When an object is created, Rapid SQL automatically opens an individual Object Editor for the newly-created object.

**To modify an existing object**

- 1 On the Explorer, navigate to the Tables node of the datasource defined previously.
- 2 Right-click the table you created in the previous exercise and select Open from the context menu.

Rapid SQL opens the Table Object Editor and you can proceed to modify the selected object.

- 3 Inspect the editor and note that it has tabs not available on the wizard you used to create the table. Close the editor window.

Proceed to [Object Documentation and Reporting](#).

**OBJECT DOCUMENTATION AND REPORTING**

Rapid SQL can provide HTML reports for all database objects. The reports are automatically displayed in the Rapid SQL application workspace, but can be saved as a new file, or referenced in the original file. All HTML reports are browser-ready and suitable for posting directly to the Web.

**To generate an HTML report**

- 1 On the Explorer, open any datasource and navigate to the Tables node.
- 2 Right-click an object and select Report from the menu. The Report dialog appears.
- 3 Enter the appropriate information in the fields provided:
  - In the Report Home Page File Name field, type the directory path and file name of the report manually, or click Browse and navigate to the path and file name. This value supports network Web server paths.
  - In the Report Title dialog, type the name of the report.
- 4 Click Execute. The HTML report is automatically generated to the path defined in the Report Home Page File Name field in Step 3 and displayed in the Rapid SQL application workspace.

Proceed to [Working With Code, Files and Data](#).

**WORKING WITH CODE, FILES AND DATA**

Rapid SQL provides many features and powerful development tools for creating and executing SQL code and working with data. This section of the tutorial will provide a high-level functionality overview of the major development features and tools within Rapid SQL.

## Setting Environment Options

The Options Editor provides an interface for you to set Rapid SQL environment parameters to meet your development requirements.

### To set development environment parameters

- 1 Select File>Options. The Options Editor dialog opens.
- 2 Click on the General Tab. General settings are displayed and modified by selecting them from the dropdown menu or via the tabbed window.
- 3 Adjust the development parameters as required and click OK. The environment options are automatically applied to all registered datasources.

## Using the Favorites Tab

The Favorites Tab provides a drag-and-drop library interface of all supported DBMS syntax, SQL syntax, built-in functions, optimizer hints, and SQL-conditional syntax. Additionally, it provides the ability to create custom folders to store commonly-used code for quick and efficient access, as needed.

### To open the Favorites Explorer

- 1 At the bottom of the Explorer pane, click Favorites. The Favorites Explorer appears.
- 2 Expand the Microsoft node and then expand the Schema subnode.
- 3 Right-click the Procedures node and select Open. The selected code opens in the SQL Editor window and is ready for execution.

### To add a custom folder to the Favorites Explorer

- 1 Open the Favorites Explorer and right-click the Favorites folder.
- 2 Select New Folder from the menu. A new folder is added to the bottom of the tree view.

### To add an existing script to the Favorites Tab

- 1 Select **New>SQL**, create a two or three-line script in the SQL editor, and then save your script.
- 2 After saving the script, right-click anywhere on the workspace and select Add to Favorites. The Favorite Properties dialog box opens.
- 3 Enter the appropriate information in the fields provided. The script is added to the Favorites Explorer.

**NOTE:** Code from the Paste SQL, Paste SQL Syntax, and Favorites Tab are available from any SQL or DDL Editing window.

Proceed to [Working with Scripts and Files](#).

## WORKING WITH SCRIPTS AND FILES

Rapid SQL extends the auto-generation of SQL code by allowing you to run your scripts across multiple databases at the same time. Additionally, Rapid SQL provides the option to execute code immediately, or schedule it to run later.

### Using the Script Execution Facility

- 1 On the Favorites Explorer, navigate to the Microsoft SQL Server node and select the Schema subnode.
- 2 Right-click the Procedures node and select Script Execution Facility. The Script Execution Facility dialog box appears.
- 3 Click the Target Tab and select the datasources against which the script will run.
- 4 Click the Output Tab and select a type of output format. For the purposes of this example, select Graphical Output.
- 5 If you want to send a notification message indicating that the script has executed, click the Notify Tab and enter the appropriate information in the fields provided.
- 6 Click Execute. Rapid SQL runs the script against the datasources specified on the Target Tab.

**NOTE:** Separate script output windows are created for each selected datasource.

### Using the File Execution Facility

Similar to the Script Execution Facility, files containing SQL scripts can be added to the File Execution Facility and executed immediately or scheduled to run later. Other than code's origin, all supporting functionality is the same.

- 1 On the Tools toolbar, click File Execution Facility.  
Rapid SQL opens the File Execution Facility dialog box.
- 2 To locate the file you want to execute, click Add. Use the **Select Files** dialog box to locate and select a file.
- 3 On the Target Tab, select the datasources to run the script against.
- 4 On the Output Tab, select the desired output option. For the purposes of this example, select Graphical Output.

**NOTE:** To enable the scheduling function for the script, you must select the File Output option.

- 5 If you want to send notification that the script has executed, on the Notify Tab, complete the target information.

- 6 Click Execute.

Rapid SQL runs the script against the target datasources.

**NOTE:** Separate script output windows are created for each selected datasource.

Proceed to [Viewing Data](#).

## VIEWING DATA

Rapid SQL provides several options for browsing data. In addition, it gives you the ability to construct even the most complex SQL statements with point-and-click ease.

### Select \* Browsing (SQL Server Example)

- 1 On the Database Explorer ensure that the **Explorer** tab is displayed and then expand the MS SQL server node.
- 2 Expand any database you know has table data, expand the Tables node, right-click a table, and then click SELECT \* FROM.

All columns and rows from the table are displayed in the active workspace.

### Retaining Datasource Explorer View Settings

- 1 On the Explorer dropdown, select Retain View Setting.

The next time you open the Explorer, it will open just as you left it. All connections that were present when you closed Rapid SQL will be reestablished.

### Datasource Explorer Bookmarks

Rapid SQL allows you to set bookmarks for frequently visited database objects.

- 1 On the Datasource Explorer, right-click any node.
- 2 Select Add Bookmark.

You can modify the bookmark name as desired.

After Bookmarks are defined you can use them to easily navigate to commonly used datasource resources via the main menu Bookmarks item.

You can use the Bookmark Manager to maintain your Bookmarks.

- 1 Select Bookmarks>Bookmark Manager.

The Bookmark Manager dialog box opens where you can rename, move, and delete bookmarks.



## Setting Keyboard Shortcuts and Hot Keys

- 1 Select Tools>Customize.

The Customize dialog box opens.

- 2 In the Customize dialog box, open the Keyboard Tab.
  - The Keyboard Tab can be used to set Keyboard shortcut hot keys for all areas of Rapid SQL functionality.

## Referencing Most Recently Used Datasources

- 1 From the File menu, select Recent Datasources, and then choose a datasource.

This will automatically place you on the datasource within the Explorer, ready to work with an active connection.

Proceed to [Session 3: Scripting](#).

# SESSION 3: SCRIPTING

By providing several code generation and assistance options, Rapid SQL makes cross-platform development easy for developers of all experience levels.

**NOTE:** The following examples build on the MS SQL Server SAMPLE\_DATASOURCE registered earlier in this tutorial. These examples can be applied to any registered datasource for any of the supported platforms.

Proceed to [Generating Code](#).

## GENERATING CODE

The Code Generation Facility can be used to create complete procedures, functions or packages revolving around views or tables.

### To open the Code Generation Facility

- 1 From the Tools toolbar, click Code Generation Facility.  
Rapid SQL opens the Embarcadero Code Generator dialog box.
- 2 In the datasource list box, select the target datasource from the drop-down list.
- 3 In the database list box, select pubs.
- 4 In the Tables/Views pane, click Tables, and then select authors.
- 5 In the Select 1 or More Where Clause Columns box, select state.

- 6 In the Select 1 or More Output Columns box, select all columns.
- 7 In the Generate box, choose Select (this is the code option).
- 8 In the Provide an Output File Name, type or browse and locate a target file where you want to save the generated script, and then click Open.
- 9 Click OK.

The DDL to create the procedure is generated and displayed in an editable window called the DDL Editor. You can edit the name of the new procedure and any of the generated code at this time.

- 10 Name the new procedure sample\_select\_authors.
- 11 To submit the DDL and create the procedure, click Execute or Step Execute in the toolbar above the editor.

The indicated file will be saved on the selected directory.

**NOTE:** No SQL statement coding is required to generate complete stored procedures and packages. If applicable, Rapid SQL allows all generated code to be previewed and edited to fit any development need.

### Right-Click Feature

Similar to the Code Generation Facility, the right-click code generation feature can be used to create complete procedures, functions, or packages revolving around views or tables.

- 1 On the Database Explorer, expand the MS SQL server node.
- 2 Expand any database, and then expand the pubs node.
- 3 Expand the Tables node, and right-click the authors table.
- 4 From the pop-up list, select Generate, Procedure, and then Select.
- 5 In the Input Columns pane, select state.
- 6 In the Output Columns pane, leave all output columns selected.
- 7 Click OK.

The DDL to create the procedure is generated and displayed in an editable window, called the DDL Editor. You can edit the name of the new procedure and any of the generated code.

To submit the DDL and create the procedure, click Execute or Step Execute.

**NOTE:** No SQL statement coding is required to generate complete stored procedures and packages. If applicable, Rapid SQL allows all generated code to be previewed and edited to fit any development need.

Proceed to [Code Assistance](#).

## CODE ASSISTANCE

Rapid SQL provides extensive, easy-to-use code assistance features for all supported DBMS platforms, throughout the application. Assistance is provided in the form of ready-to-use code templates and blocks of syntactically correct code.

### Paste SQL

- 1 On the Database Explorer, expand the MS SQL server node.
- 2 Expand any database, and then expand the pubs node.
- 3 To open a SQL Editor window, from the main toolbar, select File, New, and then SQL (Ctrl+N).
- 4 On the Edit toolbar, select Paste SQL Statement.  
The Paste SQL dialog box opens.
- 5 In the Tables pane, select the authors table.
- 6 In the Columns pane, select all columns.
- 7 In the Statement column, click Select as the generation option.
- 8 To copy the generated code to the SQL Editor workspace, click Paste Statement.

**NOTE:** You can use the statement as is, or modify the code as needed.

### Paste SQL Syntax

- 1 On the Database Explorer, expand the MS SQL server node.
- 2 Expand any database, and then expand the pubs node.
- 3 To open a SQL Editor window, from the main toolbar, select File, New, and then SQL (Ctrl+N).
- 4 On the Edit toolbar, select Paste SQL.  
The SQL Syntax for SQL Server dialog box opens.
- 5 In the SQL Statements: pane, click a target template.  
The template syntax appears in the Syntax pane.
- 6 To copy the code into the SQL Editor workspace, click Paste.

**NOTE:** You can add your own code to complete the needed operation.

Proceed to [Session 4: Building a Database Project](#).

## SESSION 4: BUILDING A DATABASE PROJECT

Rapid SQL provides an excellent team development environment that allows you to reverse engineer live database objects into off-line SQL source code files that can then be easily added to a version control system (VCS). Rapid SQL's seamless VCS integration offers all version control operations, such as get, check-out, check-in, history, and diff.

Proceed to [Creating a New Rapid SQL Project](#).

### CREATING A NEW RAPID SQL PROJECT

The example below reverse engineers the table objects from the Microsoft SQL Server pubs database into a Rapid SQL project, and then adds the project to version control using Microsoft Visual Source Safe.

To create a new project:

- 1 On the main toolbar, select File, New, and then Project.

The New Project dialog box opens.

- 2 In the Name text box, type sample\_project.
- 3 In the Location text box, type or browse and locate a directory that contains a VSS database.
- 4 Optionally, enter a description.
- 5 Under Initialize New Project, select From Database.
- 6 Click OK.

The New Project Reverse Engineering Wizard opens.

- 7 In the Select a Datasource box, select the datasource you created.
- 8 Click Next.

The next panel of the wizards opens.

- 9 In the Select Database to reverse-engineer window, select pubs.
- 10 Click Next.

The next panel of the wizard opens.

- 11 From the Owners list box, select dbo.
- 12 In the Object Type box, right-click to unselect all options, and then select Tables.
- 13 In the Extract Scope box, select Selected Objects Only.
- 14 Click Next.

The next panel of the wizard opens.

15 In the Object Type box, select Tables from the drop-down list, and then select the authors, discounts, and employee tables.

16 In the Options for Tables box, right-click to Unselect All.

17 Click Next.

The next panel of the wizard opens.

18 For Object Ownership, select Retain.

19 Click Next.

The Preview panel of the wizard opens so you can review the project you are creating.

20 Click Execute.

You have now successfully created a working database project. At this point, the project is available to be added to Version Control. This can be done by clicking Yes when prompted and following the dialog, or by right-clicking the project on the Project Explorer. Select Yes for the purpose of this example.

Proceed to [Adding a Project to Version Control \(Sample - Microsoft Visual Source Safe\)](#).

## ADDING A PROJECT TO VERSION CONTROL (SAMPLE - MICROSOFT VISUAL SOURCE SAFE)

When a project is created Rapid SQL will automatically prompt you to add the project files to the selected VCS (see Setting Environment Options below). The dialog you see is displayed below:

1 In the Visual SourceSafe Login dialog box, enter the Username and Password.

2 Type or browse and locate the target database, and then click Open.

The Add to SourceSafe Project dialog box opens.

3 To name your project, in the Project text box type "sample\_project".

4 Click OK.

You receive the message below to indicate that your project was successfully placed in the VSS.

Proceed to [Section 5: Visual Query Builder and Data Editor](#).

## SECTION 5: VISUAL QUERY BUILDER AND DATA EDITOR

Rapid SQL gives you the ability to construct complex SQL statements with point-and-click ease using the Visual Query Builder.

### Visual Query Builder

- 1 On the Database Explorer, expand the MS SQL Server node, and then select pubs.
- 2 Expand the pubs node, and then select Tables.
- 3 Right-click the authors table and select Build Query.

The authors table is automatically added to the Query Builder workspace.

- 4 Right-click the titleauthor table and select Add.

**NOTE:** Note that the tables are automatically identified as being joined by any columns with the same name and datatype.

- 5 To expose the Visual Query Building clauses and options, open the DML Tab.
- 6 Right-click on any clause to initiate a change to the clause.
- 7 Click the Execute button on the toolbar above the Build Query workspace.

The query executes and the lower window displays the results. The generated SQL is also available in the lower window.

**NOTE:** Any visual query builder session can easily be saved to a file for later use.

### Live Data Editor

- 1 On the Database Explorer, expand the MS SQL Server node, and then select pubs.
- 2 Expand the pubs node, and then select Tables.
- 3 Right-click the authors table and select Edit Data.  
Rapid SQL opens the Data Editor Filter dialog box.
- 4 In the Columns box, select Add All. Optionally, you can add a WHERE clause that will filter for only the desired data.
- 5 Click **OK**. The editing window opens.

**NOTE:** Rapid SQL builds the SQL to retrieve the data to be edited in the lower window.

- The editing window has LIVE and BATCH modes. LIVE mode commits your changes each time you move to a new row. BATCH mode will allow you to move within the window and commit your changes when ready.
- Changes made in BATCH mode can be cancelled by selecting Reload Data.

Proceed to [Section 6: SQL Debugging and Profiling](#).

## SECTION 6: SQL DEBUGGING AND PROFILING

### SQL Debugging

The SQL Debugger is another database productivity tool that lets you debug SQL Server, Oracle, Sybase or IBM DB2 LUW for Open Systems stored procedures as well as Oracle functions. SQL Debugger simplifies the task of finding coding errors.

- 1 On the Database Explorer, expand the MS SQL Server node, and then select pubs.
- 2 Expand the pubs node, and then select Procedures.
- 3 Right-click the target Procedure and select Debug.
  - If the procedure/function contains input variables, the Procedure Execution window prompts for entry of these values.
- 4 Enter the value(s) for the input variable(s) and press Continue.

**NOTE:** Pressing the Debug button or selecting Debug from the Command menu invokes The SQL Debugger.

**TIP:** Rapid SQL allows the user to save the input variable values to a file for later use. This is very helpful for procedures/functions with many input variables that need to be run repeatedly.

The application opens the SQL Debugger Interface.

### SQL Profiling- Oracle Only

The SQL Profiler within Rapid SQL provides the ability to capture the metrics of various PL/SQL programmable objects as they are executed in the database. It quickly identifies performance bottlenecks by first calculating the overall runtimes of objects like Oracle packages, and then computing the amount of time each line of PL/SQL code spends executing. Information is presented in an easily viewed, drill-down format.

To start a profiling session:

- 1 From the Tools menu select SQL Profiler, Start.

**OR**

From the SQL Profiler toolbar, click the Start Profiling button Execute.

The Profiler Start dialog box opens.

- 2 Enter a name for the profiling session, or select an existing name from the drop-down list.
- 3 Click OK.

The Profiler session is now active.

- 4 Select the programmable object you want to capture metrics for, and click Execute.

- 5 When finished, from the SQL Profiler toolbar, click Stop.

The SQL Profiler Stop dialog box opens.

- Stop & Analyze – Stops the profiling session and presents the detail profiling information.
- Stop – Stops the profiling and leave the user on the current screen.
- Cancel – Cancels the stop command.
- Help – Displays SQL Profiler Help

- 6 Click Stop.

- 7 Expand the PL/SQL Code Profiling section.

- 8 Right-click the profile session, and then select Run Summary.

The Run Summary appears.

- 9 Right-click a session and select Run Detail.

The Run Detail screen appears and allows you to view the metrics for the run in both a graphical and text format.

- 10 To drill down further into the data, right-click a unit, select Unit Detail.

**NOTE:** You can scroll through the Source window to view the execution times for each statement.

To compare two cases:

- 1 From the Run Summary screen, select the two cases you wish to compare (SHIFT+CLICK to select the second case), and then right-click and select Compare.

## ADDITIONAL EVALUATION RESOURCES

Embarcadero Technologies provides a variety of resources to help support your evaluation and selection of a development tool for your organization.

### Web site

Visit our Web site for current product and company information, educational materials and supporting information. Visit [www.embarcadero.com](http://www.embarcadero.com)

### Electronic Documentation

Detailed reference documentation is available on the Rapid SQL Evaluation CD or online at [www.embarcadero.com/support](http://www.embarcadero.com/support)



**Online FAQ**

The Rapid SQL online FAQ provides answers to commonly asked questions regarding licensing, installation and other helpful topics.

**E-mail Support**

You can contact Rapid SQL support engineers, consultants and engineers directly by sending inquiries to [support@embarcadero.com](mailto:support@embarcadero.com)

**Telephone Support**

We encourage you to call us anytime you would like help or have questions during your evaluation. Please call 415.834.3131 ext. 2, Monday to Friday, 6:00am - 6:00pm PST, Saturday and Sunday, 8:00am - 5:00 PST.



# APPLICATION BASICS

Application Basics is designed to situate you within the application and to provide information about what Rapid SQL offers in the way of design, navigation, and application features. The information presented here is high-level and conceptual. For information on how to use Rapid SQL, see [Using](#).

Application Basics is divided into two sections, the table below describes each section:

Section	Description
<a href="#">Product Design</a>	This section describes the Rapid SQL user interface.
<a href="#">Specifying Rapid SQL application and feature options</a>	This section describes how to customize Rapid SQL's configuration to suit your specific needs.

## PRODUCT DESIGN

The Rapid SQL window opens with the [Database Explorer](#) on the left, the [Workspace](#) on the right, and all [toolbars](#) docked at the top of the application. The [Output Window](#) is not automatically displayed. Rapid SQL also offers you a number of desktops, or workspaces, that you can toggle among while you work.

## DATABASE EXPLORER

Rapid SQL organizes the wealth of information pertaining to your servers through its Database Explorer. The Database Explorer provides a fast and efficient way to access your database objects and scripts. The Database Explorer is a separate window containing a tree object that you can select and expand. The tree object organizes and nests subjects as branches. By expanding or collapsing the tree, you can efficiently browse multiple datasources. The Database Explorer window is dockable so that you can maneuver through the application efficiently.

The Database Explorer includes three tabs:

- [Explorer](#)
- [Favorites](#)
- [Project](#)
- [VC Files](#)

## EXPLORER TAB

The Explorer Tab provides a visual method for browsing, accessing and manipulating your database objects. The Explorer Tab lets you:

- Connect, disconnect, and browse the objects in the supported datasources on your network.
- Drag objects to the Rapid SQL workspace.
- Create new objects.

**NOTE:** You can assign datasources to defined categories such as production or development. This lets you color-code and label Datasource Explorer user interface elements associated with the datasource. For details, see [Categorizing datasources using color-coding and labelling](#).

## Datasource Node

When you click the datasource node of the Explorer Tab tree, Rapid SQL lists all the databases available for that datasource. When you click a database node, you can view all the database object types available for that datasource. When expand an object type node, Rapid SQL displays available objects. You can collapse any portion of a datasource to concentrate on a particular portion of your database.

For more information, see:

[Organizing the Explorer](#)

[Creating New Objects from the Explorer](#)

[Extracting DDL from the Explorer](#)

[Displaying Dependencies from the Explorer Tab](#)

## ORGANIZING THE EXPLORER

Rapid SQL contains functionality that lets you configure how objects are organized in the [Explorer Tab](#). The table below explains the available options:

Functionality	Description
Organize by Object Type	Select to display objects by object type for all users in the same list. This display mode cuts performance in databases that contain many objects.
Organize by Owner	Select to display objects by owner. Most efficient if you are working with databases containing a high number of objects.
Show Only My Objects	Select to display the objects you own in the Explorer Tab. Available if you are organizing the Explorer Tab by object type.
Show System Objects	Select to display system objects.
Full Refresh	Select to refresh.
Expand All Groups	Select to expand all groups.
Collapse	Select to collapse all datasources or collapse all groups.
Retain Group View Settings	Select to retain the current state of the Explorer Tab so that the it opens the same way the next time you start Rapid SQL.

**TIP:** You can also set these options on the Explorer tab of the Options Editor. For details, see [Explorer Options](#).

### Organizing the Explorer by Object Owner

- 1 On the *Explorer* Tab, click the *Explorer* list above the list of datasource groups.
- 2 Click *Organize by Owner*.

Rapid SQL dynamically reorganizes the display of the Explorer Tab, sorting database objects by object owner.

### Organizing the Explorer by Object Type

To organize the Explorer by Object Type, do the following:

- 1 On the *Explorer* Tab, click the *Explorer* list above the list of datasource groups.
- 2 Click *Organize by Object Type*.

Rapid SQL dynamically reorganizes the display of the Explorer Tab, sorting database objects by type.

### Showing or Hiding System Objects

To show or hide system objects, do the following:

- 1 On the *Explorer* Tab, click the *Explorer* list above the list of datasource groups.
- 2 Click *Show System Objects* to show all system objects.
- 3 Click the *Show System Objects* again to hide all system objects.

## CREATING NEW OBJECTS FROM THE EXPLORER

Rapid SQL lets you create new database objects from the [Explorer Tab](#). When you create an object from the Explorer Tab, Rapid SQL does the following:

- Automatically brings up the Object Attachment Facility which associates the new object with an object type.
- Opens a DDL Editor containing a shell script based on the object type you select.
- Automatically ties that object to the current database.

### Creating a New Object from the Explorer

To create a new object from the Explorer, do the following:

- 1 In the *Explorer* Tab, right-click the object type node for the object to create and then click *New* to open the wizard associated with that object.

The Object Type Information dialog box displays. The object type and owner are automatically assigned based on the object type you selected in the previous step.

- 2 Enter the name of the object.
- 3 Click *OK*.

Rapid SQL opens an editor containing a template for the object type.

## EXTRACTING DDL FROM THE EXPLORER

Rapid SQL's [Explorer Tab](#) provides a quick, convenient method for extracting object DDL. Rapid SQL lets you extract the DDL for multiple or individual database objects. SQL scripts are the text of statements that are used to create database objects. Rapid SQL also lets you run SQL scripts against other databases to recreate the database objects. You can also extract scripts to create a record of how an object was created. Rapid SQL lets you extract the CREATE statements for any object into a DDL Editor. If you are extracting the DDL of an Oracle procedural logic object, an SQL Editor automatically opens instead, saving you the trouble of inserting PL/SQL tags.

Rapid SQL offers two ways to extract DDL from the Explorer Tab:

### Using the Shortcut Menu

- 1 Select one or more objects in the *Explorer* Tab. To select contiguous objects, hold SHIFT and select multiple objects. To select noncontiguous objects, hold ALT and select specific objects.
- 2 Right-click to display the shortcut menu and then click *Extract*.

## Dragging Objects

- 1 Select one or more objects in the *Explorer* Tab. To select contiguous objects, hold SHIFT and select multiple objects. To select non-contiguous objects, hold ALT and select specific objects.
- 2 Drag the script(s) to the desired position in the Rapid SQL workspace.

## DISPLAYING DEPENDENCIES FROM THE EXPLORER TAB

Rapid SQL lets you open a result set window, directly from the [Explorer Tab](#), to display dependencies for an object.

### Displaying Dependencies from the Explorer Tab

- 1 On the *Explorer* Tab, double-click the object node.  
Rapid SQL displays the list of objects.
- 2 Right-click the target object and then click Dependencies.  
Rapid SQL opens the Dependencies window which lists all dependent objects.

**NOTE:** Rapid SQL does not retrieve dependencies if objects are created out of order.

## FILTERING IN THE EXPLORER


For ease of navigation, Rapid SQL lets you filter the Explorer tree to display only selected nodes and objects. Rapid SQL offers the following filtering options:

- Real-time modification of the tree to show only nodes with names matching a typed character string. For details, see [Simple, on-the-fly, character-based tree filtering](#).
- Permanent object filters, enabled and disabled for individual data sources, that restrict display based on object name and owning schema (and optionally, object type). For details, see [Object Name/schema Filtering](#).
- The ability to restrict object types displayed for each DBMS platform. For details, see [Node, or Object Type, Filtering](#).

Two default filters applying to owned and system objects are available. For details, see [Using the Default Filters](#). Default and custom filters are available in the Explorer tree **Filters** node for each data source. For details, see [Filters Node](#).

### SIMPLE, ON-THE-FLY, CHARACTER-BASED TREE FILTERING

The Filter box at the top of the Explorer provides a quick, ad hoc way to restrict the nodes displayed.



### To restrict the Explorer tree display to only user-defined nodes whose name contains a particular character string

- Type one or more characters in the Filter box.


The display is updated to show only nodes whose name contains the typed character string and their parent nodes.

**NOTE:** The Clear button in the Filter box (X) deletes the contents of the Filter box and restores the unfiltered tree display.

#### OBJECT NAME/SCHEMA FILTERING

Object filters let you restrict Explorer tree display for a data source by name, owning schema, or name/schema combinations. These filters are defined at the DBMS platform level and enabled or disabled on a data source by data source level. Multiple object filters can be created for a DBMS platform, letting you enable combinations of filters for a single data source, ad hoc enabling and disabling of individual filters depending on your requirements, and so on.

Each filter consists of one or more ANDed conditions based on name or schema. For example you could create a two-condition filter that restricts the Explorer tree for a data source to display only objects belonging to the Schema **QA** and whose name starts with the string **Test**.



Each filter condition is a simple expression that tests object names for equality (**Equals**, **Not Equals**), inclusion in a specified list (**In**, **Not In**), or using pattern matching (**Like**, **Not Like**). Optionally, an object filter can also be used to restrict display of specified object types on a data source.



### To create an object filter that can be enabled for all data sources for a DBMS platform

- 1 Right-click in the left pane of the Datasource Explorer tree and select **Filter** from the context menu. The **Filters** dialog opens.
- 2 In the **View** area, select **By Filter**.
 

**NOTE:** The current instructions create a filter at the DBMS level. To create a filter for a single data source, select the **By Datasource** view, select a specific data source, and then proceed to [step 4](#).
- 3 Expand the icon corresponding to the DBMS for which you want to create a filter and ensure that the icon is selected.
- 4 Click **New** and type a name for the new filter.
- 5 On the **Predicates** tab, create one or more name-based or schema-based conditions, all of which must be satisfied when this filter is enabled. Objects that do not satisfy the conditions will be filtered from view. Use the following guidelines in constructing conditions:
  - The **Like** and **Not Like** operators take pattern-matching **Values**. Use any pattern matching characters and conventions supported by the DBMS for which you are defining the filter.
  - The **In** and **Not In** operators take a comma-delimited list such as **CO,CA,CU**, as a **Value**.
  - Click the associated **New** button to add new predicates.
- 6 Optionally, on the **Database Elements** tab, select the check boxes for all object types to which this filter is to be applied. The filter does not affect the display of any deselected object types.
- 7 Click **OK**.

By default, when you create an object filter for a DBMS platform, it is enabled for each registered data source of that DBMS type.

### To disable or enable an object filter for a data source

- 1 Right-click in the left pane of the Explorer tree and select **Filter** from the context menu. The **Filters** dialog opens.
- 2 In the **View** area, select **By Datasource**.
- 3 Expand the icon corresponding to the target DBMS and then expand the icon corresponding to the specific data source for which you want to for enable or disable a filter.
- 4 Enable an object filter by selecting the associated check box or disable the object filter by deselecting the check box.
- 5 Click **OK**.

**NOTE:** You can also enable and disable filters using controls found under the **Filters** node for each data source in the Explorer tree.

You can also use the **Filters** dialog to select and edit an object filter, rename or create a clone of an existing object filter, and delete filters.

#### NODE, OR OBJECT TYPE, FILTERING

Node filtering lets you hide specified Datasource Explorer tree nodes for a DBMS platform. For example, if you do not typically work with credential-related objects such as users or logins, you may wish to filter those out of the Datasource explorer tree.

#### To change the hidden/displayed status of object types for a DBMS platform

- 1 Right-click in empty space in the left pane of the Explorer tree and select **Filter** from the context menu. The **Filters** dialog opens.
- 2 Click the **Node Filters** tab.
- 3 Expand the icon corresponding to the DBMS for which you want to create a node filter, and continue to expand nodes until all nodes for which you want to change the hidden/displayed status, are visible.
- 4 Ensure that all nodes corresponding to objects types that are to be displayed, are selected. Similarly, ensure that all nodes corresponding to objects types that are to be hidden, are deselected.
- 5 Click **OK**.

Hidden/displayed status is retained for the DBMS until you explicitly open the Filters dialog and change the current settings.

#### USING THE DEFAULT FILTERS

The **Filters** node for each registered data source includes two default Explorer tree filters, **Ignore System Objects** and **Show Only My Objects**. These filters can be enabled and disabled for each data source but cannot be edited or deleted.

#### FILTERS NODE

The Filters node contains default Explorer tree filters and any custom filters you create for that data source or create across that DBMS type. Controls let you enable or disable the default and custom filters. Options for creating and maintaining custom filters are provided. For details, see [Filtering in the Explorer](#).

#### FAVORITES TAB

Rapid SQL provides the Favorites Tab for designating and accessing favorite scripts. On the Favorite Scripts Tab you can do the following with frequently used SQL scripts:

- View
- Navigate

- Save
- Recall
- Execute

**TIP:** Sample Favorite Scripts are installed for Microsoft SQL Server, Oracle, and Sybase Adaptive Server.

For more information, see [Favorite Properties Dialog Box](#).

#### FAVORITE PROPERTIES DIALOG BOX

The table below describes the options and functionality on the Favorite Properties dialog box:

Option	Description
Description	Lets you enter a description.
File Name	Lets you type or browse and locate a file.
Hot Key	Lets you enter a Hot Key. Hot Keys must begin with CTRL or be stand-alone function keys. Rapid SQL automatically places the CTRL command in front of any character.
File Type	Lets you select a file type.

#### SUBSYSTEM NODE

**NOTE:** The Subsystem node is available for IBM DB2 LUW for OS/390 and z/OS only.

The Subsystem node displays detailed information about the DB2 subsystem. Subsystem objects include:

- [Connections](#)
- [DDF](#)
- [Parameters](#)

#### CONNECTIONS

Connections displays the current connections to the subsystem.

**NOTE:** The information is read-only.

For more information, see [Subsystem Node](#).

## DDF

DDF (distributed data facility) displays the information regarding the status and configuration of the DDF, a set of DB2 LUW for OS/390 components through which DB2 LUW for OS/390 communicates with another RDBMS.

**NOTE:** The information is read-only.

For more information, see [Subsystem Node](#).

## PARAMETERS

Parameters displays the DB2 subsystem parameters.

**NOTE:** The information is read-only.

For more information, see [Subsystem Node](#).

## PROJECT TAB

The Project Tab provides a visual method for browsing, accessing, and manipulating your projects. The Project Tab lets you:

- Open, close, modify, and browse projects.
- Drag project items to the Rapid SQL workspace.
- Create new projects.

**NOTE:** The Project Tab is available after you create or open a project.

For more information on working with projects, see [Project Management](#).

## VC FILES TAB

The VC Files Tab displays open version control files. This tab displays files listed in the \*.xml file for the version control files. The tab displays a checkbox icon for files that are checked-out.

For more information, see [Version Control](#).

## RAPID SQL WINDOWS

The Rapid SQL interface includes several windows to help you develop your program. The windows include:

### *Topics*

- [Describe Window](#)

- [Output Window](#)
- [Browsers](#)
- [Workspaces](#)
- [Preview Dialog Boxes](#)

## DESCRIBE WINDOW

Rapid SQL offers a floating Describe window for procedures, tables, views, and functions (Oracle and IBM DB2 LUW for Open Systems only). In the Describe window, you can view columnar information (for tables and views) or input parameter information (for procedures and functions).

### Opening the Describe Window

Rapid SQL offers three ways to open the Describe window:

- 1 In an editor, right-click an object and then click Describe from Cursor.
- 2 On the *Explorer* Tab, select an object and then click Describe.
- 3 On the *Explorer* Tab or in an editor, select an object and then press CTRL+D.

### Using the Describe Window

In the Describe window:

- 1 Click the Name list and then click a name to view a list of types of objects in the database.
- 2 Click the Owner list and then click an owner to view a list of all owners of objects in the database.
- 3 Click the Type list and then click a type to view columnar information (for tables and views) or input parameter information (for functions and procedures).

## OUTPUT WINDOW

Rapid SQL incorporates a separate window to capture all messages returned by the server and to log entries about the progress of operations started by the application.

For more information, see [Configuring the Output Window](#)

### CONFIGURING THE OUTPUT WINDOW

Rapid SQL lets you display, hide, or dock the [Output Window](#) anywhere in the application.

### Displaying the Output Window

- 1 On the *View* menu, click *Output*.

OR

On the *Main* toolbar, click *Output*.

Rapid SQL displays the Output Window.

### Hiding the Output Window

- 1 On the *View* menu, click *Output*.

OR

On the *Main* toolbar, click *Output*.

OR

Right-click the *Output Window* and then click *Hide*.

Rapid SQL hides the Output Window.

### Docking the Output Window

Right-click the *Output Window* and then click *Docking View*.

Rapid SQL docks the Output Window to the bottom of the application frame.

### Undocking the Output Window

Right-click the *Output Window* and then click *Docking View*.

Rapid SQL displays the Output Window as a floating window in the application.

### MESSAGES IN THE OUTPUT WINDOW

The [Output Window](#) lets you save, print, copy, and clear server messages.

### Saving Server Messages

- 1 Right-click the *Output Window* and then click *Save*.

Rapid SQL opens the Save As dialog box.

- 2 Enter the location and name of the file in the *File Name* box.

**NOTE:** Rapid SQL defaults the file extension to .msg.

- 3 To save the file, click *OK*.

## Printing Server Messages

- 1 Right-click the *Output Window* and then click *Print*.

Rapid SQL opens the Print dialog box.

**NOTE:** Rapid SQL prompts you with information on the size of the print job before opening the Print dialog box.

- 2 Configure your print job.
- 3 Click *OK* to print the file.

## Copying Server Messages

- 1 Right-click the target *Server Messages* and then click *Copy*.

Rapid SQL copies the selected text to the Microsoft Windows Clipboard.

- 2 Paste the contents of the clipboard into target applications.

## Clearing Server Messages

- 1 Right-click the *Output Window* and then click *Clear*.

Rapid SQL clears your Server Messages.

## BROWSERS

Browsers are a flexible environment where you can examine, extract, and execute database objects and their dependencies. Browsers provide the means to view objects types across multiple database platforms and connections. You can simultaneously view and work with objects from InterBase/Firebird, Oracle, Microsoft SQL Server, Sybase Adaptive Server, and IBM DB2 LUW for Open Systems.

The benefit of using Browsers is the ability to see detailed information about specific object types. You can also print, search, copy, and sort the contents of a Browser window.

### Topics

- [Browser Toolbar](#)
- [Opening Browsers](#)
- [Browser Object Types](#)
- [Extracting DDL from Browsers](#)
- [Displaying Dependencies from Browsers](#)
- [Refreshing Browsers](#)

## BROWSER TOOLBAR

You can place the floating *Browser* toolbar anywhere on the Rapid SQL workspace.

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

## OPENING BROWSERS

Browsers let you view all types of database objects, including the SQL procedures used to build them. Browsers let you:

- Copy database objects
- Modify database objects
- Test database objects

The ability to browse dependencies is especially useful, particularly when modifying SQL code in procedures and triggers. For example, in a situation where a trigger enforces a rule that does not let you update a particular box, you can use the Browser to do the following:

- Browse the triggers to find the offending trigger.
- Extract the DDL for that trigger into one window.
- In another window, drop the offending trigger, make your update to the box, then execute the corrected trigger DDL to replace the trigger in the database.

## Opening a Browser Window

Rapid SQL offers two ways to open a Browser:

- 1 On the *Browse* menu, click the target object type.

OR





In the workspace, right-click, click *Browser*, then click the target object.

Rapid SQL opens a Browser.

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

## BROWSER OBJECT TYPES

Rapid SQL's Browsers read the appropriate object types for specific databases. A select statement is issued against the appropriate systems table based on the requested object to bring back a listing of the objects in the database. The table below contains a list of the Browsers available for each database platform:

Object Type	DB2 LUW	DB2 z/OS	InterBase/Firebird	SQL Server	Oracle	Sybase
Aliases						



Object Type	DB2 LUW	DB2 z/OS	InterBase/Firebird	SQL Server	Oracle	Sybase
Blob Filters			✓			
Check Constraints	✓	✓	✓	✓	✓	✓
Clusters					✓	
Databases	✓	✓		✓	✓	✓
Database Links					✓	
Defaults				✓		✓
Directories					✓	
Domains			✓			
Encrypted Keys			✓			
Event Monitors	✓					
Exceptions			✓			
Extended Procedures				✓		✓
External Functions			✓			
Foreign Keys	✓	✓	✓	✓	✓	✓
Functions	✓	✓			✓	
Generators			✓			
Groups				✓		✓
Indexes	✓	✓	✓	✓	✓	✓

Object Type	DB2 LUW	DB2 z/OS	InterBase/Firebird	SQL Server	Oracle	Sybase
Libraries					✓	
Packages	✓	✓			✓	
Plans		✓				
PL/SQL Code Profiling					✓	
Primary Keys	✓	✓	✓	✓	✓	✓
Procedures	✓	✓	✓	✓	✓	✓
Profiles					✓	
Roles			✓		✓	
Rules				✓		✓
Rollback Segments					✓	
Segments				✓		
Sequences					✓	
Shadows			✓			
Snapshots					✓	
Snapshot Logs					✓	
Materialized Views					✓	
Materialized View Logs					✓	
Synonyms					✓	

Object Type	DB2 LUW	DB2 z/OS	InterBase/Firebird	SQL Server	Oracle	Sybase
Tables	✓	✓	✓	✓	✓	✓
Tablespaces					✓	
Triggers	✓	✓	✓	✓	✓	✓
Types					✓	
Type Bodies					✓	
Unique Keys	✓	✓	✓	✓	✓	✓
User Datatypes	✓	✓		✓		✓
User Messages				✓		
Users				✓	✓	✓
Views	✓	✓	✓	✓	✓	✓

## WORKING WITH BROWSERS

Browsers offer a versatile method of browsing and managing the contents of your databases. To help you maintain and organize your databases, you can:

- Print the contents of a Browser
- Search the contents of a Browser
- Copy the contents of a Browser
- Sort the contents of a Browser

### Printing Browsers

- 1 Open a Browser for the desired object type.
- 2 On the *File* menu, click *Print* to open the *Print* dialog box.
- 3 In the *Name* box, click the list, which contains a list of local and network printers that you can access (if you do not see any listed, then your computer is not configured for any printers).

- 4 Click the target printer.
- 5 In the *Print Range* box, click the appropriate option button to indicate print range.
- 6 In the *Number of copies* text box of the *Copies* box, click the *Up* or *Down* arrow or enter the number of copies.
- 7 Click *OK*.

Rapid SQL prints the selection.

### Searching Browsers

- 1 Open a Browser for the desired object type.
  - 2 On the *Edit* menu, click *Find*.
- Rapid SQL opens the Find box.
- 3 In the *Find What* text box, enter the search string.
  - 4 To make the search case sensitive, select the *Match Case* check box.
  - 5 To specify the direction to search, in the *Direction* box, click the *Up* or *Down* option button.
  - 6 Click *Find Next*.

Rapid SQL finds the next occurrence of your search string.

### Copying Browsers

- 1 Open a Browser for the desired object type.
- 2 Select the objects to copy.
- 3 On the *Edit* menu, click *Copy*.
- 4 Place the pointer at the position where you want to paste the objects, and then on the *Edit* menu, click *Paste*.

### Sorting Browsers

- 1 Open or [Create](#) a Browser for the desired object type.
- 2 Double-click the column header for the column of data to sort and Rapid SQL lists the contents of the column in ascending order.
- 3 Double-click the column header again and Rapid SQL lists the contents of the column in descending order.

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

## EXTRACTING DDL FROM BROWSERS

For each database type, Rapid SQL provides an appropriate Browser. The Browsers are mutually exclusive object windows, showing only objects of a given type. If you connect to multiple datasources, you have access to a number of objects that are not available based on the database platform. Rapid SQL includes intelligence to determine the valid object types in the underlying datasource.

### Using the Main Menu

- 1 On the *Browse* menu, click the target object type.

Rapid SQL opens the appropriate Browser:

- 2 In the Browser, double-click the target object type to extract the object type DDL into a *DDL Editor*.

### Using the Browser Toolbar

- 1 On the *Browser* toolbar, click Tables.

Rapid SQL opens the Table Browser:

- 2 Click the scroll bar arrow to locate the target table.
- 3 Double-click the target table.

Rapid SQL extracts the schema DDL into a DDL Editor.

### Using the Shortcut Menu

- 1 Right-click an open area of the workspace, click Browsers, and then click the target object type.
- 2 In the Browser, double-click the target object type.

Rapid SQL extracts the schema DDL into a DDL Editor.

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

## DISPLAYING DEPENDENCIES FROM BROWSERS

You can display object dependencies for an object from its corresponding object Browser. Rapid SQL displays the dependencies in a separate result set window.

### Displaying Dependencies

Rapid SQL offers three ways to display dependencies from Browsers:

- 1 Open a Browser for an object type.
- 2 In the Browser, click the target object.

- 3 On the *Object* menu, click *Dependencies*.

OR

On the *Browser* toolbar, click *Dependencies*.

OR

Right-click the target object and then click *Dependencies*.

Rapid SQL displays dependencies in a separate window.

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

## REFRESHING BROWSERS

Rapid SQL lets you refresh and display the results of a Browser operation.

### Refreshing the Browser

- 1 On the *Object* menu, click *Refresh*.

OR

On the *Browser* toolbar, click *Refresh*.

OR

Right-click the Browser workspace and then click *Refresh*.

Rapid SQL refreshes the results of the browser operation.

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

## WORKSPACES

Workspaces are a convenient way to maximize your desktop. You can use workspaces to multiply the amount of scripting, script execution, and development resources you have available at any one time. Rapid SQL lets you open and use several workspaces at one time. Using more than one workspace lets you:

- Execute long running scripts in one workspace while working in other workspaces.
- Develop strategies for working on scripts and result sets in one workspace while other scripts reside in one or more of the other workspaces.

## **Toggling Between Workspaces**

Rapid SQL offers two ways to toggle between workspaces:

- 1 On the *Main* toolbar, click *Workspace*.

OR

Right-click the current workspace and then click the target workspace.

Rapid SQL brings the target workspace forward.

For more information, see:

[Managing Workspaces](#)

## **MANAGING WORKSPACES**

Rapid SQL provides you with three default workspaces. You manage the workspaces in the Workspace dialog box. Using the Workspace dialog box you can:

- Differentiate between workspaces by changing the background color or wallpaper.
- Toggle among workspaces.
- Create, delete, rename, and specify the order of workspaces.

## **Managing Workspaces**

The Workspace dialog box lets you manage all open windows in your workspace.

- 1 On the *Windows* menu, click *Windows*.

Rapid SQL displays the Workspace dialog box. Any open windows in the current workspace display in the list.

The table below describes the options and functionality on the Workspace dialog box:

Option	Description
Activate	Sets the focus onto the window you have selected in the list and closes the Workspace dialog box.
OK	Closes the Workspace dialog box and accepts any changes you have made to the windows in the current workspace.
Save	Saves the contents of the window you have selected in the list. You are prompted to provide a name and location for the file you are saving if you have not done so already.
Close Window	Closes the window you have selected from the list. If you have not saved the contents of the window, you are prompted with a save file alert.
Help	Initiates and displays this Help topic in the Rapid SQL Help.

## SET SORT COLUMNS DIALOG BOX

The Set Sort Columns dialog box lets you sort multiple columns, and sort column identification, in the Right Pane of the application.

For more information, see [Completing the Set Columns Dialog Box](#).

### COMPLETING THE SET COLUMNS DIALOG BOX

To complete the Set Columns dialog box, do the following:

- 1 In the right pane of the application, right-click a column heading and select Sort Multiple Columns.  
Rapid SQL opens the Set Sort Columns dialog box.
- 2 In Column Name select the column(s) to use to sort the information in the right pane of the application.
- 3 Click the right arrow to move the column(s) to the Column Name box.
- 4 Click the up and down arrows to change the order of the sort columns.

For more information, see [Set Sort Column Dialog Box](#).

## PREVIEW DIALOG BOXES

Before executing any code, Rapid SQL offers Preview dialog boxes to let you confirm actions before execution. In the Preview dialog boxes, you can:

- Preview the code to execute.
- View the SQL of the code on your database.
- Create a report detailing the affect of executing code on your database.
- Schedule execution of the code.
- Save the code to execute.
- Open your e-mail program with the code to execute as an attachment.
- Print the code to execute.

Topics:

- [Preview](#)
- [Impact Analysis](#)
- [Report](#)



## MENUS

Rapid SQL offers two context-sensitive menus to let you access all the application's features. The Main Menu is always on the top of the application window. The shortcut menu is accessible from almost anywhere in the application. Right-click to view the available shortcut menu. Rapid SQL lets you customize the Tools menu to help you tailor the application to your needs.

### Topics

- [Main Menu](#)
- [Shortcut menus](#)
- [Customizing General User Interface Appearance](#)

## MAIN MENU

Rapid SQL's features can all be accessed from the Main Menu by clicking the menu name and selecting from the submenu. The menus are context sensitive and change based on the tasks you want to perform. The table below describes the Rapid SQL menus:

Menu Item	Description
File	Create, open, close, print, send, and save script files and result sets. Set application options and defaults.
Datasource	Create, modify, select, connect to, and disconnect from datasources. Access the database search facility.
Project	Available only when a project is open. Configure project management, build projects, and use version control functions.
Browse	Browse any object type a datasource connection.
Logfile	Activate/deactivate, open, set options, and flush the Rapid SQL application log.
View	Arrange the Rapid SQL environment. Display or hide the Database Explorer, toolbars, Output Window, Describe window, activate full-screen mode.
Tools	Choose any of Rapid SQL's tools, such as Database Search and the Visual Diff Utility. Customize and add tools of your own.
Bookmarks	Access and manage bookmarks.
Help	Access HTML Help.
Query	Available only when an Editor is open. Execute and set options for your SQL scripts.
Object	Available only when a browser is open. Execute, view dependencies, extract, and refresh objects in a database.
Edit	Available only when an Editor is open. Edit and manipulate the text in your scripts.
Format	Available only when a Result Window is active. Format the contents of result sets.
Window	Cascade and tile open windows. Toggle among open windows.

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

## EXPLORER BOOKMARKS

The Bookmarks menu lets you access and manage explorer bookmarks. Explorer bookmarks let you quickly access nodes in the Database Explorer.

### Creating Explorer Bookmarks

- 1 On the Database Explorer, right-click the target node, and then select Add Bookmark.

Rapid SQL opens the Add Friendly Bookmark Name dialog box.

- 2 Type the explorer bookmark name.

- 3 Click OK.

Rapid SQL displays the explorer bookmark under the Bookmarks menu. Explorer bookmarks are organized by platform.

### Editing Explorer Bookmarks

- 1 On the Main Menu, select Bookmarks.

- 2 Select Bookmark Manager.

Rapid SQL opens Bookmark Manager.

- 3 To rename the explorer bookmark, select the target explorer bookmark, and then click Rename.

Rapid SQL opens the Edit Bookmark Name dialog box.

- 4 Type the new explorer bookmark name.

- 5 Click OK.

- 6 To delete an explorer bookmark, select the target explorer bookmark, and then click Delete.

**TIP:** To add explorer bookmarks without using the Add Friendly Bookmark Name dialog box, select Do not show 'Add Friendly Bookmark Name' dialog option.

## SHORTCUT MENUS

Rapid SQL incorporates context-sensitive menus to give you another way to access object functionality. These menus mirror the functionality that you can access from application toolbars or the main menu.

## Opening Shortcut Menus

- 1 Right-click anywhere on the Rapid SQL desktop to open the appropriate shortcut menu.

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

## TOOLBARS

Rapid SQL toolbars change to reflect the element of the application you are using. The toolbars contain icons that are the fastest way to access commonly used features of Rapid SQL. You can move the toolbars to horizontal or vertical positions anywhere on the screen, and you can toggle them off and on by using the shortcut menu when the pointer is positioned over one of Rapid SQL's toolbars. For more information, see the following topics:

- [Available Toolbars](#)
- [Using Toolbar Viewing Options](#)
- [Moving Toolbars](#)

**NOTE:** For related information, see [Customizing General User Interface Appearance](#).

## AVAILABLE TOOLBARS

The following list represents Rapid SQL's toolbars:

### *Datasource Toolbar*



### *Registration Toolbar*



### *Main Toolbar*



### *Windows Toolbar*



### *Tools Toolbar*



### *Project Toolbar*



### *SQL Edit Toolbar*



### **Browsers Toolbar**



## USING TOOLBAR VIEWING OPTIONS

Rapid SQL offers standard Windows toolbar options such as docking, floating, and positioning toolbars. The only application-specific viewing option is the hiding or display of the individual toolbars.

### **To hide or display a toolbar**

- 1 On the **View** menu, select **Toolbars** and then select the specific toolbar you want to display or hide.

For information on the toolbars available, see [Available Toolbars](#).

## MOVING TOOLBARS

- 1 Click the pointer at the right edge of any toolbar.
- 2 Drag the toolbar to the new position.

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

## CUSTOMIZING GENERAL USER INTERFACE APPEARANCE

Rapid SQL lets you choose from a set of general visual application styles, dictate hiding or display of particular items, and select preferences for specific user interface elements.

### To customize the general look and feel of Rapid SQL

- 1 On the **View** menu, select **Toolbars** and then select **Customize**. The **Customize** dialog opens.
- 2 Use the following table as a guide to understanding and setting options on tabs of the Customize dialog:

Tab	Settings and tasks	
<b>Toolbars</b>	Select the toolbars you want to display in the application. For information on the toolbars available, see <a href="#">Available Toolbars</a> .	
<b>Application Visual Style</b>	Select a visual style such as Microsoft Windows XP or one of the .NET options from the dropdown. Depending on your selection, the following options may or may not be enabled: <b>Use default Windows XP colors</b> , <b>OneNote style tabs</b> , <b>Docking Tab Colors</b> , <b>Allow MDI Tab Swapping</b> , <b>Enable Smart Docking</b> , <b>Enable Tab Menu</b> , and <b>3D Rounded Docking Tabs</b> .	
	<b>Menu animations</b>	Lets you specify a menu animation style of UNFOLD, SLIDE, or FADE.
	<b>Menu Shadows</b>	Displays shadowed 3D effects.
<b>Tools</b>	Lets you define external applications to run from the <b>Tools</b> menu of Rapid SQL: the text displayed on the Tools menu command ( <b>Menu contents</b> ), the path and file name of the executable ( <b>Command</b> ), optional <b>Arguments</b> , and an optional <b>Initial Directory</b> . For detailed information on providing arguments, see <a href="#">Specifying an Argument for a Tools Menu Command</a> .	
<b>Keyboard</b>	<b>Category</b>	Select a general category for a hot key for the command.
	<b>Commands</b>	Select a hot key command, based on the general category.
	<b>Description</b>	Displays the command description.
	<b>Set Accelerator for</b>	Select application area where you want new hot key to be active.
	<b>Current Keys</b>	Displays current hot key.
	<b>Press New Shortcut Key</b>	Press keyboard key or an unassigned F key.

Tab	Settings and tasks	
<b>Options</b>	<b>Show ScreenTips on toolbars</b>	Select to display a ScreenTip when you hover your mouse over a button. For example, when you hover your mouse over the New button, Rapid SQL displays the ScreenTip "New."
	<b>Show shortcut keys in ScreenTips</b>	Select to display a shortcut key in the ScreenTip when you hover your mouse over a button. For example, when you hover your mouse over the New button, Rapid SQL displays the ScreenTip "New (CTRL+N)."

- Click **Apply** at any time to implement any changes you have made and when finished, click **Close**.

## SPECIFYING AN ARGUMENT FOR A TOOLS MENU COMMAND

You can specify an argument to be passed to a program for newly added commands by choosing one of Rapid SQL's [predefined arguments](#) or entering a command-line argument.

The table below provides scenarios of how to use command-line arguments:

Command	Argument	Description
NOTEPAD.EXE	\$\$FilePath\$\$	Starts Microsoft Notepad displaying the contents of the \$\$FilePath\$\$ argument.
ISQL.EXE	-U\$\$CurUserID\$\$ -P\$\$CurPwd\$\$ -S\$\$CurConString\$\$ -i\$\$FilePath\$\$	Starts ISQL, automatically connects to the current datasource using the current user name and password, and executes the contents of \$\$FilePath\$\$.
SQLPLUS.EXE	\$\$CurUserID\$\$/ \$\$CurPwd\$\$@\$\$CurConString\$\$ @\$\$FilePath\$\$	Starts SQL*Plus, connects to the current datasource using the current user name and password, and executes the contents of \$\$FilePath\$\$.

The table below provides scenarios of how to use Rapid SQL's predefined arguments:

Argument	Description
\$\$FilePath\$\$	The complete filename of the current source (defined as drive+path+filename); blank if a non-source window is active.
\$\$FileDir\$\$	The directory of the current source (defined as drive+path); blank if a non-source window is active.
\$\$FileName\$\$	The filename of the current source (defined as filename); blank if the non-source window is active.
\$\$FileExt\$\$	The filename extension of the current source; blank if a non-source window is active.
\$\$CurLine\$\$	The current cursor line position within the active window.
\$\$CurCol\$\$	The current cursor column position within the active window.

Argument	Description
\$\$CurText\$\$	The current text (the word under the current cursor position, or the currently selected text, if there is one).
\$\$CurDir\$\$	The current working directory (defined as drive+path).
\$\$CurDatasource\$\$	The name of the current datasource as defined in Rapid SQL.
\$\$CurUserID\$\$	The name of the current datasource user.
\$\$CurPwd\$\$	The current datasource password.
\$\$CurConString\$\$	The current connection string or server name.

### PREDEFINED ARGUMENTS

Rapid SQL provides a number of predefined arguments that you can pass to programs that you have added to the Tools menu. The table below lists the available predefined arguments:

Argument	Description
\$\$FilePath\$\$	The complete filename of the current source (defined as drive+path+filename); blank if a non-source window is active.
\$\$FileDir\$\$	The directory of the current source (defined as drive+path); blank if a non-source window is active.
\$\$FileName\$\$	The filename of the current source (defined as filename); blank if the non-source window is active.
\$\$FileExt\$\$	The filename extension of the current source; blank if a non-source window is active.
\$\$CurLine\$\$	The current cursor line position within the active window.
\$\$CurCol\$\$	The current cursor column position within the active window.
\$\$CurText\$\$	The current text (the word under the current cursor position, or the currently selected text, if there is one).
\$\$CurDir\$\$	The current working directory (defined as drive+path).
\$\$CurDatasource\$\$	The name of the current datasource as defined in Rapid SQL.
\$\$CurUserID\$\$	The name of the current datasource user.
\$\$CurPwd\$\$	The current datasource password.
\$\$CurConString\$\$	The current connection string or server name.

**NOTE:** Arguments are case-sensitive.

### KEYBOARD SHORTCUTS

Rapid SQL provides a number of keyboard shortcuts to help you expedite your tasks. The table below lists the taxes and related shortcuts:

General Editing	Keyboard Command
Delete one character to the left	BACKSPACE

General Editing	Keyboard Command
Delete one character to the right	DELETE
Cut selected text to the Clipboard	CTRL+X
Undo the last action	CTRL+Z
Redo the last undo operation	CTRL+Y
Copy text	CTRL+C
Paste the Clipboard contents	CTRL+V

To Extend a Selection	Keyboard Command
One character to the right	SHIFT+RIGHT ARROW
One character to the left	SHIFT+LEFT ARROW
To the end of a word	CTRL+SHIFT+RIGHT ARROW
To the beginning of a word	CTRL+SHIFT+LEFT ARROW
To the end of a line	SHIFT+END
To the beginning of a line	SHIFT+HOME
One line down	SHIFT+DOWN ARROW
One screen up	SHIFT+PAGE UP
To the beginning of a document	CTRL+SHIFT+HOME
To the end of a document	CTRL+SHIFT+END
To include the entire document	CTRL+A

To Move the Insertion Point	Keyboard Command
One character to the left	LEFT ARROW
One character to the right	RIGHT ARROW
One word to the left	CTRL+LEFT ARROW
One word to the right	CTRL+RIGHT ARROW
One line up	UP ARROW
One line down	DOWN ARROW
To the end of a line	END
To the beginning of a line	HOME
One screen up (scrolling)	PAGE UP
One screen down (scrolling)	PAGE DOWN
To the end of a document	CTRL+END
To the beginning of a document	CTRL+HOME



Bookmarks	Keyboard Command
Toggle bookmark on/off	CTRL+F2
Go to next bookmark	F2
Go to previous bookmark	SHIFT+F2

Splitter Windows	Keyboard Command
Go to next pane	F6
Go to previous pane	SHIFT+F6

Debugger Operations	Keyboard Command
Start Debugging	CTRL+F5
Stop Debugging	SHIFT+F5
Step Over	F10
Step Into	F11
Run to Cursor	CTRL+F10
Step Out	SHIFT+F11
Describe from Cursor	CTRL+D
Insert or Remove Breakpoint	F9
Toggle (Enable or Disable) Breakpoint	CTRL+F9
Edit Breakpoint	ALT+F9
Go	F5
Restart	CTRL+SHIFT+F5

Debugger Windows	Keyboard Command
Open or Close Watch Window	ALT+3
Open or Close Variables Window	ALT+4
Open or Close Call Stack Window	ALT+5
Open or Close Dependency Tree Window	ALT+6

Other Windows	Keyboard Command
Go to the Result Tab	CTRL+ALT+R
Go to the Query Tab	CTRL+ALT+Q

Other Windows	Keyboard Command
Open the Describe window (for highlighted object)	CTRL+D
Toggle between Workspaces	CTRL+W
Toggle between Datasource Explorer and ISQL Window.	CTRL+ALT+E

## FULL SCREEN MODE

Rapid SQL has full screen mode capabilities so you can conceal the application framework and use the entire monitor area. Full screen mode hides any other applications running on the computer and uses every available pixel for the application. Main menu functionality is accessible through keyboard commands when you use full screen mode.

### Activating Full Screen Mode

Rapid SQL offers two ways to activate full screen mode:

- 1 On the *View* menu, click *Full Screen*.

OR

On the *Main* toolbar, click Full Screen.

Rapid SQL expands the application to fit the entire monitor area.

**NOTE:** The Full Screen mode icon is a stand-alone floating toolbar.

### Dismissing Full Screen Mode

- 1 Click *Full Screen* to expand the application to fit the entire monitor area.

**TIP:** If you closed the Full Screen mode toolbar, right-click the top of the Rapid SQL desktop to bring the toolbar back.

- 2 Click *Full Screen* to restore the application to the default size.

# SPECIFYING RAPID SQL APPLICATION AND FEATURE OPTIONS

Rapid SQL lets you customize the application configuration for your specific needs. All Rapid SQL application settings are available in the Options Editor, which is organized in a tabular format based on feature sets.

## To specify options for the Rapid SQL application or a particular feature

- 1 On the *File* menu, click *Options*.

Rapid SQL opens the Options Editor.

- 2 Click the tab corresponding to the feature you want to customize:

<a href="#">Browsers Options</a>	Lets you specify appearance of browser windows.
<a href="#">Code Analyst Options</a>	Lets you set dependency profiling level, allram and threshold, and oracle options for Code Analyst.
<a href="#">Connection Options</a>	Specifies the timeout parameters, packet size for a connection, and ANSI to OEM settings.
<a href="#">Data Editor Options</a>	Specifies settings for Data Editor.
<a href="#">Data Transfer Options</a>	Sets default directories when performing data unload and load and data export and import operations.
<a href="#">Explorer Options</a>	Sets defaults for the organization of objects in the Datasource Explorer.
<a href="#">Datasource Options</a>	Specifies general datasource management options and Datasource Explorer preferences.
<a href="#">DDL Extract Options</a>	Specifies whether or not Rapid SQL should include DROP statements when extracting the schema of different database object types.
<a href="#">Debug Options</a>	Sets the duration of your debug initialization and debug session, enable or disable DBMS_OUTPUT, and enable the refresh option.
<a href="#">Directories Options</a>	Sets the default directories for placing the output of different operations such as HTML reports or schema extractions.
<a href="#">General Options</a>	Sets defaults for automatic login, restoring the last session, and other general application options.
<a href="#">Grid Properties (Results window) Options</a>	Dictates the physical appearance of the results window grid.
<a href="#">ISQL Options</a>	Sets defaults for the maximum allowable errors before aborting the execution of a SQL script, executing selected text, and the position of Query and Results tabs.
<a href="#">Java Options</a>	Specifies load Java files and drop Java files.
<a href="#">JDBC Options</a>	Lets you specify Java Virtual Machine options in addition to the JDBC driver location for each platform.
<a href="#">Logging Options</a>	Sets defaults for SQL Logging.
<a href="#">MySQL Utilities Options</a>	Specifies paths to MySQL utilities.
<a href="#">Oracle Utilities Options</a>	Specifies the location of the Oracle Utilities.
<a href="#">Perf Center Options</a>	Specifies Performance Center's integration with Rapid SQL.
<a href="#">Query Builder Options</a>	Specifies global settings for Query Builder.

<a href="#">Results (ISQL) Options</a>	Specifies auto format result sets, sets the display and format of Result Windows, and the mail file type and default fonts.
<a href="#">SMTP Mail Options</a>	Lets you specify defaults for your outgoing mail notifications.
<a href="#">Version Control Options</a>	Lets you select which version control system you want Rapid SQL to use as the underlying version control system
<a href="#">Warnings Options</a>	Activates specific warnings when undesirable actions are attempted against a database.

3 Set feature options on the tab and then click OK.

**NOTE:** If there is an open document, Rapid SQL opens the **Update Document Statement Properties** dialog box. The Update Document Statement Properties dialog box lets you override changes you made to a current document or documents with new setting you made in the Options Editor.

## BROWSERS OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Browsers** tab.

The Browsers Tab of the Options Editor lets you set a number of options for your browser windows, including browser appearance, mail file type options, and default object owner. The table below describes Browser options:

Interface Element	Option	Description	Default
<b>Window</b>	<b>Show Toolbar</b>	Toggles the Browser window toolbar. Commands can also be activated from the shortcut menu that displays when you right-click in an open Browser window.	Selected
	<b>Show Status Bar</b>	Toggles the status bar at the bottom of the Browser window. The status bar displays the cell location of the current focus and the number of rows in the Browser window itself.	Selected
	<b>Detailed Listing</b>	Toggles detail columns for a given object type. Each object contains object specific details, such as creation date, segment, and so on. When Detailed Listing is turned off, only the names of objects and the numbers of rows of data for a table are displayed.	Not selected
<b>Appearance</b>	<b>Text Color</b>	Sets the text color for all Browsers of that object type. Select the row appropriate to a particular object type, then click a color in the Text column. Remember that text colors appear best against contrasting background colors, such as black on white.	Available
	<b>Background Color</b>	Sets the background color for all Browsers of that object type. Select the row appropriate to a particular object type, then click a color in the Background column. Remember that text colors appear best against contrasting background colors, such as black on white.	Available
<b>Browser File</b>	<b>Mail File Type</b>	Specifies what file type you want the browser file saved as when sent as e-mail to another user on a MAPI compliant mail system. The valid types are the proprietary Rapid SQL Results type, Tab delimited, Comma separated, and HTML, which formats the results in a simple HTML table.	Results

Interface Element	Option	Description	Default
<b>Default Owner</b>	<b>All Owners</b>	Indicates that the Browser Window should default to displaying database objects in the browsers for all owners in the database.	Not Selected
	<b>Current User</b>	Indicates that the Browser Window should default to displaying database objects owned exclusively by the current user.	Selected
	<b>Specific Owner</b>	Indicates that the Browser Window should default to displaying only the database objects owned by a specific owner. If you select this option, you must provide an owner name in the box.	Not Selected

## CODE ANALYST OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can set preferences for Code Analyst. For information on how to use Code Analyst to identify time-consuming lines of code, see [Code Analyst](#).

The table below describes Code Analyst options:

Group	Options and descriptions	
<b>Dependency profiling level</b>	<b>All Dependencies</b>	Code Analyst will use all levels of dependencies when profiling the session.
	<b>Dependencies up to level</b>	Lets you specify how many levels of dependencies the Code Analyst will use when profiling the session.
<b>Oracle</b>	<b>List Package objects</b>	Code Analyst lists the procedures within a package and functions that are found in the Oracle database.
	<b>Use SQL Profiler</b>	Code Analyst uses Oracle's DBMS_Profiler package to collect time metrics.  Code Analyst displays the actual run time on the database, and does not include the time it takes to get to the server.
<b>Alarms and Thresholds</b>	<p>Lets you set an alarm for each type of object that is collected. For each supported <b>Object Type</b> (procedure, function, package, trigger, type), you can provide the following settings:</p> <p><b>Alarm %</b> - lets you specify the percentage of total run time a line of code takes when an alarm appears. If the object took more than specified percent of total time, Code Analyst alerts the user by changing the color of the text</p> <p><b>Ignore 100%</b> - lets you ignore lines of code that take all of the total run time</p>	
<b>Show Code Analyst Confirmation Dialog</b>	When you create or execute a session, Code Analyst displays a message that the Code Analyst will run longer than the actual code. You can also select the "Please do not show me this dialog again" option in the dialog box.	

## CONNECTION OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can configure a number of server parameters for interacting with your datasources on the Connection tab.

**NOTE:** Setting Connection Options are available for Microsoft SQL Server and Sybase ASE.

The table below describes the options and functionality on the Connection Tab:

Option	Description	Default
<b>Login Timeout</b>	Specifies the number of seconds that the application should wait for a response to a connection request from server. If server does not respond within the specified period, the application aborts the connection and returns an error message.	30
<b>Query Timeout</b>	Specifies the number of seconds that the application should wait for a response to a query from the server. If the server does not respond within the specified period, the application terminates its query process and returns an error.	0
<b>Packet Size</b>	Specifies the network packet size to be used when communicating with the server, in bytes.	512
<b>Max Connections (ctlib)</b>	The maximum number of connections allowed by the ctlib setting on the client.	0
<b>Client Character Set</b>	Character set of client computer.	Local character set
<b>Host Name</b>	Name of the client computer.	Local name
<b>Use Quoted Identifiers</b>	If you plan to use delimited identifiers, select this option.	Not selected

## DATA EDITOR OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Data Editor** tab.

**NOTE:** Option parameters set in the Options Editor override options set within Query Builder and Data Editor.

The table below describes the options and functionality on the **Data Editor** Tab:

Interface Element	Option	Description	Default
<b>Default Execution Mode</b>	<b>Live Mode</b>	lets you execute changes one row at a time. <b>Batch Mode</b> lets you make unlimited changes prior to execution. Batch Mode offers three sub-options: <b>Ignore errors-continue processing</b> , <b>Prompt on Error</b> , and <b>Stop after error(s)</b> (number of errors allowed before stopping execution)	Live Mode
<b>Data Editor File</b>	<b>Mail File Type</b>	Sets the default mail output style as Results, Tab Delimited, Comma Separated, or HTML.	Results

Interface Element	Option	Description	Default
	<b>Include column titles when saving</b>	Includes column titles when saving.	Not selected
<b>Grid Font</b>	Customizes font style, and size for the Data Editor and the Results Grid.		Available
<b>Printer Font</b>	Sets font style, and size for printing output.		Available
<b>Auto Format (Best Fit)</b>	Fits formatting to match your desktop.		Selected
<b>Begin and End Transaction Statements</b>	Adds a beginning and ending transaction on each statement.		Selected
<b>Default Date/Time Format</b>	Displays the current date/time format and lets you customize the date/time display.		Results
	<b>Use Calendar Control as default</b>	If selected, Rapid SQL uses the Calendar Control window.	Not selected
	<b>2 digit year system setting warning</b>	If selected, Rapid SQL sends a warning when you use a two digit year system setting.	Selected
<b>Confirmation Dialog Options</b>	Enabling <b>Show Delete Confirmation Dialog</b> , <b>Show Update Confirmation Dialog</b> , or <b>Show Update LOB Confirmation Dialog</b> lets you display a confirmation dialog box when you use a delete command, update a row, or update a LOB, respectively..		Selected

## DATA TRANSFER OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Data Transfer** tab.

You can configure Rapid SQL to use default directories when performing data unload and load and data export and import operations. Setting a default directory saves time because it describes a single reference point for loading or unloading, exporting and importing table or view data files.

The table below describes the options and functionality on the **Data Transfer** Tab:

Option	Description
<b>Data Unload</b>	Specifies the name and location of the default directory in the Data Unload box.
<b>Data Load</b>	Specifies the name and location of the default directory in the Data Load box.
<b>Oracle Export</b> (for Oracle)	Specifies the name and location of the default directory in the Oracle Export box.
<b>Oracle Import</b> (for Oracle)	Specifies the name and location of the default directory in the Oracle Import box.

## DATASOURCE OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Datasource** tab. The table below describes the options and functionality on the **Datasource** Tab:

Option	Description	Default
<b>Datasource Storage Management</b>	<p>Rapid SQL offers both registry-based and file-based datasource catalog management. Prior to changing methods, see <a href="#">Understanding the Datasource Catalog</a>.</p> <p>The following options are offered:</p> <p><b>Windows Registry datasource catalog</b> - The datasource catalog is stored and managed locally, in the Windows registry. This user can add and delete datasources and edit the registration details of existing resources. Any changes are retained across startups of Rapid SQL.</p> <p><b>File based datasource catalog</b> - The datasource catalog is stored locally and is file-based (%APPDATA%\Embarcadero\Data Sources\ in a default installation). This user can add and delete datasources and edit the registration details of existing resources. Any changes are retained across startups of Rapid SQL.</p> <p><b>Network shared datasource catalog</b> - The datasource catalog is file-based, and obtained from the location specified in the <b>Location of shared datasources file</b> control. Changes such as new datasources, deleted datasources or edited datasource registrations are lost as soon as this user shuts down Rapid SQL.</p> <p>Changing methods does not delete datasource definitions stored using the current method. When you change methods, Rapid SQL does not automatically convert existing catalog definitions to the new method</p>	Windows Registry
<b>Recent Datasource List Contains</b>	Lets you specify number of datasources to display when you select <b>Recent Datasources</b> from the <b>File</b> menu.	8
<b>Default to Alias Usage When Defining New Datasources</b>	For DBMS that support alias usage, defaults the <b>Connection Information</b> tab to alias usage selection.	Unselected
Datasource Category Settings	<p>Lets you select the optional user interface elements that are to be color-coded or labelled for datasources that have been assigned a category:</p> <p><b>Color the status bar using the datasource's category color</b></p> <p><b>Color the tabs using the datasource's category color</b></p> <p><b>Include category short name in the datasource toolbar combo</b></p> <p>For an introduction to datasource categorization and related tasks, see <a href="#">Categorizing datasources using color-coding and labelling</a>.</p>	N/A

## DDL EXTRACT OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **DDL Extract** tab. The DDL Extraction Tab of the Options Editor lets you specify common and DBMS-specific DDL Extract options.



## Common options

The table below describes the options and functionality on the **DDL Extract** Tab, main view, of the Options Editor:

Option	Description
<b>Extract to multiple windows</b>	Select to extract the schema for each object into separate DDL windows. NOTE: This option only works when you extract DDL for multiple objects.
<b>Extract in dependency order</b>	This is the default. When you select a number of objects as part of an extraction job, this option ensures objects will be extracted in the proper dependency order. If the option is not selected, errors may result when you execute the script. It's also true, however, that loading the dependencies can add significant overhead when you are extracting numerous objects.
<b>Script Use Statement</b>	This option is for Sybase. Optimizes extraction through generating Use statements in the generated script.
<b>Grouping Extracted DDL</b>	There are two option as to how the DDL is generated for an extract operation: <b>Group the Drop and Create DDL for each object together</b> - Objects are extracted one at a time. In the DDL generated by the extract, the Drop statement for a selected or referenced object is followed by the Create statement for that object. <b>Group DDL statements by type - all Drop statements, then all Create statements</b> - In the DDL generated by the extract, Drop statements for all selected or referenced objects are grouped to precede Create statements for those objects. This method ensures that all dependencies are respected. For more information on DROP options and dependencies, see <a href="#">DBMS-specific DDL Extract Options</a> .

## DBMS-specific DDL Extract Options

The DBMS-specific views of the DDL Extract tab let you specify:

- The object types to include DROP statements for
- The default dependent object types for each object type included in extraction/migration operations

You can choose to include DROP statements before you perform ad hoc DDL extractions. You can use this feature to modify and to re-compile database objects. To recompile a database object, drop it before recreating it. This option drops any existing objects of the same name before recreating the object. The data in the existing table is not saved when you specify a DROP statement for extracted DDL.

**CAUTION:** Because dropping an object is a destructive action, you should carefully consider including drop statements before activating this option.

You can also specify dependent object types included in extraction/migration operations. Each DBMS-specific view of the DDL Extract tab lets you specify the object dependencies for each object type. Dependent object types selected on that view are by default selected when you run an extraction/migration operation and can be overridden using those wizards. For more information see [Extract](#).

## DEBUG OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Debug** tab. The Debug Tab of the Options Editor lets you set the duration of your debug initialization and debug session, enable or disable DBMS\_OUTPUT, and enable the refresh option.

The table below describes the options and functionality on the **Debug** Tab:

Tab	Option	Description	Default
<b>General</b>	<b>Dependency Tree Option</b>	Lets you specify pre-fetch levels.	Pre-Fetch All Dependencies
<b>Profiler</b>	<b>Profiler Time Unit</b>	Lets you select a unit of milliseconds, seconds or minutes.	Milliseconds
	<b>Save Profiler Reports</b>	Lets you save profiler reports and type or browse for the report path.	Not Selected
<b>Oracle</b>	<b>Initialization Timeout (seconds)</b>	Specifies the point at which the application stops trying to initialize the debugger. If it cannot initialize the debugger in the specified time, it displays message in the Debug Output window.	60
	<b>Debug Session Timeout (seconds)</b>	Specifies the point at which the application terminates your debug session due to idle time.	7200
	<b>Enable DBMS Output</b>	Enables the Oracle built-in package, DBMS_OUTPUT, letting you send messages from stored procedures, packages, and triggers.	Selected
	<b>Refresh Dependencies for each run</b>	Refreshes the dependencies each time you run the debugger.	Not selected
	<b>Compile with Debug Option</b>	When debugging a procedure against an Oracle datasource, Rapid SQL performs a compilation before debugging. During a compilation, Oracle invalidates referencing objects. For example, if procedures A2 and A3 both reference procedure A1, debugging procedure A1 in Rapid SQL, will result in procedures A2 and A3 being marked as invalidated. The <b>Compile with Debug Options</b> settings let you control compilation of dependent objects while debugging.	Compile dependent options
<b>DB2</b>	<b>Debug Session Timeout (seconds)</b>	Specifies the point at which the application terminates your debug session due to idle time.	300
	<b>Compile with Debug Option before Debug Session</b>	Lets you specify options.	Prompt Always

## DIRECTORIES OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Directory** tab. Rapid SQL lets you configure Directory Options to use default directories when performing certain operations. You can set the default directories for:

- Wizard operations

- Report generation
- Schema extraction
- HTML templates for customizing reports

The table below describes the options and functionality on the **Directories** Tab:

Option	Description	Default
Wizard Definitions	Specifies the name and location of the default directory for saving wizard operations. This option applies to wizard operations in which you have the option to save the definition file, such as data unload/load or schema migration operations.	C:\Documents and Settings\user name\Application Data\Embarcadero\RapidSQL\DefFiles
HTML Reports	Specifies the name and location of the default directory for the output from generated HTML reports.	C:\Documents and Settings\user name\Application Data\Embarcadero\RapidSQL\Report
Schema Extraction	Specifies the name and location of the default directory for placing the output from schema extraction operations.	C:\Documents and Settings\user name\Application Data\Embarcadero\RapidSQL\Extract
HTML Template	Specifies the name and location of the default directory where DBArtisan can find the HTML template on which to base HTML reports. This feature lets you customize your HTML reports.	C:\Program Files\Embarcadero\RSQL770\HTML\Tpl
User SQL Scripts	Specifies the name and location of the default directory for SQL Scripts.	C:\Documents and Settings\user name\Application Data\Embarcadero\RapidSQL\UserSQLScripts
Job Config Files	Specifies the location of any ETSQLX job configuration files you have set on your local machine.	C:\Documents and Settings\davidt\My Documents\Embarcadero\RapidSQL\Directories\ETSQLXJobCfg
#include Files	Specifies the name and location of the directory searched for files specified by a #include directive in the ISQL editor, Procedure Object Editor, or Package Body Object Editor if there are no paths specified on the Datasource Properties tab of the Datasource Registration Editor. For more information, see <a href="#">Registering or editing datasources</a> .	C:\Documents and Settings\davidt\My Documents\Embarcadero\RapidSQL\Directories\IncludeFiles

## EXPLORER OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Explorer** tab. The Explorer Tab of the Options Editor lets you configure how objects are organized in the Datasource Explorer.

The table below describes the options and functionality on the **Explorer** Tab:

Option	Description	Default
Explorer organization	Lets you select a default grouping option of <b>Organize by Object Owner</b> (groups objects, by object type, for each user) or <b>Organize by Object Type</b> (groups objects by object type for all users in the same list). Grouping by object owner is most efficient if you are working with databases containing a high number of objects.  You can change the organization manually during a Rapid SQL session. For details, see <a href="#">Explorer Tab</a> .	Owner
<b>Show System Objects</b> and <b>Show Only My Objects</b> (available if you are organizing the Datasource Explorer by object type)	Lets you select a default Datasource Explorer tree display options for system objects and objects that you own. Available options in the two groups are:  <b>Never</b> - the relevant objects are not displayed by default. <b>Always</b> - the relevant objects are always displayed by default. <b>Use datasource filter</b> - The relevant default datasource filter is enabled. For more information, see <a href="#">Filters Node</a> .  You can change the <b>Show System Objects</b> and <b>Show Only My Objects</b> settings manually during a Rapid SQL session.. For details, see <a href="#">Using the Default Filters</a> .	N/A
<b>Options</b>	Lets you enable or disable the following settings:  <b>Refresh after Object Commands</b> - Refreshes the Explorer automatically after an object has been modified or created.  <b>Retain View Setting on Close</b> - Select to retain the current state of the Datasource Explorer so that it opens the same way the next time you start Rapid SQL.	N/A

## GENERAL OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **General** tab. The General tab of the Options Editor lets you specify general application options. The table below describes the options and functionality on the **General** Tab:

Tab	Option	Description	Default
<b>Main tab</b>	<b>Confirm on Exit</b>	Instructs Rapid SQL to issue a message confirming the operation before exiting the application.	Selected
	<b>Max Editors in a Single Open Operation</b>	Specifies the maximum number of editors allowable from a single Open operation.	5
	<b>Max Entries in Output Window</b>	Specifies the maximum number of messages that can appear in the Output Window before the contents are flushed. This option conserves memory resources. You can clear the Output window or raise the maximum number allowed at any time.	1500

Tab	Option	Description	Default
<b>Oracle</b>	<b>Data Dictionary View Usage</b>	Lets you synchronize Rapid SQL use of the Oracle Data Dictionary DBA and ALL/USER views with the privileges associated with the role used in creating and using the datasource.  <b>Role-based Views</b> - Rapid SQL will use views based on the user's role. That is, users with DBA privileges will use DBA views. <b>DBA Views</b> - Rapid SQL always uses DBA views. <b>All Views</b> - Rapid SQL always uses ALL views. <b>Use Datasource Configuration for Views</b> - lets you set up Data Dictionary usage on a view-by-view basis when editing or registering a datasource. For more information, see <a href="#">Registering or editing datasources</a> .	
	<b>Preserve Case in Object Identifiers</b>	Preserves case of the database object.	Not selected
<b>Sybase and SQL Server</b>	<b>Rename action style</b>	Lets you specify whether user-defined objects or datatypes are renamed using an extended ALTER or using the sp_rename procedure.	extended alter
<b>InterBase/Firebird</b>	<b>Preserve Case in Object Identifiers</b>	Preserves case of the database object.	Not selected

## GRID PROPERTIES (RESULTS WINDOW) OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the Grid Properties tab.

The Grid Properties tab of the Option editor lets you set preferences for the layout and appearance of the grid in an Results window. The table below describes the options and functionality on the **Grid Properties** tab:

Interface Element	Options and descriptions		Default
<b>Titles and Gridlines</b>	<b>3D-Buttons</b>	Enables or disables a 3-D appearance on row and column headings in the grid.	Set
	<b>Horizontal Gridlines</b>	Enables or disables ruling between rows of the grid.	Set
	<b>Vertical Gridlines</b>	Enables or disables ruling between columns of the grid.	Set
	<b>Mark Headings</b>	Enables or disables highlighted row and column headings.	Set
<b>Preview</b>	Displays a preview of the settings currently selected in the Titles and Gridlines group.		

## ISQL OPTIONS

The ISQL options are available on the following tabs:

- [ISQL Options - Main Tab](#)
- [ISQL Options - Oracle Tab](#)

- [ISQL Options - DB2 Tab](#)
- [ISQL Options - Sybase Tab](#)
- [ISQL Options - SQL Server Tab](#)
- [ISQL Options - MySQL Tab](#)
- [ISQL Options - DB2 OS390 Tab](#)
- [ISQL Options - InterBase Tab](#)
- [ISQL Options - Editor Tab](#)
- [ISQL Options - Auto format Tab](#)
- [ISQL Options - Code Assist Tab](#)

## ISQL OPTIONS - MAIN TAB

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the main tab of the **ISQL** pane. Use the following table as a guide to modifying settings on this tab:

<b>Max Errors Before Aborting Execution</b>	Sets the maximum number of errors allowed before aborting the execution of a script. A zero indicates that the option is inactive and that you recognize no limit to the number of errors allowed. Rapid SQL refers to this value when step executing SQL scripts.
<b>Execute Selected Text</b>	Executes a portion of a selected SQL script.
<b>Check Syntax When Executing</b>	For DB2, required to execute DB2 call statements in the ISQL Window.
<b>Automatically lock connection</b>	When disabled, a prompt to commit or rollback the transaction is displayed when you close the ISQL editor window. Enabling this options disables the prompts and locks the connection automatically.
<b>Prompt to lock database connection</b>	Locks the database connection on execution.
<b>Tabs</b>	Sets the appearance of the ISQL Window tabs to either the top or bottom of the ISQL Window.
<b>File Association</b>	Specifies whether the application should open an unknown file type automatically into an ISQL Window or prompt you with a message that Rapid SQL does not recognize the file type.
<b>Query Plan Layout</b>	Sets the default orientation of a graphical query plan. For more information, see <a href="#">Query Options</a> .

## ISQL OPTIONS - ORACLE TAB

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Oracle** tab of the **ISQL** pane. Use the following table as a guide to modifying settings on this tab:

Option	Description
<b>Enable DBMS Output</b>	Lets you specify <b>Buffer</b> size. 0 is the default.
<b>Auto-Commit changes</b>	Applies auto commit status changes to all open windows.
<b>View xmltype as clob</b>	When enabled, xmltype columns are displayed as CLOBs in the Results grid and the Data editor. Without this option selected, SELECT statements that qualify xmltype columns produce an <b>OciTypeBinder</b> conversion error. With the option selected, the SELECT submitted by Rapid SQL is modified to include a getclobval() method call.
<b>Default Query Plan</b>	Lets you select the default display, tree-based or graphical, when you generate a query plan. For more information, see <a href="#">Viewing a graphical Query Plan (Oracle)</a> .
<b>Load Query Options</b>	<p>If this option is not enabled, Rapid SQL loads a set of default query options that customize your execution environment, and periodically sends those settings to the server. Enabling this option and specifying an XML file, forces Rapid SQL to load a previously saved query options file each time an ISQL windows opens against this DBMS platform. This lets you override Rapid SQL's default query options.</p> <p>You can also enable this option when saving a query options file. For information on other load and save options, specific query options offered, manually updating query options, and conditions for which query options are sent to the server, see <a href="#">Query Options</a>.</p>

## ISQL OPTIONS - DB2 TAB

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **DB2** tab of the **ISQL** pane. Use the following table as a guide to modifying settings on this tab:

Option	Description
<b>Set Isolation Level</b>	Sets the default for the value of the <b>Isolation Level</b> option on the Query Options dialog. For details, see <a href="#">Query Options</a> .
<b>Auto-Commit changes</b>	Applies auto commit status changes to all open windows.
<b>Create Explain plan tables if required</b>	If set to TRUE, Explain Plan tables are created, as necessary. If set to FALSE and you don't manually create tables, Explain Plan operations will fail.
<b>Create explain plan tables on the SYSTOOLS schema</b>	If set to TRUE, Explain Plan tables are created on the SYSTOOLS schema. If the tables already exist in the user's default schema, Rapid SQL continues to use those tables. Refer to DB2 documentation for a listing of Explain Plan tables that must be deleted in order to use the SYSTOOLS option. If set to FALSE, Explain Plan tables are created under the user's default schema.

Option	Description
<b>Load Query Options</b>	<p>If this option is not enabled, Rapid SQL loads a set of default query options that customize your execution environment, and periodically sends those settings to the server. Enabling this option and specifying an XML file, forces Rapid SQL to load a previously saved query options file each time an ISQL windows opens against this DBMS platform. This lets you override Rapid SQL's default query options.</p> <p>You can also enable this option when saving a query options file. For information on other load and save options, specific query options offered, manually updating query options, and conditions for which query options are sent to the server, see <a href="#">Query Options</a>.</p>

## ISQL OPTIONS - SYBASE TAB

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Sybase** tab of the **ISQL** pane. Use the following table as a guide to modifying settings on this tab:

Option	Description
<b>Auto-Commit changes</b>	Applies auto commit status changes to all open windows.
<b>Load Query Options</b>	<p>If this option is not enabled, Rapid SQL loads a set of default query options that customize your execution environment, and periodically sends those settings to the server. Enabling this option and specifying an XML file, forces Rapid SQL to load a previously saved query options file each time an ISQL windows opens against this DBMS platform. This lets you override Rapid SQL's default query options.</p> <p>You can also enable this option when saving a query options file. For information on other load and save options, specific query options offered, manually updating query options, and conditions for which query options are sent to the server, see <a href="#">Query Options</a>.</p>

## ISQL OPTIONS - SQL SERVER TAB

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **SQL Server** tab of the **ISQL** pane. Use the following table as a guide to modifying settings on this tab:

Option	Description
<b>Enable Set Query Options</b>	Sets the default value of the <b>Send Set Options</b> setting on the Query Options dialog. For details, see <a href="#">Query Options</a> . Changing this value does not affect ISQL sessions currently open.
<b>Auto-Commit changes</b>	Applies auto commit status changes to all open windows.



Option	Description
<b>Load Query Options</b>	<p>If this option is not enabled, Rapid SQL loads a set of default query options that customize your execution environment, and periodically sends those settings to the server. Enabling this option and specifying an XML file, forces Rapid SQL to load a previously saved query options file each time an ISQL windows opens against this DBMS platform. This lets you override Rapid SQL's default query options.</p> <p>You can also enable this option when saving a query options file. For information on other load and save options, specific query options offered, manually updating query options, and conditions for which query options are sent to the server, see <a href="#">Query Options</a>.</p>

## ISQL OPTIONS - MYSQL TAB

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **MySQL** tab of the **ISQL** pane. Use the following table as a guide to modifying settings on this tab:

Option	Description
<b>Load Query Options</b>	<p>If this option is not enabled, Rapid SQL loads a set of default query options that customize your execution environment, and periodically sends those settings to the server. Enabling this option and specifying an XML file, forces Rapid SQL to load a previously saved query options file each time an ISQL windows opens against this DBMS platform. This lets you override Rapid SQL's default query options.</p> <p>You can also enable this option when saving a query options file. For information on other load and save options, specific query options offered, manually updating query options, and conditions for which query options are sent to the server, see <a href="#">Query Options</a>.</p>

## ISQL OPTIONS - DB2 OS390 TAB

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **DB2 OS390** tab of the **ISQL** pane. Use the following table as a guide to modifying settings on this tab:

Option	Description
<b>Load Query Options</b>	<p>If this option is not enabled, Rapid SQL loads a set of default query options that customize your execution environment, and periodically sends those settings to the server. Enabling this option and specifying an XML file, forces Rapid SQL to load a previously saved query options file each time an ISQL windows opens against this DBMS platform. This lets you override Rapid SQL's default query options.</p> <p>You can also enable this option when saving a query options file. For information on other load and save options, specific query options offered, manually updating query options, and conditions for which query options are sent to the server, see <a href="#">Query Options</a>.</p>

## ISQL OPTIONS - INTERBASE TAB

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Interbase** tab of the **ISQL** pane. Use the following table as a guide to modifying settings on this tab:

Option	Description
<b>Load Query Options</b>	<p>If this option is not enabled, Rapid SQL loads a set of default query options that customize your execution environment, and periodically sends those settings to the server. Enabling this option and specifying an XML file, forces Rapid SQL to load a previously saved query options file each time an ISQL windows opens against this DBMS platform. This lets you override Rapid SQL's default query options.</p> <p>You can also enable this option when saving a query options file. For information on other load and save options, specific query options offered, manually updating query options, and conditions for which query options are sent to the server, see <a href="#">Query Options</a>.</p>

## ISQL OPTIONS - EDITOR TAB

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Editor** tab of the **ISQL** pane. The table below describes the options and functionality on the **Editor** Tab:

Group	Option	Description	Default
<b>Window</b>	<b>Show Toolbar and Show Status Bar</b>	Enables/disables these ISQL Window user interface elements.	Selected
	<b>Maximize on new or open</b>	Indicates that Rapid SQL should maximize the ISQL Editor. If you already have an active MDI Window that is maximized, the default behavior is to maximize a new child window. To deactivate this option, ensure you do not have any active MDI Windows, such as the Explorer.	Selected
	<b>File Tracking</b>	Indicates that the ISQL Editor should use the File Tracking Facility to monitor the status of a file. If a file has been modified and saved outside the application, the application loads the most current version of the file into the IISQL Editor based on the options set for Auto-Reload File (see above.)	Selected
	<b>Auto-Reload File</b>	If File Tracking is enabled, indicates that the application should automatically reload a file that has been externally modified without prompting you. If you turn this option off, Rapid SQL prompts you before reloading your file if external changes have been saved.	Not selected
	<b>Auto-Save File</b>	<p>Indicates that files in the ISQL Editor should automatically be saved at the indicated time interval. Backup files are saved to <b>C:\Documents and Settings\username\Application Data\Embarcadero\RApidSQL\Backup\</b>.</p> <p>Backup files use a naming convention of the form <b>~ETxxxx.tmp</b>., See the registry under <b>HKEY_CURRENT_USER\Software\Embarcadero\RApidSQL\version\Backup</b> for a listing of backup file names and the path of the file to which they correspond.</p>	Selected (5 minutes)

Group	Option	Description	Default
<b>Command History</b>	<b>Save File Before Overwriting</b>	Specifies the action you want the application to take when selecting a command from the Command History box. You have the option to be reminded to save a file before overwriting (Ask First), to automatically save a file before overwriting (Always), or to automatically overwrite the file with the command (Never).	Ask First
	<b>Save Most Recent</b>	Specifies the number of commands you want to save in the Command History list in the top of the ISQL Window toolbar. The maximum value is 99.	15
<b>Printing</b>	Lets you select common printer options.		
<b>Appearance</b>	<b>Enable Syntax Highlighting</b>	Sets syntax highlighting on so that all keywords and comments are colored for easier reading and debugging.	Selected
	<b>Show Line Numbers</b>	Places line numbers in the left column of an ISQL Window.	Selected
	<b>Enable Outlining</b>	Enables and disables outlining.	Selected
	<b>Enable Text Wrapping</b>	Enables and disables a typical text wrap feature.	Selected
	<b>Editor Font</b>	Sets the font face, style, and size displayed in the editor.	Available
	<b>Syntax Colors</b>	Sets syntax coloring for keywords, comments, quotes, and default text for various file types and scripts from the Syntax Coloring dialog box.	Available
<b>Formatting</b>	<b>Auto Indent</b>	Sets automatic indentation for each carriage return, new line in your SQL script.	Selected
	<b>Expand Tabs and Tab Size</b>	Sets tabs as the specified number of spaces in result sets.	Selected (4)
<b>Clipboard Line Endings</b>	<b>None</b>	When copying text from the editor, no substitution is made for end-of-line characters. This provides the fastest cut and paste operation.	Selected
	<b>Line Endings</b>	When copying text from the editor, end-of-line characters are replaced with platform-specific equivalents (Windows - CR LF, UNIX - LF, Macintosh - CR). This lets you copy text into environments such as Notepad, that require platform-specific end-of-line characters.  This option avoids unnecessary characters and incorrect display, but cut-and-paste operations take more time to complete.	

## ISQL OPTIONS - AUTO FORMAT TAB

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Auto Format** tab. of the **ISQL** pane.

The **Auto Format** tab lets you specify the style and spacing of SQL statements in an ISQL window when you choose to auto format SQL in the ISQL editor.

On opening, the Preview area shows an SQL statement formatted according to the current Auto Format settings. Click **Edit** to open the **Auto Format Options** dialog. Use the following table as a guide to understanding and modifying the settings in this dialog:

Setting	Description
<b>Keywords</b>	Lets you select a character case (UPPERCASE, LOWERCASE, INITIALCAPS, or NOCHANGE) formatting treatment for SQL keywords.
<b>Right Margin</b>	Specifies the maximum number of characters per line.
<b>New line before keyword</b>	If enabled, a new line is forced before every SQL keyword. If disabled, lines are only forced for statement type and clause type keywords.
<b>BEGIN..END block</b>	If enabled, empty BEGIN..END blocks occupy a single line.
<b>Stack lists</b>	If enabled, a new line and indenting is forced for each item in comma-separated list such as argument lists.
<b>List indent size</b>	If <b>Stack lists</b> is enabled, this control lets you specify the number of spaces each list item will be indented from the current offset.
<b>Parenthesis indent size</b>	Lets you specify the number of indenting spaces for forced lines following an open parenthesis.
<b>Conditions format style</b>	Specifies how conditions in clauses are formatted: CONDITIONS_WRAPPED - lines are not forced before conditions. CONDITIONS_STACKED_WITH_LEADING_OPERATORS - a new line is forced for each condition in a clause, with logical operators, if present for a line, displayed at the start of the line. CONDITIONS_STACKED_WITH_TRAILING_OPERATORS - a new line is forced for each condition in a clause, with logical operators, if present for a line, displayed at the end of the line.
<b>Conditions stack threshold</b>	For <b>Conditions format style</b> selections that specify stacking, this value specifies the minimum number of conditions that must appear in a clause before conditions are stacked.
<b>THEN statements</b>	If enabled, simple THEN clauses are kept to a single line.

## ISQL OPTIONS - CODE ASSIST TAB

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **ISQL - Code Assist** tab. This tab lets you activate and select options for three coding assistance and error-detection features available from the ISQL editor:

- **Code Complete** - lets you insert or replace object names, selected from suggestion lists, as you edit a script. For more information, see [Code Complete](#).
- **Semantic validation** - detects object name references to objects not present in the datasource index. For more information, see [Semantic Validation](#).
- **Real-time syntax checking** - on-the-fly syntax checking is performed as you type. For more information, see [Syntax Checking](#).

The **ISQL - Code Assist** tab has the following settings:

Group/Option		Description
<b>Code Complete</b>	<b>Enable Code Complete</b>	Lets you enable or disable the Code Complete feature. When disabled, Code Complete cannot be invoked automatically or manually.  Because of the overhead in maintaining internal resources used in implementing this feature, a slight performance increase can result from disabling the feature.
	<b>Enable auto-activation and Auto-activation delay</b>	When auto-activation is disabled, the Code Complete feature must be invoked manually. When auto-activation is enabled, the Code Complete feature is invoked automatically each time the interval between keystrokes exceeds the specified <b>Auto-activation delay</b> .
	<b>Insert single proposals automatically</b>	Specifies that if a Code Complete suggestion list would contain only a single suggestion, that suggestion is inserted automatically.
	<b>Fully qualify completions automatically</b>	Specifies that Code Complete results are returned fully qualified, rather than the minimum required to identify the object.
<b>Code Validation</b>	<b>Enable Real-time syntax checking</b>	If enabled, syntax-checking is performed as you type. If disabled, syntax checks must be initiated manually.
	<b>Severity levels for semantic validation problems</b>	Lets you select the severity level (ERROR, WARNING, or IGNORE) associated with detected semantic errors. The WARNING option provides a contrast to syntax errors, which are always flagged with a severity level of ERROR. The IGNORE setting disables the feature. For more information, see <a href="#">Semantic Validation</a> .
	<b>Perform syntax and semantic validation on files smaller than</b>	Places an upper limit on the size of files for which automatic syntax checking and semantic validation are active.
<b>Code Assist Resources</b>	<b>Refresh</b>	Each time you open an ISQL editor window, Rapid SQL obtains schema information for the object types supported for the Code Complete and Semantic Validation features. That information is updated with any modifications that you make and persists until the window is closed. The <b>Refresh</b> button ensures that the information is current in case other users or applications are modifying the schema while the current SQL editor session window is open.
<b>Restore defaults</b>	Restores all settings on the tab to the original defaults.	

## JAVA OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Java** tab. The table below describes the options and functionality on the **Java** Tab:

**NOTE:** Options differ by platform.

Interface Element	Option	Description	Default
<b>Load Java files</b>	<b>Use the DBMS_JAVA package</b>	Lets you schedule your SQL.	Selected
	<b>Use batch file (oracle_home/bin/loadJava.bat)</b>	Uses batch file (oracle_home/bin/loadJava.bat)	Not selected
	<b>Default Encoding option</b>	Leave blank to use the default.	Default
<b>Drop Java files</b>	<b>Use the DBMS_JAVA package</b>	Lets you schedule your SQL.	Selected
	<b>Use batch file (oracle_home/bin/dropJava.bat)</b>	Uses batch file (oracle_home/bin/dropJava.bat)	Not selected

## JDBC OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **JDBC** tab.

Rapid SQL requires a JDBC connection. You have the ability, with this panel of the Options Editor, to set or change options that apply to platform-specific JDBC drivers and to the Java virtual machine (JVM) that is running on the client. The table below describes the options and functionality of the **JDBC** tab:

Option	Description	Default
JVM Options:		
<b>Initial Heap Size</b>	Set the size, in MB, for the repository where live and dead objects comingle with free memory. If memory runs out, JVM executions stop so the garbage can be collected and expunged. Each platform responds differently, so trial and error can help you maximize performance	64 MB
<b>Maximum Heap Size</b>	Set the upper limit for your heap size.	64 MB
<b>Additional Options</b>	Add options here ONLY in consultation with Embarcadero Technical Support	N/A
Default JDBC Driver		
<b>DBMS Type</b>	Select the DBMS appropriate to the driver you want to review/adjust from the drop-down list.	
<b>Search</b>	Click ... to open the JDBC Driver Editor where you can see what drivers for the given platform are available to you. You can assign the JDBC Driver you selected as a default for all platform-specific datasources.	
Sybase		

Option	Description	Default
<b>JConnect metadata procedures are installed on all servers</b>	NOTE: Sybase documentation recommends that if you have transactional components that use jConnect caches, you need to make sure the jConnect Database Metadata stored procedures are installed on the database server. Please consult Sybase documentation for more information.	

## LOGGING OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Logging** tab. The Logging tab lets you set defaults that specify the behavior and placement of SQL Logging and Output Logging.

The table below describes the options and functionality on the **Logging** Tab:

Option	Description	Default
<b>Log all SQL Statements to a File</b>	Indicates that the application should log all of the SQL that it executes to a file. SQL logging provides an audit trail for Rapid SQL. You can examine this log to determine the SQL executed to complete a task.	Not selected
<b>Logfile Path</b>	If you choose to log the SQL generated by the application, specify the drive, directory, and file name.	None
<b>Max File Size</b>	Specifies the maximum size for the logfile. When the logfile reaches this threshold, it automatically starts deleting lines in the logfile (starting with the oldest statements) to remain within the specified size limit.	1024 KB
<b>Truncate</b>	Empties the entire contents of the logfile.	Not available
<b>Log all Output Messages to a File</b>	Indicates that the application should log all server messages sent to the Output window. This type of logging lets you monitor only messages issued by the server versus all SQL logged by the application. You can examine this log to determine all messages the server issued.	Not selected
<b>Logfile Path</b>	If you choose to log the server messages generated in the Output window, specify the drive, directory, and file name.	None
<b>Max File Size</b>	Specifies the maximum size for the output logfile. When the output logfile reaches this threshold, it automatically starts deleting lines in the file (starting with the oldest statements) to remain within the specified size limit.	1024 KB
<b>Truncate</b>	Empties the entire contents of the output logfile.	Not available

## MYSQL UTILITIES OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **MySQL Utilities** tab. Rapid SQL integrates with MySQL utilities. For Rapid SQL to access these utilities, you need to specify their location in the MySQL Utilities Tab of the Options Editor. You can use the MySQL Dump and Import Utilities. The table below describes the options and functionality on the **MySQL Utilities** Tab

Option	Description
mysqldump	Specifies a path for the MySQL dump utility, mysqldump.exe. By default, MySQL installs this utility in the C:\mysql\bin directory.
mysqlimport	Specifies a path for the MySQL import utility, mysqlimport.exe. By default, MySQL installs this utility in the C:\mysql\bin directory.

## ORACLE UTILITIES OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Oracle Utilities** tab.

Rapid SQL integrates with multiple Oracle utilities. For Rapid SQL to access these utilities, you need to specify their location in the Oracle Utilities Tab of the Options Editor. You can use the Oracle Export, Import Utilities, and SQL \* Loader. The table below describes the options and functionality on the **Oracle Utilities** Tab:

Option	Description
Export	Specifies a path for the Oracle Export utility. By default, Oracle installs to C:\Orant\Bin directory.
Import	Specifies a path for the Oracle Import utility. By default, Oracle installs to C:\Orant\Bin directory.
SQL*Loader	Specifies a path for the SQL * Loader utility. By default, Oracle installs to C:\Orant\Bin directory.
Default Bind Size	Specifies the bind size. Default is set to 70KB.

## PERF CENTER OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Performance Center** tab. Rapid SQL lets you customize Performance Center's integration with Rapid SQL. The table below describes the options and functionality on the Perf Center Tab:

Option	Description	Default
<b>Computer Name or IP Address</b>	Specify where the Performance Center Server is installed.	localhost
<b>Port Number</b>	Specify the port for the Apache server or IIS web server.	80
<b>Test</b>	Verifies the settings.	



Option	Description	Default
<b>Connect to the server using</b>	Specify if you want Rapid SQL to open the <b>Web Client</b> or the Performance Center file ( <b>PerfCntr.exe</b> ) within Rapid SQL.	Web Client

## QUERY BUILDER OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the Query Builder tab. Option parameters set on the Options Editor elicit a prompt if there are different options set on an open individual session. Global options override properties set within individual Query Builder sessions.

The table below describes the options and functionality on the **Query Builder** Tab:

Interface Element	Option	Description	Default
<b>Code Generation</b>	<b>Generate Use Database statement</b>	Adds a line of SQL code indicating which database or instance is used in the statement.	Selected
	<b>Generate owner names</b>	Adds a line of SQL code showing the table owner name as part of the query.	Selected
	<b>Include Row Count limits</b>	Includes the output row limit set in the Execution settings.	Not selected
	<b>Generate SQL/92 if supported by DBMS</b>	SQL/92 is a standard for relational database management systems.	Not selected
<b>Execution</b>	To set the maximum number of rows in your result set, type the number in the dialog box. This lessens congestion of server processes when queries execute by setting row count limits.		1000 rows
<b>General</b>	<b>Show Column Data types in Query Diagram</b>	Reveals the data type in each column for tables in the SQL Diagram pane.	Not selected
	<b>Confirm on Item delete</b>	Opens a Confirm Delete dialog box when an item is deleted.	Selected
	<b>Auto Populate Views</b>	Checks syntax every time an execute statement, refresh or copy statement begins.	Not Selected
	<b>Auto Format</b>	Automatically sets style and spacing of display.	Selected
<b>Auto Join</b>	<b>Run Automatically</b>	Automatically detects names and data types, and create joins for multiple tables.	Selected
	<b>Require Indexes</b>	Joins only indexed columns. Requires an indexed column for joins.	Selected
	<b>Require same data type</b>	Automatically joins columns with the same data type.	Selected
<b>Syntax Checker</b>	<b>Automatic Syntax Check</b>	Automatically checks SELECT and CREATE VIEW statements for errors.	Selected
	<b>Warn on non index join</b>	Returns a warning when it detects a join against a non-indexed column, or a column not participating in a primary key	Not selected
<b>Display</b>	Lets you sets the style, size, and color of <b>Column Font</b> and <b>Title Font</b> . Also lets you set the background <b>Table Color</b> for the SQL Diagram Pane.		Available

## REPORTS OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Reports** tab. The **Reports** tab lets you:

- Append a date and time stamp to report file names. This prevents existing reports from being overwritten.
- Save reports in date-specific folders.
- Save reports in PDF format as opposed to HTML.
- Let you include a provided image and title to reports.

## RESULTS (ISQL) OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Results** tab. The table below describes the options and functionality on the **Results** Tab:

Interface Element	Options and descriptions	
<b>Results Option pane</b>	<b>Result Window</b>	<p><b>Single Window</b> - Displays all results in one tabbed result window. Multiple result sets together in the window. Single Window and Multiple Windows options are mutually exclusive.</p> <p><b>Multiple Windows</b> - Displays multiple result sets one result set per window.</p> <p><b>Attached to Editor</b> - Sets results as tabbed windows attached to the ISQL window. Used in conjunction with Single Window option or Multiple Window option. Attached and Unattached options are mutually exclusive.</p> <p><b>Unattached</b> - Sets results appear separate from the ISQL Window. Used in conjunction with Single Window option or Multiple Windows option.</p> <p><b>Reuse Window</b> - Sets new result sets to overwrite any current result sets in an open Result Window. Only valid for Single and Attached to Editor combination.</p>
	<b>Results File</b>	<p><b>Mail File Type</b> - Selects the file type to use when sending result sets via a MAPI-compliant mail package. Valid formats include the proprietary Results type, Tab delimited, Comma separated, and HTML.</p> <p><b>Schedule File Type</b> - Selects the schedule file type. Valid formats include Tab delimited, Comma separated, and HTML.</p> <p><b>Include column titles when saving</b> - Includes column titles when saving a result set. If this option is turned off, column titles Rapid SQL does not save result sets.</p>
	<b>Result Set Options</b>	<p><b>Default Rowcount</b> - Lets you limit the number of rows returned to the result window of the ISQL window (default 0).</p> <p><b>Sybase and SQL Server: Text Size (bytes)</b> - Lets you specify the text size (default 8192).</p> <p><b>Oracle:LONG Size (bytes)</b> - Lets you specify the LONG size (default 8192).</p> <p><b>LOB Preview: Text Size (bytes)</b> - Specifies the length of the preview of LOB column data (default 4096).</p>
	Sort Type	Lets you select a result sorting: <b>Alphanueric</b> or <b>Lexicographic</b> .

Interface Element	Options and descriptions	
<b>Format pane</b>	<b>Column Formatting</b>	<p><b>Auto Format (Best Fit)</b> - Sets column widths automatically to accommodate the longest piece of data in a column. Large queries depend on the longest row formatting, so activating this option can affect performance.</p> <p><b>Use pre-defined column</b> - Lets you select column type and character length</p>
	<b>Enable Date/Time Format</b>	Lets you select the date/time format.
	<b>Format</b>	<p><b>Standard Grid</b> - Displays all result sets in a standard grid format. Rapid SQL displays result sets in grid format in IISQL Editors that are opened after you have selected this option. Rapid SQL does not display IISQL Editors that are already open.</p> <p><b>HTML</b> - Displays all result sets as HTML tables. Rapid SQL displays result sets in HTML format in ISQL Editors that are opened after you have selected this option. Rapid SQL does not display ISQL Editors that are already open.</p> <p><b>ASCII Text</b> - Displays all result sets as ASCII Text. Rapid SQL displays result sets in ASCII Text format in ISQL Editors that are opened after you have selected this option. Rapid SQL does not display ISQL Editors that are already open.</p> <p><b>Grid Font</b> and <b>Printer Font</b> buttons - Opens a Font dialog box, letting you select the font, style, and size for the result sets grid or printed result sets.</p>

## SMTP MAIL OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **SMTP Mail** tab. The SMTP Mail Tab of the Options Editor lets you specify outgoing notification e-mail message options.

The table below describes the options and functionality on the **SMTP Mail** tab:

Option	Description	Default
<b>Send messages through SMTP</b>	Enables SMTP messaging and makes the other controls on this tab available.	Unselected
<b>Name</b>	Name that appear as the e-mail sender.	Name you specified during installation.
<b>E-mail Address</b>	Address to send e-mails from Rapid SQL.	E-mail address you specified during installation.
<b>Authentication</b>	Lets you specify authentication options.	None
<b>User Name</b>	User name for authentication.	Not available
<b>Password</b>	Password for authentication.	Not available

Option	Description	Default
<b>Host Name</b>	SMTP server for outgoing messages. For Microsoft Outlook, select Tools, and then Accounts. On the Mail Tab, select target account, and then click Properties. On Servers Tab, copy the Outgoing Mail(SMPT) and paste.	Host Name you specified during installation.
<b>Port Number</b>	Port number you connect to on your outgoing SMTP server.	25
<b>Test</b>	Opens an SMTP Configuration Test e-mail addressed to your e-mail address. Click Send Mail to send the e-mail.	Available
<b>Bind to</b>	Your IP address the message is bound to.	ANY_IP_ADDRESS
<b>Encoding</b>	E-mail encoding.	Western Europe (ISO)
<b>Send messages Mime encoded</b>	Messages encoded using Multipurpose Internet Mail Extensions (MIME) support enriched content and attachments.	Selected
<b>Send all messages as HTML</b>	Messages include text formatting.	Selected
<b>Auto Connect to the Internet</b>	Rapid SQL connects to internet at launch.	Selected

## VERSION CONTROL OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Version Control** tab. The **Version Control** Tab of the Options Editor lets you select which version control system you want Rapid SQL to use as the underlying version control system. The table below describes Version Control options:

Item	Description	Default
Specify Version Control System	Lists the Version Control systems that integrate with Rapid SQL. Select the option button for the product you want Rapid SQL to use.	None
Optional Fields	If you select Microsoft Visual SourceSafe, Rapid SQL uses your default SourceSafe database unless you override this setting by entering the path for the SourceSafe file associated with the desired database. To use a SourceSafe database other than the default, provide a user name, password, and the location of a scrsafe.ini file.	Optional

## WARNINGS OPTIONS

After opening the Options editor (see [Specifying Rapid SQL application and feature options](#)), you can make changes to the **Warnings** tab. Rapid SQL lets you issue warning messages to users whenever improper operations are attempted in a database. Warning messages differ by platform.

The tables below describes the options of the **Warnings** Tab:

Option	Description	Platform
<b>Create an index on the same tablespace as the associated table</b>	Issues a warning message whenever a user does not create an index on a different tablespace than the table. This makes it possible for the server to read the index and table data in parallel.	DB2
<b>Create an index on the same tablespace as the associated table</b>	Issues a warning message whenever a user does not create an index on a different tablespace than the table. This makes it possible for the server to read the index and table data in parallel.	Oracle
<b>Create an object in the SYSTEM tablespace</b>	Issues a warning message whenever a user tries to create or place an object on the SYSTEM tablespace.	Oracle
<b>Create a user with default or temp tablespace as the SYSTEM tablespace</b>	Issues a warning message when a user is created with a default or temp tablespace on the SYSTEM tablespace.	Oracle
<b>Create an object in the master database</b>	Issues a warning message when an object is created in the master database.	Microsoft SQL Server and Sybase ASE
<b>Create a table or index on the default segment</b>	Issues a warning message when a table or index is created on the default segment.	Microsoft SQL Server and Sybase ASE
<b>Create a non-clustered index on same segment as the associated table</b>	Issues a warning message when a non-clustered index is created on the same segment as the associated table.	Microsoft SQL Server and Sybase ASE



# USING

Using shows how to use each component of the application. Using breaks each process and procedure into simple step-by-step instructions.

Using is divided into sections. The table below describes each section:

Section	Description
<a href="#">Datasource Management</a>	This section describes the purpose of datasources, the process of establishing datasource connections, and managing your datasources and datasource groups.
<a href="#">Supported Objects</a>	This section describes the different platform objects the application supports.
<a href="#">Modifying objects using editors</a>	This section describes Object editors. Editors let you modify existing objects.
<a href="#">Creating objects</a>	This section describes Object wizards. Editors let you create new objects.
<a href="#">Functionality</a>	This section describes the different platform object functionalities.
<a href="#">SQL Scripting</a>	This section describes the SQL scripting environment, that lets you write, debug, test and deploy solid SQL code for your database applications.
<a href="#">Permissions Management</a>	This section describes the Permissions Management features.
<a href="#">Project Management</a>	Describes Rapid SQL's project-based features.
<a href="#">Version Control</a>	Describes how Rapid SQL integrates with version control software.





# DATASOURCE MANAGEMENT

After installing Rapid SQL, you must set up [datasources](#) to establish reusable connections to your database servers. A datasource is a database connection profile that is similar to the connection information that you have stored in your SQL.INI or WIN.INI files. Rapid SQL stores information about the datasource specification in the system registry and provides a visual interface for maintaining it.

The Datasource Registration Wizard guides you through the required steps to establish a connection to your server and makes the process of setting up datasources, easier. For details, see [Registering or editing datasources](#).

This section describes the function of datasources, the process of establishing datasource connections and managing your datasources and datasource groups.

## DATASOURCES

A datasource is a database connection profile. A datasource includes:

- Name
- Connection String
- Default User ID
- Optional Password Specification

All database operations are performed through datasources. You can create a datasource profile for each database instance (or database server) in your enterprise, and you can create multiple datasource profiles with different logins for a single database. The table below describes the data items stored for each datasource:

Data Item	Description
Name	A unique, user-defined name for the datasource.
Connection String	For Oracle: the SQL*Net connect string, for Sybase ASE: the database server name.
Default User	Default user ID to use when logging in to the datasource. Can be null.
Default Password	Default password to use when logging in to the datasource. This is encrypted. Can be null.
Auto-Connect Flag	If Yes, then automatically login using default user and password. If No, open Login dialog box.
Default Database	SYBASE ASE ONLY: Database to automatically use after logging in.

### Available Functionality

Rapid SQL offers the following functionality for the Datasource Explorer:

- [Change Group](#)

- [Connect](#)
- [Disconnect](#)
- [Drop](#)
- [Registering or editing datasources on page 115](#)
- [New LUW Datasource](#)
- [Register Datasource](#)
- [Unregister Datasource](#)

#### Related Topics

- [Selecting Datasources](#)
- [Viewing Datasource Properties](#)

## UNDERSTANDING THE DATASOURCE CATALOG

A Datasource Catalog is a collection of defined datasources. On startup, Rapid SQL loads a Datasource Catalog using one of three methods, specified using the Options editor:

- **Windows Registry datasource catalog** - The datasource catalog is stored and managed locally, in the Windows registry. This user can add and delete datasources and edit the registration details of existing resources. Any changes are retained across startups of Rapid SQL.

This method is compatible with other Embarcadero products using registry-based datasource definition storage. If more than one of those products is installed on the same machine, they can share the same Datasource Catalog.

- **File based datasource catalog** - The datasource catalog is stored locally and is file-based. Each datasource definition included in the catalog is stored in a single file (.dsd file suffix) in the %APPDATA%\Embarcadero\Data Sources\ folder, in a default installation. This user can add and delete datasources and edit the registration details of existing resources. Any changes are retained across startups of Rapid SQL.

This method is compatible with other Embarcadero products using file-based datasource definition storage. If more than one of those products is installed on the same machine with Rapid SQL, all can share the same Datasource Catalog.

- **Network shared datasource catalog** - The datasource catalog is file-based, and obtained from a user-specified location such as a network share. At each startup, Rapid SQL imports definitions from the network shared file (.etsds file suffix).

Changes such as new datasources, deleted datasources or edited datasource registrations are lost as soon as this user shuts down Rapid SQL. This method allows centralized maintenance of the Datasource Catalog and sharing of a single catalog among multiple users across a network.

**NOTE:** The definition files used in both the **File based datasource catalog** and **Network shared datasource catalog** methods do not store user ID and password credentials. Credentials, links to definition files in the current user's catalog, and folder structure are stored in the %APPDATA%\Embarcadero\Data Sources\metadata\dsuri.xml file.

For information on choosing a method and specifying the network location of a Shared Datasource Management definition file, see [Datasource Options](#).

Regardless of the catalog storage method, the Datasource Catalog can be built manually by explicitly registering datasources or using automated methods such as the Discover Datasources feature. For details, see [Automatically Discovering Datasources](#) and [Registering or editing datasources](#). Keep in mind that changes made with **Network shared datasource catalog** in effect, are lost on shutdown of Rapid SQL.

In order to set up a **Network shared datasource catalog** scheme, you can build a catalog first using the **File based datasource catalog** or **Windows Registry datasource catalog** method. You can subsequently use the **Manage Datasources** facility to export your **Network shared datasource catalog** file to the location that others will use to load the file. For details, see [Importing and Exporting Datasource Definitions](#).

More generally, the **Manage Datasources** facility's import/export functions can be used to exchange datasource definitions between users employing the different storage methods. Even more simply, these functions save time in allowing datasource definitions to be defined once and then shared with multiple users.

## AUTOMATICALLY DISCOVERING DATASOURCES

The first time you run Rapid SQL a dialog box displays, giving you the option to Auto-Discover all configured datasources. If you click Yes, the Rapid SQL Auto-Discover feature searches the DBMS configuration files on your computer and automatically discovers all the datasources that you are licensed for. For example, if you have a cross-platform license, Discover Datasources finds all unregistered datasources. If you have an Oracle only license, Discover Datasources finds all unregistered Oracle datasources.

**NOTE:** Microsoft SQL Server datasources are registered through a Windows NT system call to your network. Provide login information (user name and password) the first time you [connect](#) to a datasource.

**NOTE:** IBM DB2 LUW for Linux, Unix, and Windows databases use ODBC/CLI or DB2 (attach) to connect. Therefore, you need the proper ODBC/CLI Connection established in order for the auto-discover feature to find your IBM DB2 LUW for Linux, Unix, and Windows databases, including registering the DB2 datasource to ODBC as a system datasource. Although your datasources are auto-discovered, provide login information (user name and password) the first time you connect to a datasource.

In addition to Auto-Discovering your database servers, the application creates Datasource Groups based on RDBMS type. Each registered datasource is placed in its respective Datasource Group. For example, all Microsoft SQL Server datasources are added to the Microsoft SQL Server Group. Each registered datasource is placed in its respective Datasource Group.

For more information on how to configure your datasources, see [Working with Datasources](#).

## CHANGING DATASOURCE GROUPS

Rapid SQL lets you change datasource groups by:

- Dragging the datasource between groups.
- Invoking the Change Group dialog box.

### Dragging and Dropping Between Groups

- 1 On the Datasource Explorer, left-click the datasource group you want to move, drag it over the folder for the new group, and release the pointer.

### Using the Change Group Dialog Box

Disconnect your datasource before changing groups.

- 1 On the Datasource Explorer, right-click the datasource you want to move, and then click *Change Group*.

Rapid SQL opens the Change Group dialog box.

- 2 In the *Select Group tree*, click new group.

### 3 Click OK.

Rapid SQL changes groups.

For more information, see [Datasource Groups](#).

## CONNECT

Rapid SQL lets you set datasources to [automatically connect](#) each time you open the application. The first time you start the application, Rapid SQL prompts you to register your datasources. During this process, you can select the Auto Connect check box, which automatically connects all registered datasource each subsequent time you open the application.

If you did not check the Auto Connect box, or if you clicked No when prompted to connect to a database after registering, you must connect manually, each time you want to access that datasource, using the [Datasource Login dialog box](#). If you later want to automatically connect your datasources, edit the datasource accordingly. See [Registering or editing datasources on page 115](#).

The table below describes the options and functionality on the Datasource Login dialog box:

Option	Description
Login ID	Lets you type the Login ID.
Password	Lets you type the password.
Auto Connect	Select to automatically connect to the datasource in the future, select the check box.

For more information, see [Completing the Datasource Login Dialog Box](#).

## COMPLETING THE DATASOURCE LOGIN DIALOG BOX

To complete the Datasource Login dialog box, do the following:

- 1 On the *Datasource* menu, click Connect.

OR

On the *Datasource* tool bar, click Connect.

OR

On the *Datasource Explorer* tool bar, click Connect.

OR

On the *Datasource Explorer*, right-click the datasource, and then click Connect.

OR

On the *Datasource Explorer*, double-click the datasource.

Rapid SQL opens the Datasource Login dialog box.

- 2 In Login ID type the Login ID.
- 3 In Password type the password.
- 4 Select Auto Connect to automatically connect to the datasource in the future.
- 5 Click Connect.

Rapid SQL opens the Datasource Landing Page in the right pane of the application.

For more information, see [Working with Datasources](#).

## DISCONNECT

When you disconnect from a server, the application immediately breaks the connection between any open ISQL Windows, the servers, and databases. Although your ISQL Windows are still visible, the connections are no longer valid. If you attempt to execute a script, Rapid SQL attempts to reconnect to a registered datasource, if available.

For more information, see [Completing the Disconnect Dialog Box](#).

## COMPLETING THE DISCONNECT DIALOG BOX

To complete the Datasource Login dialog box, do the following:

- 1 On the *Datasource* menu, click Disconnect.

OR

On the *Datasource* tool bar, click Disconnect.

OR

On the *Datasource Explorer* tool bar, click Disconnect.

In the right pane of the *Datasource Explorer* window, right-click the datasource, and then click Disconnect.

Rapid SQL opens a dialog box asking if you want to commit all pending transactions for that connection or to rollback all before disconnecting. You cannot disconnect if there is an uncommitted transaction.

- 2 Click Yes.

Rapid SQL confirms you want to disconnect and closes the dialog box.

## DISCOVER DATASOURCE

Rapid SQL discovers datasources residing on your system that are not currently registered datasources through a Windows NT system call to your network. The [Discover Datasource dialog box](#) includes a list, which includes the name of the server or instance and the type of DBMS of all unregistered datasources found on your network or local machine. Once discovered, you have the option to register datasources.

### Completing the Discover Datasources Dialog Box

- 1 On the *Datasource* menu, click *Discover Datasource*.

Rapid SQL opens the Discover Datasources dialog box.

- 2 Select the *check box* next to the datasource you want to register.
- 3 Click Select All to select all the datasources on the list.
- 4 Click Register.

Rapid SQL registers the datasource or datasources selected.

- 5 Click OK.

Rapid SQL closes the Rapid SQL message.

For more information, see:

[Datasources](#)

## [Working with Datasources](#)

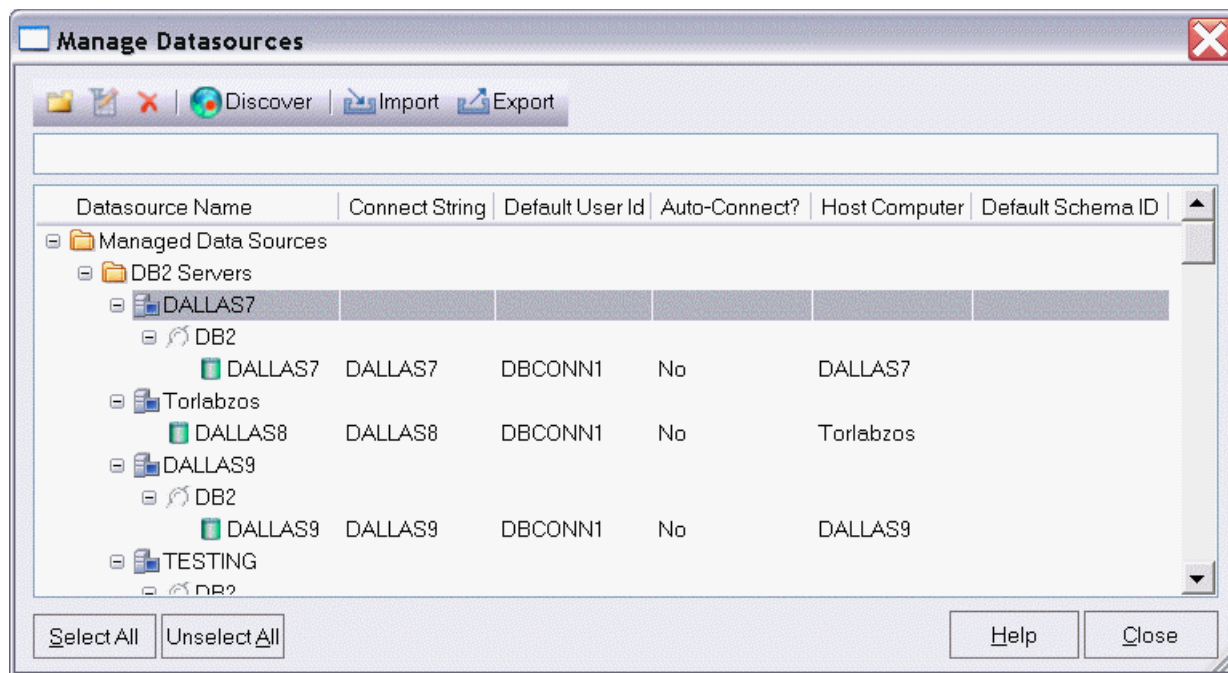
### MANAGING DATASOURCES

The Manage Datasources dialog box lets you manage datasources throughout your enterprise from a single vantage point. It provides concise, relevant information on your datasources in a simple grid format. It provides centralized access to common datasource operations such as adding, modifying, deleting, discovering, importing, and exporting datasources.

#### To open the Manage Datasources dialog box:

- 1 On the *Datasource* menu, click *Manage Datasources*.

Rapid SQL opens the Manage Datasources dialog box.



The grid format lets you access and view datasource specifications. datasources are grouped according to the default, DBMS-based folder structure and then within any custom datasource groups you have set up. The table below describes information available for each datasource entry:

Column	Description
Datasource Name	Uses an explorer-like interface to display all registered datasources and their groups. You can navigate this column in the same manner as the datasource explorer, by clicking on nodes to expand or collapse your view of the datasources.
Connect String	Displays the full connection string for the datasource.
Default User Id	Displays the Default User ID for the datasource.
Auto-Connect?	Indicates whether the Auto Connect feature is turned on or off.



Column	Description
Host Computer	Displays the name of the Host Computer if one has been configured.
Default Schema	Displays view default schemas for your DB2 datasources.

In addition, you can take the following actions:

- Click the New Datasource button to register a datasource. For more information, see [Registering or editing datasources](#).
- Select a datasource and click the Edit Datasource button to edit that datasource definition. For more information, see [Registering or editing datasources](#).
- Select a datasource and click the Remove Datasource button to unregister that datasource definition. For more information, see [Unregistering Datasource](#).
- Click the **Discover** button to locate all configured datasources on your network and automate the process of registering them. For more information, see [Automatically Discovering Datasources](#).
- Import or export datasource definitions. For details, see [Importing and Exporting Datasource Definitions](#).

## IMPORTING AND EXPORTING DATASOURCE DEFINITIONS

On startup, Rapid SQL loads a Datasource Catalog, a collection of datasource definitions that the user can work against. How the catalog is stored and whether changes to the catalog are temporary or permanent, depends on the datasource storage method used. For an introduction, see [Understanding the Datasource Catalog](#).

Regardless of the storage method used, Rapid SQL lets you add datasources to the currently loaded catalog by importing datasource definitions from the following sources:

- Windows Registry - if you previously used the **Windows Registry datasource catalog** storage method but are now using one of the file-based methods
- **File based datasource catalog** storage method files (.dsd file suffix)
- **Network shared datasource catalog** storage method files (.etsds file suffix)

Similarly, you can export the currently loaded Datasource Catalog definitions to **File based datasource catalog** or **Network shared datasource catalog** storage method files.

The ability to export and import datasource definitions provides a number of practical advantages. For example:

- Simple sharing among users - if you have a large number of datasources to register, you can have one user walk through the process of registering each datasource and then export those definitions. Other users can then import the resulting datasource definition files.

- Sharing definitions among users employing different datasource storage methods - datasource definitions or entire catalogs created using the **Windows Registry datasource catalog** or **File based datasource catalog** methods can be exported for use by users employing any of the three methods.

#### To export one or more datasource definitions:

- 1 On the *Datasource* menu, click *Manage Datasources*.

Rapid SQL opens the **Manage Datasources** dialog box.

- 2 Select a single datasource, multiple datasources, or a default or custom datasource group folder.

**NOTE:** Datasource group folder information is not exported. Datasources are exported as individual files with no higher level grouping.

- 3 Click **Export**. An **Export Data Sources Folder** dialog opens.

- 4 Take one of the following actions:

- To export datasource definitions to individual files that can be imported on a datasource-by-datasource basis, select **Export Data Sources to Individual Data Source Files** and use the associated control to locate and select a target folder.
- To export datasource definitions to a single file for use by installations employing the Shared Datasource Managed catalog storage method, select **Export Data Sources To a Shared Data Source Management File** and use the associated control to locate and select a target folder and provide a file name. For information on how to designate a user as a Shared Datasource Managed catalog storage method, see [Datasource Options](#).

- 5 Click **OK**.

**NOTE:** User ID and password credentials are not exported to the datasource definition file.

After export, other users can then import the datasource definitions.

#### To add one or more datasource definitions to a catalog:

- 1 On the *Datasource* menu, click *Manage Datasources*.

Rapid SQL opens the **Manage Datasources** dialog box.

- 2 Select the datasource group folder where the files or files are to be imported.
- 3 Click **Import**. An **Import Data Sources** dialog opens.

4 Take one of the following actions:

- If you previously used the Windows Registry catalog storage method but are currently using the Locally Managed or Shared Datasource Managed method, select **Import Previously Registered Embarcadero Data Sources From Registry** to load any datasources registered when you were using the Windows Registry catalog method.
- To import datasource definitions from individual datasource definition files (.dsd file suffix), select **Import Previously Registered Embarcadero Data Sources From File(s)**, click **Add**, and use the **Select datasource definitions to import** dialog to locate and select the files to be added.
- To import datasource definitions from a single Shared Datasource Managed storage method file (.etsds file suffix) file, select **Import Data Sources From a Shared Datasource Management File** and use the associated control to locate and select the target file.

5 Click **OK**.

After importing a datasource definition file, you can either edit the datasource registration to provide security parameters or provide credentials when connecting to the datasource.

## MANAGING DATASOURCE PROPERTIES

The [Datasource Properties dialog box](#) displays the name, type, version, status and mode of the datasource. The Datasource Properties box also lets you view the middleware or connectivity software that is being used to establish a particular datasource connection. You can use this information to troubleshoot connectivity problems, determining vital information such as the server version, connectivity library used, and library version and date.

### Completing the Datasource Properties Dialog Box

- 1 On the Datasource Explorer, click a datasource with an established connection.
- 2 On the *Datasource* menu, click *Properties*.

Rapid SQL opens the Datasource Properties dialog box.

For more information, see:

[Datasources](#)

[Working with Datasources](#)

## REGISTERING OR EDITING DATASOURCES

Each database instance must be registered. Whether you are registering a new data source, registering a new data source based on an existing data source definition, or editing preexisting connection information, you use the Datasource Registration dialog box.

**To register a new data source**

- On the Explorer, right-click **Managed Datasources**, and select **Register Datasource** from the context menu.

**To register a new data source based on an existing definition**

- In the left pane of the explorer, select the existing data source and then from the **Datasource** menu, select **Register Like**.

**To edit registration information**

- On the **Explorer**, right-click a data source, and select **Edit Datasource Registration** from context menu.

The table below describe the fields you will encounter when registering a data source or editing the connectivity information:

Panel/Tab	Option	Description
<b>Database Type</b>		Select the radio box corresponding to the DBMS type for the data source. If there is currently a DBMS node selected in the Explorer, that database type is automatically selected.
<b>Connection Information</b>		For detailed descriptions of the options on this tab, see <a href="#">Data Source connection information on page 119</a> .
<b>Advanced</b> (only available on clicking the <b>Advanced</b> button on the <b>Connection Information</b> tab)	<b>JDBC Driver</b>	You should see the JDBC Driver that's appropriate to the platform. If you select ..., the JDBC Driver Editor opens letting you browse to a new or possibly undetected JDBC driver. For details, see <a href="#">JDBC Driver Editor</a> .  NOTE: If you select a driver other than the default, you'll be queried about whether you want that driver to be used for data sources you register going forward. Data sources already registered will not be affected.
	<b>Connection URL</b>	The connection URL cannot be edited--it's for reference only.
	<b>URL Check</b>	Indicates whether the connection URL is valid or not.

Panel/Tab	Option	Description
<b>Security Parameters</b>	<b>User ID</b>	The User ID that Rapid SQL will use to connect to the data source.
	<b>Password</b>	The password associated with the User ID.
	<b>Connect As:</b> (Oracle)	When relevant, choose the appropriate user/administrator level.
	<b>Domain:</b>	For MS SQL only. Identify the domain if the user has restricted access.
	<b>Auto connect?</b>	Spares the user from reentering the password every time you connect.
	<b>Connect using Windows Authentication or Connect using OS Authentication</b> (IBM DB2 for Windows, Unix, and Linux, MySQL, Oracle, SQL Server)	Login to the server is verified using Windows/OS authentication
	<b>Connect using Kerberos Authentication</b> (Sybase)	If this option is selected, login to the server is verified using Kerberos authentication.  <b>NOTE:</b> This is only available if the <b>Use alias information in the SQL.INI file/Sybase Server</b> option on the <b>Connection Information</b> tab is selected. For details, see <a href="#">Data Source connection information</a> .
	<b>Test Connection</b>	When you test the connection, you're testing both the native and JDBC connectivity. If the connection fails, read the error message and backtrack as necessary.
<b>Datasource Properties</b>	<b>#Include Search Directories</b>	Enter one or more paths on this datasource, which will be searched for files in conjunction with use of the <b>#include</b> directive in the ISQL editor, Procedure Object Editor, or Package Body Object Editor. Separate multiple paths using semicolons. For example:  c:\myscripts;c:\Program Files\Scripts For more information on use of the #include directive, see <a href="#">Preprocessing #define and #include Directives</a> .  Note that if there are no entries specified here, the directory specified on the <b>Directories</b> tab of the Options editor will be searched. For more information, see <a href="#">Directories Options</a> .
	<b>Database device default path</b>	The database device default path is used when creating a new database device.

Panel/Tab	Option	Description
<b>Oracle View Config</b> (Oracle only)		Lets you manually configure Rapid SQL access to Oracle Data Dictionary DBA or ALL/USER views to conform to the view privileges associated with a specific role.  <b>NOTE:</b> In order to enable manual configuration of Oracle Data Dictionary views for Oracle datasources and make this panel available, you must select <b>Use Datasource Configuration for Views</b> in the Options Editor ( <b>File &gt; Options &gt; General &gt; Oracle</b> ). For more information, see <a href="#">General Options</a> .  Only DBA views with an ALL/USER equivalent are offered on this panel. You can use the following methods to select DBA or ALL/USER type views on a view-by-view basis:
	Manual selection using the <b>Use DBA View</b> check boxes	Selecting the <b>Use DBA View</b> check box associated with an individual view makes the DBA group version of that view available to Rapid SQL. Deselecting the <b>Use DBA View</b> check box associated with an individual view makes the ALL/USER group version of that view available to Rapid SQL.
	<b>Reset to All</b>	Deselects all <b>Use DBA View</b> check boxes, making only ALL views accessible to Rapid SQL.
	<b>Reset to DBA</b>	Selects all <b>Use DBA View</b> check boxes, making DBA views accessible to Rapid SQL.
	<b>Resolve Automatically</b> (only available when editing a connected datasource)	Queries the datasource and selects the <b>Use DBA View</b> check box for each DBA view to which the current user has access.
	Changes take effect as soon as you click the <b>Finish</b> button.	
<b>Datasource Group</b>	<b>Select the datasource group folder</b>	Select the folder that is to contain this datasource.
	<b>Select a category</b>	If you want to use category-based user interface item color-coding and labelling for this datasource, this dropdown lets you select an existing category to assign this datasource to that category. You can also select <b>Customize</b> to manage your categories. For details, see <a href="#">Customizing Datasource Categories</a> .  For an introduction to datasource categorization and related tasks, see <a href="#">Categorizing datasources using color-coding and labelling</a> .

After clicking **Finish**, you are prompted to connect to the newly registered data source.

**NOTE:** On initial connection to a Sybase Cluster edition data source, you are offered the choice to connect to the entire cluster. If you select **NO**, your connection scope is an instance. This choice sets the default connection scope for this data source. You can subsequently change the connection scope for the data source by editing the data source.

**NOTE:** After you successfully connect to a data source, you can see the name of the host and the connection string information displayed at the bottom of the Rapid SQL window. For Oracle and Sybase ASE, should a question arise, you can see if perhaps you set up the data source NOT to use the alias file. You can also use these cues to remind yourself that if you changed the alias file, that change will not be reflected here unless you very specifically made a change to the defaults in the Options Editor-JDBC tab.

## DATA SOURCE CONNECTION INFORMATION

The following table lists the **Connection Information** tab options, DBMS-by-DBMS, available when registering or editing a data source.

DBMS	Option	Description
IBM DB2	<b>Server</b>	The name of the server.
	<b>Schema ID</b>	This is an optional field. If you want to include the current schema as part of the data source connection properties, you can set the implicit schema for unqualified object references to a schema different from the user's login id.
	<b>Function Path</b>	This is an optional field. If you choose, you can enter the appropriate function path name so the search scope of functions will not be limited to the IBM DB2 library.
	<b>Datasource Name:</b>	This field is automatically populated with the server name, but you can rename it to whatever you want.
Interbase/Firebird	<b>Host name</b>	The name of the machine where the database resides.
	Port	Optionally, provide a port number.
	<b>Database file path</b>	The full path to the database file, a .gdb file for Interbase or a .fdb file for Firebird.
	<b>Datasource Name</b>	Optionally, override any value in this field with a value to be displayed in the Rapid SQL Datasource Explorer tree.
ODBC	<b>Connect String</b>	A valid ODBC connect string.
	<b>Datasource Name</b>	The name that will be displayed in the Rapid SQL Datasource Explorer tree.
SQL Server	<b>Use Network Library Configuration</b>	When you select this, you can connect to a Net-Library that is listening on the server that is configured to either the Named Pipes or TCP/IP protocols.
	<b>Alias</b>	This box is enabled only when you opt to use the Network Library Configuration
	<b>Protocol: TCP/IP or Named Pipes</b>	You need to select one or the other to register a SQL Server database.
	<b>Host</b>	Enabled only when you have not selected an Alias from the Network Library.
	<b>Port (optional)/ Pipe Name</b>	Depending on your section, you can optionally indicate the port for TCP/IP or the pipe name depending on your means of connection.
	<b>Default DB (optional)</b>	Optionally give the default name for the database.
	<b>Datasource Name</b>	The field is automatically populated with the host name, but you can rename to whatever you want.
MySQL	<b>Server</b>	Enter the name of the host, for example doctest01.
	<b>Port (optional)</b>	The default port is 3306. You do not need to enter this information.
	<b>Default Database</b>	You must enter the name of the default database.
	<b>Datasource Name</b>	This field is automatically populated with the server name, but you can rename it to whatever you want.

DBMS	Option	Description
Oracle	<b>Use TNS Names Alias/Oracle Alias</b>	If this option is checked, look for the Alias you want to register from the drop-down list. The remaining fields will automatically populate. If it is unchecked, you must manually enter the requisite information.
	<b>Host</b>	Manually enter the name of the host machine.
	<b>Port</b>	The default is 1521, but you can change it to wherever the Oracle listener is set up.
	<b>SID/Service Name / Type: SERVICE_NAME or SID</b>	Enter the SID or Service Name to correspond with the option you select.
	<b>Instance Name</b>	The specific name used to identify the Oracle instance (the SGA and the Oracle processes).
	<b>Datasource Name</b>	The field is automatically populated with the host name, but you can change it.
Sybase	<b>Use alias information in the SQL.INI file/Sybase Server</b>	If this option is selected, choose a name from the drop-down list in the Sybase Server field.
	<b>Use SSL encryption</b>	If this option is selected, the JDBC connection will be established using SSL encryption
	<b>Host</b>	If you selected the Alias option, this field automatically populates. Otherwise, you need to manually enter the name of the host.
	<b>Port</b>	There is no default for this optional field.
	<b>Default DB (optional)</b>	Optionally give the default name for the database.
	<b>Datasource</b>	This field automatically populates with the server name. You can change it to whatever you want.
	<b>Connection scope</b> (Cluster Edition only)	Lets you specify a connection to the entire cluster or to a single instance.

For context information, see [Registering or editing datasources](#).

## JDBC DRIVER EDITOR

The following table lists the controls available on the JDBC Driver Editor dialog. You use this dialog when defining or editing a data source.

Required Information	Description
Driver Type	The database platform.
Driver List	You see what platform-specific drivers are available to you and which driver (highlighted) is currently associated with the server you are connecting to.
Driver Name	
Driver Class	The name of the class the driver uses for creating connections to the server.
Version	The driver's version.



Required Information	Description
Driver supports the use of native aliases	Rapid SQL cannot detect whether the driver you are using can support native aliases. You need to know by checking the driver's documentation.
Required jars	Rapid SQL detects the required jars and lists the path to them in the little window. You can select or add new jar files as required by your driver.

## CUSTOMIZING DATASOURCE CATEGORIES

Rapid SQL lets you customize your datasource category scheme. A datasource category has the following configurable components:

- **Category name** - The name displayed as a selection in the **Select a category** dropdown
- **Short Name** - The abbreviation shown in the Datasource toolbar dropdown when a categorized datasource is selected.
- **Color** - The color used to denote a categorized datasource in Datasource Explorer tree icons and window tabs.

**NOTE:** Display of specific user items used to denote categories is configurable. For more information, see [Datasource Options](#).

### To customize your datasource categories

- 1 On the **Datasource** menu select **Register Datasource** and then select the **Datasource Group** panel.
- 2 From the **Select a category** dropdown select **Customize**. The **Manage Datasource Categories** dialogs opens.
- 3 Take one of the following actions:
  - Create a new category by clicking **New** and selecting or providing a name, short name and color combination.
  - Edit an existing category by selecting the category, clicking **Edit** and modifying the name and color combination.

**NOTE:** The short name for a category cannot be edited.

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

## SELECTING DATASOURCES

The Rapid SQL [Select Datasource dialog box](#) lets you select a datasource and connect to it.

## Completing the Select Datasource Dialog Box

- 1 On the *Datasource* menu, click *Select* to open the *Select Datasource* dialog box.
- 2 Click the *Datasource* list box, and then click the target datasource.
- 3 Click *Connect* to connect to the datasource.
- 4 Click the *Database* list, and then click the target database.
- 5 Click *OK* to close the *Select Datasource* dialog box.

For more information, see:

[Datasources](#)

[Working with Datasources](#)

## UNREGISTERING DATASOURCE

Rapid SQL lets you unregister datasources when you no longer need them.

**TIP:** Removing a datasource from Rapid SQL does not delete the physical database. It simply removes the datasource definition, and connection information, from the Rapid SQL catalog.

- 1 On the *Datasource* menu, click *Unregister*.

OR

On the *Registration* tool bar, click *Unregister*.

OR

On the *Datasource Explorer*, right-click the datasource, and then click *Unregister Datasource*.

Rapid SQL opens a dialog box.

- 2 Click *Yes*.

Rapid SQL confirms you want to unregister the datasource.

**NOTE:** The datasource manager is shared across Embarcadero's database management products. When you remove a datasource in any of Embarcadero's database management tools the datasource is removed across all relevant products.

For more information, see:

[Datasources](#)

[Working with Datasources](#)

## DATASOURCE GROUPS

Rapid SQL lets you define datasource groups to organize the datasources in your enterprise. Datasource Groups behave as folders in the Windows Explorer, allowing you to group related datasources together. If you manage or monitor many Microsoft SQL Servers, datasource groups are a great mechanism for alleviating desktop clutter.

Anywhere that datasources are presented in a hierarchical tree format, datasource group folders expand to display one or more contained datasources. Upon installation of the first Embarcadero database management product, an initial datasource group is called Managed Datasources. You can rename this group.

### Available Functionality

Rapid SQL offers the following functionality for Datasource Groups:

- [Removing a Datasource Group](#)
- [New Datasource Group](#)
- [Rename Datasource Group](#)

## REMOVING A DATASOURCE GROUP

Rapid SQL lets you remove database groups that you no longer need to access, or that have become obsolete. Keep the following in mind when removing datasource groups:

- If a datasource group contains datasources or other datasource groups, those items are deleted when you delete the containing group.
- The default datasource groups correspond to the DBMS platforms supported by Rapid SQL. If you delete a default group, it will be recreated if you subsequently register a datasource for the associated platform.

### To remove a datasource group

- 1 Right-click the datasource, and then click Remove. Rapid SQL prompts you to verify.
- 2 Click Yes.

For more information, see [Datasource Groups](#)

## NEW DATASOURCE GROUP

Rapid SQL lets you define datasource groups to organize the datasources in your enterprise. Datasource Groups behave as folders in the Windows Explorer, allowing you to group related datasources together. If you manage or monitor many Microsoft SQL Servers, datasource groups are a great mechanism for alleviating desktop clutter.

Anywhere that datasources are presented in a hierarchical tree format, datasource group folders expand to display one or more contained datasources. Upon installation of the first Embarcadero database management product, an initial datasource group is called Managed Datasources. You can rename this group.

### Completing the New Datasource Group Dialog Box

- 1 On the *Datasource Explorer*, right-click the datasource group folder, and then click New.

Rapid SQL opens the New Datasource Group dialog box.

- 2 In the *Datasource Group Name* box, type the new name of the datasource group.

- 3 Click OK.

Rapid SQL closes the New Datasource Group dialog box.

## RENAME DATASOURCE GROUP

Rapid SQL lets you:

- Rename a datasource group.
- Change the members of a datasource group.
- Change the group to which a datasource belongs.

### Completing the Rename Datasource Group

- 1 On the *Datasource Explorer*, right-click the datasource group folder, and then click *Rename*.

Rapid SQL opens the Rename Datasource Group.

- 2 In the *Datasource Group Name* box, type the new name of the datasource group.

- 3 Click OK.

Rapid SQL closes the Rename Datasource Group dialog box.

4 On the *Datasource* menu, click Shutdown. On the *Datasource* menu, click Configure.

Option	Description
Normal	Shuts down the server in an orderly fashion. disables logins, and waits for currently executing Transact-SQL statements and stored procedures to finish.
Immediate	Shuts down the server immediately. does not perform checkpoints in every database. The server terminates all user processes and rolls back any active transactions.

OR

On the Utilities tool bar, click Configure to open the Server Configuration dialog box.

OR

On the Datasource Explorer tool bar, click Configure to open the Server Configuration dialog box.

Option	Description
Edit Button	Click to modify the target parameter. Opens the <a href="#">Edit Configuration dialog box</a> .
New Value	Lets you type the value for the parameter.

Option	Description
New Value	Lets you type the value for the parameter.

On the *Datasource* menu, click Session Recording.

OR

On the *Main* tool bar, click Session Recording. On the *Datasources* menu, click Session

Option	Description
Session Name	Lets you type the name of the session.
Session File	Lets you type the location and session file name or click the browse button. uses *.ses file extension for session files.

Recording.

OR

On the *Main* tool bar, click Session Recording. On the *File* menu, click Open.

OR

On the *Main* tool bar, click Open. On the *Datasource* menu, click Find Object.

OR

On the *Main* tool bar, click Find.

OR

On the Datasource Explorer, right-click a database, and then click Find Object.

### Changing an ODBC Datasource to a Native Datasource

Option	Description
Object to Find	Lets you type the target string of text. You can also click the list to choose from a list of up to ten previous search strings.
Type of Object	Lets you select the target database object type.
Object Owner	Lets you select the target database object owner.
Search Direction	Lets you select a search direction: From Beginning Down Up
Case-Sensitive Search	Select to perform the search with the same capitalization as the search string.
Find Entire String Only	Select to perform the search using the entire search string, not partial strings.
Search My Objects Only	Select to perform the search only on your database objects.

**CAUTION:** The section below refers only to Microsoft SQL Server connectivity.

Microsoft SQL Server uses ODBC to connect to Microsoft SQL Servers. Rapid SQL requires native connectivity. To connect through Rapid SQL, register your Microsoft SQL Server(s) using native connectivity in the Microsoft SQL Server Client Utility.

To change your ODBC servers to native connectivity, do the following:

- 1 Open the Microsoft SQL Server Client Network Utility dialog box, CLICONFG.exe.
- 2 On the General Tab, click Add.

Microsoft SQL Server opens the Add Network Library Configuration dialog box.

The table below describes the options and functionality on the Add Network Library Configuration dialog box:

Option	Description
Server alias	In the box, type the unique name of the server.
Network libraries	In the box, click the appropriate option button to specify the network library that connects to the server.
Computer name	In the box, type the name of the target computer.

Option	Description
Port number	In the box, type the port number of the target computer.

3 Click OK.

Rapid SQL returns to the Client Network Utility dialog box.

4 In the Client Network Utility dialog box, click Apply.

Rapid SQL adds the server.

5 Open Rapid SQL and the [Datasource Registration Wizard](#).

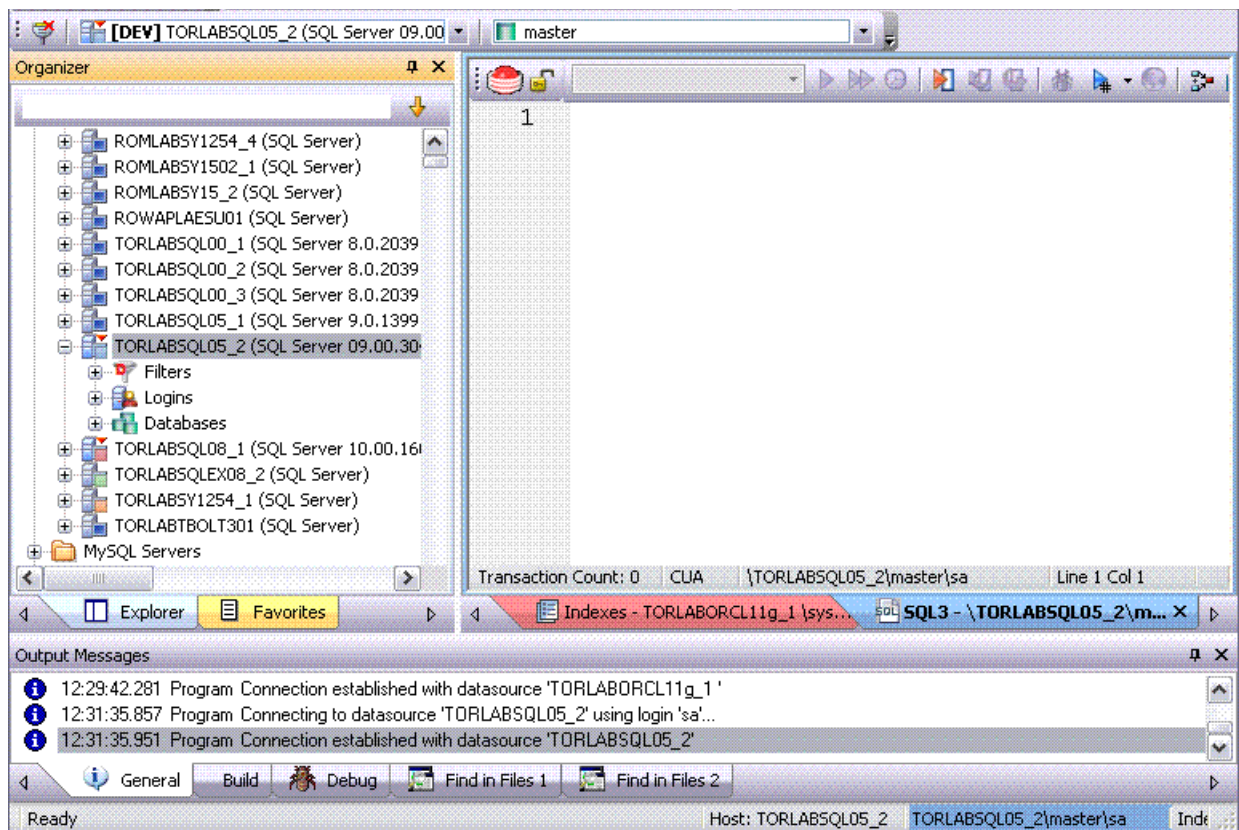
For more information, see [Datasources](#).

## CATEGORIZING DATASOURCES USING COLOR-CODING AND LABELLING





Rapid SQL provides a means to visually distinguish between datasources used for different purposes in your organization. You can assign a category, such as development or production, when creating or editing a datasource. This color-codes or labels particular user interface elements associated with that datasource.

When you assign a category to a datasource, the Datasource Explorer icons for the datasource shows a distinctive color-coded category indicator. Optionally, you can customize your categorization scheme to include the following user interface elements:

- A short name label included in the Datasource toolbar combo/dropdown when a categorized datasource is selected
- Color-coded tabs on windows associated with actions against a categorized datasource, editor or ISQL windows for example
- Color-coding of the status bar at the bottom of the Rapid SQL window



The default categories [and short name labels] and the associated colors are as follows:

Default Category	Short Name Label	Color
Production	PROD	
Development	DEV	
QA	QA	
Test	TEST	

In addition to the defaults, you can create custom category/color code/label category combinations.

To work with color-coded datasource resources, you must be familiar with the following tasks:

- Enabling the specific user interface items that can be color-coded or labelled. For details, see [Datasource Options](#).
- Customizing your category scheme. For details, see [Customizing Datasource Categories](#).
- Assigning a category to a datasource when creating or editing it. For details, see [Datasource Registration](#).



## SUPPORTED OBJECTS

Rapid SQL lets you manage database objects across different database platforms. The table below indicates the objects that Rapid SQL supports by platform:

Objects	DB2 LUW	DB2 z/OS	InterBase Firebird	Oracle	SQL Server	Sybase
<a href="#">Aliases</a>	X	X			X	X
<a href="#">Blob Filters</a>			X			
<a href="#">Check Constraints</a>	X	X	X	X	X	X
<a href="#">Clusters</a>				X		
<a href="#">Database Links</a>				X		X
<a href="#">Databases</a>					X	X
<a href="#">DBRM</a>		X				
<a href="#">Defaults</a>					X	X
<a href="#">Domains</a>			X			
<a href="#">Encryption Keys</a>			X			
<a href="#">Exceptions</a>			X			
<a href="#">External Functions</a>			X			
<a href="#">Extended Procedures</a>						X
<a href="#">Foreign Keys</a>	X	X	X	X	X	X
<a href="#">Functions</a>	X	X		X	X	X
<a href="#">Generators</a>			X			
<a href="#">Indexes</a>	X	X	X	X	X	X
<a href="#">Java Classes</a>				X		
<a href="#">Java Resources</a>				X		
<a href="#">Java Sources</a>				X		
<a href="#">Libraries</a>				X		
<a href="#">Logins</a>					X	
Materialized Query Tables	X					
<a href="#">Materialized Views</a>				X		
<a href="#">Materialized View Logs</a>				X		
<a href="#">Outlines</a>				X		
<a href="#">Package Bodies</a>				X		
<a href="#">Packages</a>	X	X		X		
<a href="#">Plans</a>		X				

Objects	DB2 LUW	DB2 z/OS	InterBase Firebird	Oracle	SQL Server	Sybase
<a href="#">Primary Keys</a>	X	X	X	X	X	X
<a href="#">Procedures</a>	X	X	X	X	X	X
<a href="#">Profiles</a>				X		
<a href="#">Rollback Segments</a>				X		
<a href="#">Rules</a>					X	X
<a href="#">Segments</a>					X	X
<a href="#">Sequences</a>				X		
<a href="#">Shadows</a>			X			
<a href="#">Snapshots</a>				X		
<a href="#">Snapshot Logs</a>				X		
<a href="#">Stogroups</a>		X				
<a href="#">Structured Types</a>	X	X				
<a href="#">Summary Tables</a>	X					
<a href="#">Synonyms</a>				X		
<a href="#">System Indexes</a>	X	X		X	X	X
<a href="#">System Tables</a>	X	X		X	X	X
<a href="#">Tables</a>	X	X	X	X	X	X
<a href="#">Triggers</a>	X	X	X	X	X	X
<a href="#">Type Bodies</a>				X		
<a href="#">Types</a>				X		
<a href="#">Unique Keys</a>	X	X	X	X	X	X
<a href="#">User Datatypes</a>	X	X			X	X
<a href="#">Users</a>	X	X	X	X	X	X
<a href="#">Views</a>	X	X	X	X	X	X

## ALIASES

**NOTE:** This object is supported by IBM DB2 LUW for Linux, Unix, and Windows, IBM DB2 LUW for OS/390 and z/OS, Microsoft SQL Server, and Sybase ASE.

Aliases let you assume the permissions of another database user without creating a separate user identity. You can use an alias when a users requires only temporary access to a database. You can also use an alias to mask a user's identity.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Drop](#)

- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)
- [Transfer Ownership](#)

## BLOB FILTERS

Rapid SQL provides support for blob filters, user-written programs that convert data stored in Blob columns from one subtype to another.

**NOTE:** This object is supported by InterBase/Firebird only.

### Creating and editing

See [Blob Filters Wizard \(InterBase/Firebird\)](#) and [Blob Filters Editor \(InterBase/Firebird\)](#).

### Other functionality/object actions supported

[Drop](#)                      [Extract](#)                      [Rename](#)                      [Report](#)

## CHECK CONSTRAINTS

**NOTE:** This object is supported on DB2 LUW, DB2 z/OS, InterBase/Firebird, Oracle, SQL Server, and Sybase.

Check constraints are data values that are acceptable in a column. They are logical expressions that verify column values meet defined acceptance criteria.

### Available Functionality

Rapid SQL offers the following functionality for this object:

**NOTE:** Functionalities differ by platform.

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)

- [Report](#)
- [Status](#)
- [Transfer Ownership](#)

## CLUSTERS

**NOTE:** This object is supported by Oracle only.

Clusters provide an optional method of storing table data. A cluster comprises of a group of tables that share the same data blocks, and which are grouped together because they share common columns and are often used together. The related columns of tables stored in a cluster are known as the cluster key.

There are two types of clusters:

- Index
- Hash

Index clusters store the cluster data together and index the cluster key, which should make them faster at retrieving a range of data rows.

Hash clusters apply hashing functions to the cluster key to determine the physical location of a data row, which should make them faster at retrieving specific data rows.

**NOTE:** To place a table on a cluster, include the ON CLUSTER syntax within the CREATE TABLE statement. Placing a table on a cluster precludes you from placing it on a tablespace or defining the associated storage parameters.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Allocate Extent](#)
- [Analyze](#)
- [Deallocate Unused Space](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)
- [Truncate](#)

## DATABASE LINKS

**NOTE:** This object is supported by Oracle and Sybase only.

Database links are named schema objects that describe a path from one database to another. Database links are implicitly used when a reference is made to a global object name in a distributed database. To use a database link, either it is public or you own it.

**NOTE:** Oracle syntax does not let you alter an existing database link. To change its definition, drop and re-create it.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)
- [Report](#)

## DATABASES

**NOTE:** Databases are available for [Microsoft SQL Server](#) and [Sybase ASE](#). For Oracle and [IBM DB2 LUW for Linux, Unix, and Windows](#), [IBM DB2 LUW for OS/390 and z/OS](#) databases are called Instances.

Databases are a collection of tables, or a collection of index spaces and tablespaces. The goals of a database system are straightforward but challenging. In general, a database aims to manage large amounts of data in a multi-user environment. It should achieve high performance while letting many users access the same information concurrently without compromising data integrity. A database also must protect against unauthorized access and provide reliable solutions for failure recovery.

For more information, see [Available Functionality](#).

### IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS INSTANCES

Databases are a collection of tables, or a collection of index spaces and tablespaces. The goals of a database system are straightforward but challenging. In general, a database aims to manage large amounts of data in a multi-user environment. It should achieve high performance while letting many users access the same information concurrently without compromising data integrity. A database also must protect against unauthorized access and provide reliable solutions for failure recovery.

For more information, see [Available Functionality](#).

## MICROSOFT SQL SERVER DATABASES

Databases are a collection of tables, or a collection of index spaces and tablespaces. The goals of a database system are straightforward but challenging. In general, a database aims to manage large amounts of data in a multi-user environment. It should achieve high performance while letting many users access the same information concurrently without compromising data integrity. A database also must protect against unauthorized access and provide reliable solutions for failure recovery.

For more information, see [Available Functionality](#).

## IBM DB2 LUW FOR OS/390 AND Z/OS INSTANCES

Databases are a collection of tables, or a collection of index spaces and tablespaces. The goals of a database system are straightforward but challenging. In general, a database aims to manage large amounts of data in a multi-user environment. It should achieve high performance while letting many users access the same information concurrently without compromising data integrity. A database also must protect against unauthorized access and provide reliable solutions for failure recovery.

For more information, see [Available Functionality](#).

## SYBASE ASE DATABASES

Databases are a collection of tables, or a collection of index spaces and tablespaces. The goals of a database system are straightforward but challenging. In general, a database aims to manage large amounts of data in a multi-user environment. It should achieve high performance while letting many users access the same information concurrently without compromising data integrity. A database also must protect against unauthorized access and provide reliable solutions for failure recovery.

### Available Functionality for Databases

Rapid SQL offers the following functionality for databases:

**NOTE:** Functionalities differ by platform.

- [Checkpoint](#)
- [Coalesce](#)
- [DBCC](#)
- [Detach/Attach](#)
- [Drop](#)

- [Extract](#)
- [Move Log](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)
- [Report](#)
- [Set Online/Offline](#)
- [Shrink](#)
- [Update Statistics](#)

## DBRM

**NOTE:** This object is supported by IBM DB2 LUW for OS/390 and z/OS only.

A Database Request Module (DBRM) is an application containing information on SQL statements extracted from the DB2 precompiler. A DBRM refers to a specific database and is stored outside of the DB2 directory.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)

## DEFAULTS

**NOTE:** This object is supported by Microsoft SQL Server and Sybase only.

Defaults promote data integrity by supplying a default value to a table column if the user does not explicitly provide one. They are reusable objects that you can bind to table columns or user datatypes.

## Available Functionality

Rapid SQL offers the following functionality for this object:

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)
- [Report](#)

## DOMAINS

Rapid SQL provides support for domains, letting you work with these global column/datatype-based objects.

**NOTE:** This object is supported by InterBase/Firebird only.

### Creating and editing

See [Domains Wizard \(InterBase/Firebird\)](#) and [Domains Editor \(InterBase/Firebird\)](#).

### Other functionality/object actions supported

[Drop](#)

[Extract](#)

[Rename](#)

[Report](#)

## ENCRYPTION KEYS

Rapid SQL provides support for encryption keys, used in encrypting and decrypting tables and columns.

**NOTE:** This object is supported by InterBase/Firebird only.

### Creating and editing

See [Encryption Keys wizard \(InterBase/Firebird\)](#) and [Encryption Keys editor \(InterBase/Firebird\)](#).



## Other functionality/object actions supported

[Drop](#)

[Extract](#)

[Rename](#)

[Report](#)

## EXCEPTIONS

Rapid SQL provides support for exceptions, letting you work with these named error messages.

**NOTE:** This object is supported by InterBase/Firebird only.

## Creating and editing

See [Exceptions Wizard \(InterBase/Firebird\)](#) and [Exceptions Editor \(InterBase/Firebird\)](#).

## Other functionality/object actions supported

[Drop](#)

[Extract](#)

[Rename](#)

[Report](#)

## EXTENDED PROCEDURES

**NOTE:** This object is supported by Sybase ASE only.

Extended Procedures are dynamic link libraries that can be used to load and execute application routines written in other programming languages, such as C or Visual Basic. Extended Procedures function and appear in the same manner as normal stored procedures in that you can pass parameters to them and obtain results.

**NOTE:** Extended Procedures can only be accessed on the Master database.

## Available Functionality

Rapid SQL offers the following functionality for this object:

- [Dependencies](#)
- [Drop](#)
- [Execute](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)
- [Report](#)

## EXTERNAL FUNCTIONS

Rapid SQL provides support for external functions, letting you work with declarations for existing functions.

**NOTE:** This object is supported by InterBase/Firebird only.

### Creating and editing

See [External Functions Wizard \(InterBase/Firebird\)](#) and [External Functions editor \(InterBase/Firebird\)](#).

### Other functionality/object actions supported

[Rename](#)

## FOREIGN KEYS

**NOTE:** This object is supported on DB2 LUW, DB2 z/OS, InterBase/Firebird, Oracle, SQL Server, and Sybase.

Foreign keys enforce referential integrity between tables by verifying the existence of foreign key values in the parent table before letting you insert or update foreign key values in the child table.

### Available Functionality

Rapid SQL offers the following functionality for this object:

**NOTE:** Functionality differs by platform.

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)
- [Report](#)
- [Status](#)
- [Transfer Ownership](#)

## FUNCTIONS

**NOTE:** This object is supported by all platforms.

Functions are subroutines that you define. Functions are useful for reusable application logic. You can use functions to determine the best methods for controlling access and manipulation of the underlying data contained in an object.

The table below describes the types of user-defined functions that Rapid SQL lets you create:

Function	Description
Column or External Table Function	You can write in a host programming language, such as C. This function can act on a table and returns a table value rather than a scalar value.
External Scalar Function	You can write in a language other than SQL, such as C++ or Java and returns a scalar value to the program. This type of function is referenced by the CREATE FUNCTION statement and can be used to perform computations on data contained in the database but cannot directly reference the data.
OleDb Function	Accesses OLE DB data in user-defined OLE DB external tables.
Sourced Function	Inherits the semantics of another function and can be an operator.
Template Function	Partial functions that do not contain any executable code. Mainly used in a federated database to map the template function to a data source function -Oracle, SQL Server, Sybase, etc. A function mapping needs to be created in conjunction with the template function.

### Available Functionality

Rapid SQL offers the following functionality for this object:

**NOTE:** Functionalities differ by platform.

- [Create Synonym](#)
- [Dependencies](#)
- [Drop](#)
- [Extract](#)
- [Generate Anonymous Block](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)
- [Transfer Ownership](#)

## GENERATORS

Rapid SQL provides support for generators, letting you work with these objects used to generate sequential numbers in a column.

**NOTE:** This object is supported by InterBase/Firebird only.

### Creating and editing

See [Generators Wizard \(InterBase/Firebird\)](#) and [Generators editor \(InterBase/Firebird\)](#).

### Other functionality/object actions supported

[Drop](#)

[Extract](#)

[Rename](#)

[Report](#)

## INDEXES

**NOTE:** This object is supported on DB2 LUW, DB2 z/OS, InterBase/Firebird, Oracle, SQL Server, and Sybase.

Indexes are optional structures associated with tables. You can create indexes specifically to speed SQL statement execution on a table. When properly used, Indexes are the primary means of reducing disk I/O. Indexes are logically and physically independent of the data in the associated table. Unique Indexes guarantee that no two rows of a table have duplicate values in the columns that define the index.

For more information, see:

[IBM DB2 LUW for Linux, Unix, and Windows Indexes](#)

[IBM DB2 LUW for OS/390 and z/OS Indexes](#)

[Microsoft SQL Server Indexes](#)

[Oracle Indexes](#)

[Sybase ASE](#)

[Available Functionality](#)

## IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS INDEXES

IBM DB2 LUW for Linux, Unix, and Windows offers two types of indexes:

- Unique
- Non-Unique

Unique Indexes guarantee that no two rows of a table have duplicate values in the columns that define the index.

For more information, see [Available Functionality](#).

## MICROSOFT SQL SERVER INDEXES

Microsoft SQL Server offers two types of indexes: clustered and non-clustered. Clustered indexes physically sort table data to match their logical order. Non-clustered indexes only order the table data logically. In a database, an index lets you speed queries by setting pointers that allow you to retrieve table data without scanning the entire table. An index can be unique or non-unique.

Microsoft SQL Server creates indexes as B-Trees, which are a series of pointers mapping index pages to their underlying data pages. As tables and, therefore, indexes grow, the number of levels in the B-Tree increases. The B-Tree of a clustered index is shorter than that of a non-clustered index because the leaf level of a clustered index is the data page.

A sound indexing strategy is critical to overall system performance. One pitfall to avoid is placing many indexes on a table without regard for their cumulative cost. Remember that indexes improve read but slow write performance because Microsoft SQL Server must update more information in the system catalog. Consequently, extra indexes can actually slow overall performance if data modification occurs frequently on the table. To determine the efficacy of indexes, you should tune your queries using SHOWPLAN and IO STATISTICS and analyze the selectivity of indexes using DBCC SHOW\_STATISTICS.

Rapid SQL separates system indexes from user-defined indexes in the Datasource Explorer to ensure that system indexes are not accidentally altered or dropped.

For more information, see [Available Functionality](#).

## ORACLE INDEXES

Oracle offers two types of indexes. The table below describes these indexes:

Index	Description
Table	A table index is defined on an individual table.
Cluster	A cluster index is defined on a set of tables physically stored together in a cluster. In an Oracle database, both table and cluster indexes use a B-tree structure.

The indexing strategy, particularly with large, active tables, is critical to overall system performance. The optimal definition and number of indexes for a given table is determined by the mix of access paths to that table performing insert, update, delete and select operations. For example, adding or changing an index can speed up your selects but slow your inserts, updates and deletes. Careful tuning and testing helps you achieve the best overall performance.

**TIP:** Indexes generally improve read operations in a database, but you should not place too many indexes on some tables. Since Oracle must maintain each index along with its referenced table, placing too many indexes on a table that is the object of much insert, update, and delete activity, can actually degrade performance.

Even when an index exists on a table, the way a SQL statement is coded can actually disallow the use of the index. To prevent this from happening, follow these rules of thumb:

- Try not to use SQL statements that include the NOT IN, NOT LIKE, <>, IS NULL operators because they typically suppress the use of indexes.
- When referencing concatenated indexes with queries, be sure the leading column in the index is used. If it isn't, the index won't be used at all.
- Avoid using functions in WHERE predicates.

If you must use functions, and you are using Oracle8i, investigate the use of function-based indexes.

For more information, see [Available Functionality](#).

## INDEX PARTITIONS

Index partitions are similar to table partitions. There are [three](#) types of partitioned indexes that Oracle supports:

- 1 [Local prefixed](#)
- 2 [Local nonprefixed](#)
- 3 [Global prefixed](#)

**NOTE:** An index cannot be partitioned if it is a cluster index or if the index is defined on a clustered table.

### Local prefixed and nonprefixed indexes

A local partitioned index has keys that refer to rows in a single table partition. A local partitioned index is automatically partitioned to mirror the underlying table. The number of partitions or subpartitions and the partition bounds for the partitioned index correspond with the partitions on the table. Oracle maintains this correspondence. If the table partitions are altered, the index partitions are altered accordingly.

A local partitioned index is prefixed if it is partitioned on the same column as the underlying table. The local partitioned index is nonprefixed if it is partitioned on a different column.

### Global prefixed indexes

A global partitioned index can refer to rows in more than one table partition or subpartition. Global partitioned indexes are more difficult to manage than local partitioned indexes because any change in the underlying table partition affects all partitions in a global index. As a result, there is increased partition maintenance.

**NOTE:** A global index can only be range partitioned but it can be defined on any kind of partitioned table.

## IBM DB2 LUW FOR OS/390 AND Z/OS INDEXES

IBM DB2 LUW for OS/390 and z/OS offers two types of indexes:

- Unique
- Non-Unique

Unique Indexes guarantee that no two rows of a table have duplicate values in the columns that define the index.

Non-Unique indexes let table rows have duplicate values in the columns that define the indexes.

For more information, see [Available Functionality](#).

## SYBASE ASE INDEXES

Sybase ASE offers two types of indexes: clustered and non-clustered. Clustered indexes physically sort table data to match their logical order. Non-clustered indexes only order the table data logically. In a database, an index lets you speed queries by setting pointers that let you retrieve table data without scanning the entire table. An index can be unique or non-unique.

Sybase ASE creates indexes as B-Trees, which are a series of pointers mapping index pages to their underlying data pages. As tables and, therefore, indexes grow, the number of levels in the B-Tree increases. The B-Tree of a clustered index is shorter than that of a non-clustered index because the leaf level of a clustered index is the data page.

A sound indexing strategy is critical to overall system performance. One pitfall to avoid is placing many indexes on a table without regard for their cumulative cost. Remember that indexes improve read but slow write performance because Sybase ASE must update more information in the system catalog. Consequently, extra indexes can actually slow overall performance if data modification occurs frequently on the table. To determine the efficacy of indexes, you should tune your queries using SHOWPLAN and IO STATISTICS and analyze the selectivity of indexes using DBCC SHOW\_STATISTICS.

## AVAILABLE FUNCTIONALITY FOR INDEXES

Rapid SQL offers the following functionality for indexes:

**NOTE:** Functionality differs by platform.

- [Analyze](#)
- [Allocate Extent](#)
- [Convert to Partitioned](#)
- [Deallocate Unused Space](#)
- [Disable Index](#)
- [Drop](#)

- [DBCC](#)
- [Estimate Size](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Place](#)
- [Rebuild Index](#)
- [Rename](#)
- [Reorganize](#)
- [Report](#)
- [Set Statistics](#)
- [Transfer Ownership](#)
- [Update Statistics](#)

## INSTANCE

**NOTE:** This object is support for IBM DB2 and Oracle.

Rapid SQL places Instance as the first level of information under the Datasource node in the Database Explorer. Instance includes:

- DB Manager Configuration
- Datasources

### Available Functionality

Rapid SQL offers the following functionality for this object:

- **New LUW Database** - for details, see [Database Wizard \(IBM DB2 LUW\)](#)

## JAVA CLASSES

**NOTE:** This object is supported by Oracle 8i or later only.



The Java Classes contain compiled Java code. Java Classes are made up of a group of data items, with associated functions that perform operations. The data items are called fields or variables; the functions are referred to as methods.

**TIP:** Oracle is shipped with a JVM(Java Virtual Machine). The JVM provided by Oracle sits atop the Oracle RDBMS and interacts directly with the RDBMS instead of the operating system.

### Available Functionality

- [Drop](#)

## JAVA RESOURCES

**NOTE:** This object is supported by Oracle only.

The Java Resources node of the *Explorer* tab offers support for browsing Java resources.

### Available Functionality

- [Drop](#)

## JAVA SOURCES

**NOTE:** This object is supported by Oracle 8i or later only.

Java Sources contain the uncompiled Java source code.

**TIP:** Oracle is shipped with a JVM(Java Virtual Machine). The JVM provided by Oracle sits atop the Oracle RDBMS and interacts directly with the RDBMS instead of the operating system.

### Available Functionality

- [Compile](#)
- [Create or Edit Java Source](#)
- [Load Java](#)
- [Drop](#)

## JOB QUEUES

**NOTE:** This object is supported by Oracle only.

Job Queues are built-in mechanisms that let you schedule a variety of SQL-based or command-line driven tasks.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Run](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Disable/Enable Job Queues](#)

## LIBRARIES

**NOTE:** This object is supported by Oracle only.

Libraries are an object type introduced in Oracle8 that represent a call to an operating system shared library. After the call is made, libraries can be used by SQL or PL/SQL to link to external procedures or functions. Libraries are only to be used on operating systems that support shared libraries and dynamic linking. Libraries serve as pointers or aliases to physical operating system shared library files and do not have existence as a physical object on their own, rather they rely on the physical existence of the files in the external operating system library to which they refer. To access the function or procedures stored in the library, you need execute privileges at the operating system level where the shared library resides.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)
- [Report](#)

## LOGINS

**NOTE:** This object is supported by Microsoft SQL Server and Sybase only.

Logins let you access your account. Your login account controls access to the server and all of the databases within it. Only the System Administrator or System Security Officer can create logins. Once you can log into a server, you need additional privileges to access user databases. Specifically, each database owner adds the login as a user or alias to the database.

**NOTE:** Logins are available for [Microsoft SQL Server](#) and [Sybase ASE](#).

## MICROSOFT SQL SERVER LOGINS

Logins let you access your account. Your login account controls access to the server and all of the databases within it. Only the System Administrator or System Security Officer can create logins. Once you can log into a server, you need additional privileges to access user databases. Specifically, each database owner adds the login as a user or alias to the database.

For more information, see [Available Functionality](#).

## SYBASE ASE LOGINS

Logins let you access your account. Your login account controls access to the server and all of the databases within it. Only the System Administrator or System Security Officer can create logins. Once you can log into a server, you need additional privileges to access user databases. Specifically, each database owner adds the login as a user or alias to the database.

For more information, see [Available Functionality](#).

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Change Password](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)

## MATERIALIZED QUERY TABLES

**NOTE:** This object is supported by IBM DB2 LUW for Linux, Unix, and Windows version 8.

A materialized query table is a table whose definition is based on the result of a query. The materialized query table typically contains pre-computed results based on the data existing in the table or tables that its definition is based on. If the SQL compiler determines that a query will run more efficiently against a materialized query table than the base table or tables, the query quickly executes against the materialized query table.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Create Insert Statements](#)
- [Create Synonym](#)
- [Drop](#)
- [Edit Data](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Refresh Summary Table](#)
- [Rename](#)
- [Report](#)
- [Transfer Ownership](#)

## MATERIALIZED VIEWS

**NOTE:** This object is supported by Oracle only.

Materialized views are used to dynamically copy data between distributed databases. There are two types of materialized views:

- Complex
- Simple

Complex materialized views copy part of a master table or data from more than one master table. Simple materialized views directly copy a single table. You cannot directly update the underlying data contained in materialized views.

**NOTE:** Materialized Views are only in Oracle 8.1 or later. If you are using Oracle 8.0 or earlier, see [Snapshots](#).

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Create Synonym](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)
- [Report](#)

## MATERIALIZED VIEW LOGS

**NOTE:** This object is supported by Oracle only.

Materialized View logs are tables that maintain a history of modifications to the master table, and they are used to refresh simple materialized views. When you create a materialized view log, Oracle automatically creates a log table to track data changes in the master table and a log trigger to maintain the data in the log table.

**NOTE:** Materialized View Logs are in Oracle 8i or later. If you are using Oracle 8.0 or earlier, see [Snapshot Logs](#).

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)

## ORACLE JOB QUEUES

**NOTE:** This object is supported by Oracle only.

Oracle job queues are built-in mechanisms in the Oracle database that let you schedule a variety of SQL-based or command-line driven tasks. Basically, they are database-driven schedulers.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)

## OUTLINES

**NOTE:** This object is supported by Oracle 8.1 or later only.

Outlines are a set of results for the execution plan generation of a particular SQL statement. When you create an outline, plan stability examines the optimization results using the same data used to generate the execution plan. That is, Oracle uses the input to the execution plan to generate an outline, and not the execution plan itself.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Change Category](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Reassign by Category](#)
- [Rebuild Outlines](#)
- [Rename](#)
- [Report](#)

## PACKAGES

**NOTE:** This object is supported by IBM DB2 LUW for Linux, Unix, and Windows, IBM DB2 for OS/390 and z/OS, and Oracle only.

Packages contain all the information needed to process SQL statements from a single source file. You can use packages to process and call batches of SQL. Depending on the platform, packages can include:

- [Procedures](#)

- [Functions](#)
- [Types](#)
- Variables
- Constants
- Exceptions
- Cursors
- Subprograms

Packages offer a number of important advantages over using standalone procedures and functions, including the ability to:

- Modify package objects without recompiling dependent database objects.
- Declare global variables and cursors that can be shared within the package.
- Grant privileges more efficiently.
- Load multiple package objects into memory at once.

Packages usually have two parts: a header or specification and a body, although sometimes the body is unnecessary. The package header declares the members of the package while the body details the logic underlying each of the package components.

### Available Functionality

Rapid SQL offers the following functionality for this object:

**NOTE:** Functionalities differ by platform.

- [Bind Package](#)
- [Create Synonym](#)
- [Dependencies](#)
- [Drop](#)
- [Explain Plan](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rebind Packages](#)
- [Report](#)

## PACKAGE BODIES

**NOTE:** This object is supported by Oracle only.

Package Bodies implement the package specification in that the package body includes the definition of every cursor and subprogram declared in the package specification. While Package Bodies are listed as a separate object in the Datasource Explorer, they are created on the Packages Editor in conjunction with Packages. For details, see [Packages Editor \(Oracle\)](#)

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Create Synonym](#)
- [Dependencies](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)

## PLANS

**NOTE:** This object is supported by IBM DB2 LUW for OS/390 and z/OS only.

A Plan is an executable application created in the bind process. It can include one or more packages or debris.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Bind Plan](#)
- [Copy Object Names](#)
- [Drop](#)
- [Free Plans](#)
- **Open** - see [Modifying objects using editors](#)
- [Rebind Plans](#)
- [Report](#)



## PRIMARY KEYS

**NOTE:** This object is supported on DB2 LUW, DB2 z/OS, InterBase/Firebird, Oracle, SQL Server, and Sybase.

Primary Keys are a set of table columns that can uniquely identify every row of a table.

### Available Functionality

Rapid SQL offers the following functionality for this object:

**NOTE:** Functionalities differ by platform.

- [Disable Index](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rebuild Index](#)
- [Rename](#)
- [Reorganize](#)
- [Report](#)
- [Status](#)
- [Transfer Ownership](#)

## PROCEDURES

**NOTE:** This object is supported on DB2 LUW, DB2 z/OS, InterBase/Firebird, Oracle, SQL Server, and Sybase.

Procedures are a reusable block of PL/SQL, stored in the database, that applications can call. Procedures streamline code development, debugging and maintenance by being reusable. Procedures enhance database security by letting you write procedures granting users execution privileges to tables rather than letting them access tables directly.

**NOTE:** Only IBM DB2 LUW for OS/390 and z/OS SQL stored procedures created by Rapid SQL, DBArtisan, or IBM's Stored Procedure Builder can be retrieved by Rapid SQL or DBArtisan.

## Sybase ASE

Procedures perform procedural logic in your Sybase ASE applications. They are batches of SQL statements that are compiled and stored in the system catalog. Procedures execute faster than embedded SQL statements because they are pre-compiled and have execution plans for use by the optimizer. When you create a procedure, Sybase ASE builds a query tree, which it stores in a system table. When you execute a procedure for the first time, Sybase ASE loads it from the system table, compiles, and optimizes it. Sybase ASE places the resulting query plan in the procedure cache where it remains on a most recently used basis. In addition to better performance, procedures yield other benefits, including easier code maintenance, additional security and reduced network traffic.

In addition, Rapid SQL now supports [extended procedures](#). Extended procedures are dynamic link libraries that can be used to load and execute application routines written in other programming languages, such as C or Visual Basic. Extended procedures function and appear in the same manner as normal procedures in that you can pass parameters to them and obtain results.

## Available Functionality

Rapid SQL offers the following functionality for this object:

**NOTE:** Functionalities differ by platform.

- [Create Synonym](#)
- [Drop](#)
- [Dependencies](#)
- [Execute](#)
- [Extract](#)
- [Generate Anonymous Block](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)
- [Report](#)
- [Transfer Ownership](#)

## PROFILES

**NOTE:** This object is supported by Oracle only.

Profiles are a mechanism for allocating system and database resources to users. Profiles let you specify limits on:

- Number of sessions
- CPU time
- Connect time
- Idle time
- Logical reads and space in the SGA's shared pool

You can assign profiles to one or more users. The database's default profile and all of its resource limits are assigned to users without a specific profile assignment.

### **Available Functionality**

Rapid SQL offers the following functionality for this object:

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)

## ROLLBACK SEGMENTS

Microsoft SQL Server 7	
Column Name	Description
Remote Server Name	Name of the server
Net Name	Reserved (name of the server)
Product Name	Product name for the remote server(e.g. Microsoft SQL Server)
Ole Provider Name	Name of the Ole provider accessing the server

Sybase ASE	
Column Name	Description
Remote Server Name	Name of the server
Network Name	Name of the network on which the Remote Server resides
Timeouts	Shows whether or not the Timeout option has been enabled. If it is set to 1 then the site handler for the remote login will timeout if there is no remote user activity for one minute.

Column Name	Description
Role Name	Name of the role
Role Type	Whether the role is a standard Microsoft SQL Server role or an application role
Created	When the role was created
Updated	When the role was last updated

Column Name	Description
Role Name	Name of the role
Password Required	Password authentication for the role. The valid values are Yes (specific password required), No (no password required), External (password is authenticated through the operating system) or Global (password is authenticated through the security domain central authority).

Column Name	Description
Role Name	Name of the role

**NOTE:** This object is supported by Oracle only.

Rollback segments manage all transactions in your Oracle databases. They maintain read consistency among concurrent users in a database and can rollback transactions. They are transaction logs that maintain discrete records of changes to data. By maintaining a history of data changes, rollback segments can rollback uncommitted transactions so that data is rolled back to their prior state.

**TIP:** Databases with large user bases and intensive data manipulation characteristics can benefit from having multiple rollback tablespaces that are placed on different server drives or file systems. This strategy aids in the reduction of I/O and rollback segment contention.

Because all changes are logged to a rollback segment, they can become a performance bottleneck as the number of concurrent users and transaction activity increases. You can tune rollback segment performance by creating additional ones to handle increased traffic. To determine whether rollback segment performance is satisfactory, you should monitor the ratio of gets to waits. If you see that the number of waits is climbing, create another rollback segment to further spread concurrency.

**NOTE:** The analysis for determining when to add a rollback segment can be quite complicated. As a rule-of-thumb, however, many DBAs simply plan on allocating a rollback segment for every four concurrent users.

The proper sizing of rollback segments is critical to their overall performance. Performance degrades whenever a rollback segment must extend, wrap or shrink in response to transaction loads. Ideally, you want to make their extents as small as possible while ensuring that each transaction fits into a single extent, which in practice is very difficult to achieve.

**NOTE:** You should carefully monitor each rollback segments for the occurrence of extends, wraps and shrinks. If you see a high level of extends and wraps, then the rollback segment is probably too small. If you see a high level of shrinks, then you should increase the rollback segment's optimal size.

## Available Functionality

Rapid SQL offers the following functionality for this object:

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)
- [Shrink](#)

## RULES

**NOTE:** This object is supported by Microsoft SQL Server and Sybase only.

Rules promote data integrity by allowing you to validate the values supplied to a table column. They are reusable objects that you can bind to table columns or user datatypes. For example, you can create a rule, bind it to a column in a table and have it specify acceptable values that can be inserted into that column.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)
- [Report](#)

## SEGMENTS

**NOTE:** This object is supported by Sybase only.

Segments are a mechanism for placing tables and indexes on specific logical partitions. You create segments on one or more fragments of a database. You can map segments to specific database fragments, which in turn reside on specific hard disks; and, mapping segments lets you increase i/o throughput by placing intensively used tables and indexes on different physical devices. You can allocate tables and indexes to segments by including placement statements at the end of CREATE TABLE or CREATE INDEX statements.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)

## SEQUENCES

Sequences are programmable database objects that generate a definable sequence of values. Once defined, a sequence can be made available to many users. A sequence can be accessed and incremented by multiple users with no waiting. A sequence can be used to automatically generate primary key values for tables. When you create a sequence, you can define its initial value, increment interval and maximum value.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Create Alias](#)
- [Create Synonym](#)
- [Dependencies](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)
- [Report](#)
- [Restart](#)
- [Transfer Ownership](#)

## SHADOWS

Rapid SQL provides support for shadows, letting you create one or more in-sync copies of the database stored on secondary storage devices.

**NOTE:** This object is supported by InterBase/Firebird only.

### Creating and editing

See [Shadows wizard \(InterBase/Firebird\)](#) and [Shadows editor \(InterBase/Firebird\)](#).

## Other functionality/object actions supported

[Drop](#)[Extract](#)[Rename](#)[Report](#)

## SNAPSHOTS

**NOTE:** This object is supported by Oracle only.

Snapshots are used to dynamically copy data between distributed databases. There are two types of snapshots:

- Complex
- Simple

**NOTE:** Snapshots are available in Oracle 8.0 and earlier. For Oracle 8i or later, see [Materialized Views](#).

Complex snapshots copy part of a master table or data from more than one master table. Simple snapshots directly copy a single table. You cannot directly update the underlying data contained in snapshots.

**NOTE:** To use snapshots, use the Oracle 7 Distributed Option.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Create Synonym](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)
- [Report](#)

## SNAPSHOT LOGS

**NOTE:** This object is supported by Oracle only.



Snapshot logs are tables that maintain a history of modifications to the master and they are used to refresh simple snapshots. When you create a snapshot log, Oracle automatically creates a log table to track data changes in the master table and a log trigger to maintain the data in the log table.

**NOTE:** Snapshot Logs are available in Oracle 8.0 or earlier. For Oracle 8i or later, see [Materialized View Logs](#).

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)

## STOGROUPS

**NOTE:** This object is supported by IBM DB2 LUW for OS/390 and z/OS only.

Stogroups are storage groups, named sets of volumes on which DB2 LUW for OS/390 data is stored.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)

## STRUCTURED TYPES

**NOTE:** This object is supported by IBM DB2 LUW for Linux, Unix, and Windows and IBM DB2 LUW for OS/390 and z/OS only.

Structured types define an abstract data type or object composed of a collection of similar types of data. For example, create an structured type that defines a full address rather than the pieces of an address, such as city, state and postal code. An structured type stores the pieces of an address in a single type, storing them in the same location and allowing the full address to be accessed and manipulated as single unit rather than multiple units.

Structured types are useful for ensuring uniformity and consistency as they are defined as single encapsulated entity that can be reused in other structured types and objects. They also offer flexibility by allowing for the creation of objects that represent real-world situations which is limited in relational objects.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Dependencies](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- [Report](#)
- [Transfer Ownership](#)

## SUMMARY TABLES

**NOTE:** This object is supported by IBM DB2 LUW for Linux, Unix, and Windows version 7.

A summary table is a table whose definition is based on the result of a query. The summary table typically contains pre-computed results based on the data existing in the table or tables that its definition is based on. If the SQL compiler determines that a query will run more efficiently against a summary table than the base table or tables, the query quickly executes against the materialized query table.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Create Synonym](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)

- [Report](#)

## SYNONYMS

**NOTE:** This object is supported by IBM DB2 LUW for OS/390 and z/OS and Oracle only.

Synonyms are an alternate name for an object in the database. Depending on the platform, you can define synonyms on tables, views, sequences, procedures, functions, packages and materialized views. There are two types of synonyms:

- Public
- Private

Public synonyms can be used by everyone in the database. Private synonyms belong to the synonym owner.

Synonyms simplify object names and let you:

- Reference an object without needing to specify its owner.
- Reference a remote object without needing to specify its database.
- Alias an object so its purpose becomes more understandable.

### Available Functionality

Rapid SQL offers the following functionality for this object:

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)
- [Report](#)

## TABLES

**NOTE:** This object is supported on DB2 LUW, DB2 z/OS, InterBase/Firebird, Oracle, SQL Server, and Sybase.

Tables are the basic unit of data storage. Tables store all the data accessible to users in rows and columns. Each column has a name, datatype and other associated properties. After you define a table, users can insert valid data into the table, which you can later query, update and delete.

**NOTE:** Available Functionality

Rapid SQL offers the following functionality for this object:

**NOTE:** Functionalities differ by platform.

- [Activate Logging](#)
- [Allocate Extent](#)
- [Analyze](#)
- [Build Query](#)
- [Clone Table](#)
- [Convert to Partitioned](#)
- [Copy Object Names](#)
- [Create Alias](#)
- [Create Clone](#)
- [Create Insert Statements](#)
- [Create Like](#)
- [Create Synonym](#)
- [Create View](#)
- [DBCC](#)
- [Deallocate Unused Space](#)
- [Dependencies](#)
- [Disable/Enable Triggers](#)
- [Drop](#)
- [Drop Clone](#)
- [Edit Data](#)
- [Estimate Size](#)
- [Exchange Data With Clone](#)
- [Extract](#)

- [Generating Packages, Procedures, and Statements from Tables and Views](#)
- [Import Data](#)
- [Indexes](#)
- [Lock](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Place](#)
- [Quiesce Tablespaces](#)
- [Rename](#)
- [Reorganize](#)
- [Report](#)
- [Viewing Schema](#)
- [Select \\* From](#)
- [Transfer Ownership](#)
- [Triggers](#)
- [Truncate](#)
- [Update Statistics](#)

## TABLESPACES

**NOTE:** This object is supported by IBM DB2 LUW for OS/390 and z/OS, IBM DB2 LUW for Linux, Unix, and Windows, and Oracle only.

Tablespaces are storage structures that act as partitions for the database. You can create a tablespace to store table data and other objects related to table performance such as indexes or large object data. Tablespaces are used to manage large complex databases. Once you have created a tablespace, you can place objects on it.

**TIP:** Create separate tablespaces for your tables and indexes and put each tablespace on a different drive or file system. Segmenting tables and their corresponding indexes in this manner helps eliminate I/O contention at the server level.

**NOTE:** IBM DB2 LUW for Linux, Unix, and Windows lets you assign a location for table or index data directly to physical storage devices. Each tablespace can also be broken down into a collection of containers which are the actual physical storage files or devices. You can then spread the data or database objects across multiple file systems, thereby giving you the necessary space for objects that require it.

Once you have created a tablespace, you can place individual tables and indexes on it. Because tablespaces map to physical drives, you can improve i/o performance by placing tables and their indexes on physically separated table spaces.

### Available Functionality

Rapid SQL offers the following functionality for this object:

**NOTE:** Functionalities differ by platform.

- [Change Status \(Tablespaces\)](#)
- [Coalesce](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Quiesce Tablespaces](#) (IBM DB2 LUW for Linux, Unix, and Windows)
- [Rename](#)
- [Report](#)
- [Switch Online](#)
- [Transfer Ownership](#)

## TRIGGERS

**NOTE:** This object is supported on DB2 LUW, DB2 z/OS, InterBase/Firebird, Oracle, SQL Server, and Sybase.

Triggers are a special type of procedure that automatically fire when defined data modification operations (insert, update or delete) occur on a target table. Triggers fire after an insert, update or delete, but belong to the same transaction as the data modification operation.

Because triggers fire automatically and are part of the same transaction as the initiating SQL statement, they are often used to enforce referential integrity in a database. For example, an insert trigger can validate that a foreign key value exists in the primary key of the parent table. Similarly, you can write a trigger to cascade the deletion or update of data in a parent table to all corresponding data rows in child tables.

While triggers are a powerful data validation mechanism, they do complicate the task of debugging and tuning SQL. Because triggers fire inside transactions, they hold open transactions, contributing to concurrency problems. Accordingly, you should move all non-critical data validation outside of triggers to shorten transaction lengths to their minimum.

## Available Functionality

Rapid SQL offers the following functionality for this object:

**NOTE:** Functionalities differ by platform.

- [Change Status \(Tablespaces\)](#)
- [Dependencies](#)
- [Disable/Enable Triggers](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)
- [Report](#)
- [Transfer Ownership](#)

## TYPES

**NOTE:** This object is supported by Oracle only.

Types define an abstract data type or object composed of a collection of similar types of data. For example, create an object type that defines a full address rather than the pieces of an address, such as city, state and postal code. An object type stores the pieces of an address in a single type, storing them in the same location and allowing the full address to be accessed and manipulated as single unit rather than multiple units.

Object types are useful for ensuring uniformity and consistency as they are defined as single encapsulated entity that can be reused in other object types and objects. They also offer flexibility by allowing for the creation of objects that represent real-world situations which is limited in relational objects.

You can choose to create a type that is incomplete, complete, a VARRAY, or a nested table or any combination of the above. An incomplete type specifies no attributes and can be used for circular references such as person - female. It lets the type be referenced before it is complete. The VARRAY type can be used to store small sets of related data. For example, if you have ten offices (each one with a different description) at a particular division in your company, you could create a VARRAY of 10 to hold the details of these offices. The values for a VARRAY type must be fixed and known and small values as they are stored in RAW format. A nested table type can be used when data is repeated for the same entity an unknown number of times and storage is a concern.

## Available Functionality

Rapid SQL offers the following functionality for this object:

- [Dependencies](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)

## TYPE BODIES

**NOTE:** This object is supported by Oracle only.

Type Bodies implement object type specification by containing the definition of every cursor and subprogram declared in the object type specification. While Type Bodies are listed as a separate object in the Datasource Explorer, they are created on the Types Editor in conjunction with [Types](#). For details, see [Types Editor \(Oracle\)](#).

## Available Functionality

Rapid SQL offers the following functionality for this object:

- [Dependencies](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)

## UNIQUE KEYS

**NOTE:** This object is supported on DB2 LUW, DB2 z/OS, InterBase/Firebird, Oracle, SQL Server, and Sybase.

Unique keys can enforce logical keys that are not chosen as the primary key. They enforce uniqueness for specified columns in a table.



## Available Functionality

Rapid SQL offers the following functionality for this object:

**NOTE:** Functionalities differ by platform.

- [Disable Index](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rebuild Index](#)
- [Rename](#)
- [Reorganize](#)
- [Report](#)
- [Status](#)
- [Transfer Ownership](#)

## USER DATATYPES

**NOTE:** This object is supported by IBM DB2 LUW for Linux, Unix, and Windows, IBM DB2 LUW for OS/390 and z/OS, Microsoft SQL Server, and Sybase only.

User-defined datatypes promote domain consistency by streamlining the definition of commonly used table columns in a database. You can build a customized datatype from system datatypes and bind defaults and rules to it to enhance integrity. When you reference the user datatype in a column, the column assumes all of the properties of the user datatype.

## Available Functionality

Rapid SQL offers the following functionality for this object:

**NOTE:** Functionalities differ by platform.

- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)

- [Report](#)
- [Transfer Ownership](#)

## USERS

**NOTE:** This object is supported on DB2 LUW, DB2 z/OS, InterBase/Firebird, Oracle, SQL Server, and Sybase.

A user is an individual with access to the DBMS.

For more information, see:

- [IBM DB2 LUW for Linux, Unix, and Windows Users](#)
- [IBM DB2 LUW for OS/390 and z/OS Users](#)
- [Microsoft SQL Server Users](#)
- [Oracle Users](#)
- [Sybase ASE Users](#)
- [Available Functionality](#)

### IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS USERS

IBM DB2 LUW for Linux, Unix, and Windows relies on your machine's operating system to define users for the server. You create and maintain users through your operating system. For example, for Windows NT, you create users in the User Manager.

For more information, see [Available Functionality](#).

### IBM DB2 LUW FOR OS/390 AND Z/OS USERS

IBM DB2 LUW relies on the operating system you are running to define users for the server. This means that users are created and maintained through your operating system. For example, if you are running Windows NT, the User Manager is the utility where users should be created.

For more information, see [Available Functionality](#).

### MICROSOFT SQL SERVER USERS

Microsoft SQL Server controls access at the database level by requiring the System Administrator or Database Owner to add a login as a database user or alias. After you create a database user, you can implement further security by the granting or revoking the privileges for that user on specific database objects. To consolidate the process of granting or revoking permissions to many users, the database owner can assign users to groups.

For more information, see [Available Functionality](#).

## ORACLE USERS

To access an Oracle database, you need a user account authenticated with a password. A user account is what Oracle uses to permit access by the user. You can assign the following optional properties to the user:

- Default tablespace
- Temporary tablespace
- Quotas for allocating space in tablespaces
- Profile containing resource limits

For more information, see [Available Functionality](#).

## SYBASE ASE USERS

Sybase ASE controls access at the database level by requiring the System Administrator or Database Owner to add a login as a database user or alias. After you create a database user, you can implement further security by granting or revoking the privileges for that user on specific database objects. To consolidate the process of granting or revoking permissions to many users, the database owner can assign users to groups.

For more information, see [Available Functionality](#).

## AVAILABLE FUNCTIONALITY FOR USERS

Rapid SQL offers the following functionality for users:

**NOTE:** Functionalities differ by platform.

- [Change Password](#)
- [Drop](#)
- [Extract](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Report](#)

## VIEWS

**NOTE:** This object is supported on DB2 LUW, DB2 z/OS, InterBase/Firebird, Oracle, SQL Server, and Sybase.

Views are SQL queries stored in the system catalog that customize the display of data contained in one or more tables. Views behave like tables because you can query views and perform data manipulation operations on them. However, views do not actually store any data. Instead, they depend on data contained in their base tables. Views let you:

- View a customized selection of data from one or more tables. As a result, you can display data more cogently to different sets of users, even though the underlying data is the same.
- Restricting access to a defined set of rows and columns.

### Available Functionality

Rapid SQL offers the following functionality for this object:

**NOTE:** Functionalities differ by platform.

- [Build Query](#)
- [Copy Object Names](#)
- [Create Alias](#)
- [Create Synonym](#)
- [Dependencies](#)
- [Describe Window](#)
- [Drop](#)
- [Extract](#)
- [Generating Packages, Procedures, and Statements from Tables and Views](#)
- **New** - see [Creating objects](#)
- **Open** - see [Modifying objects using editors](#)
- [Rename](#)
- [Report](#)
- [Select \\* From](#)
- [Transfer Ownership](#)
- [Viewing Schema](#)

## MODIFYING OBJECTS USING EDITORS

An Object editor is a tabbed dialog box that stores information about existing server and database types and objects, and lets you modify those items. For an introduction to object editors, see [Overview and common usage of object editors](#).

**NOTE:** Availability of editors for specific object type varies from DBMS to DBMS. Similarly, functionality offered for an editor for an object type common to two or more DBMS, will differ from DBMS to DBMS.

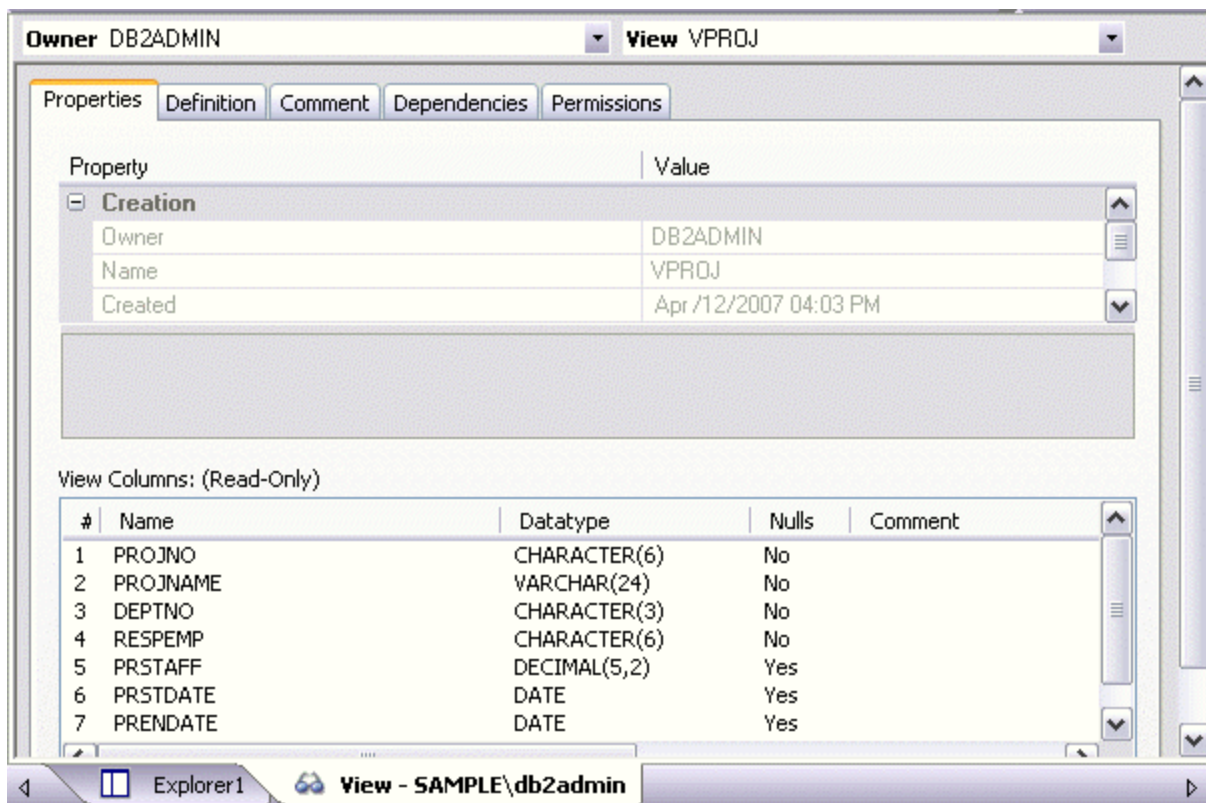
See the following topics for information the object editors available for each supported platform:

- [IBM DB2 for Linux, Unix, and Windows Object Editors](#)
- [IBM DB2 for z/OS Object Editors](#)
- [InterBase/Firebird Object Editors](#)
- [Microsoft SQL Server Object Editors](#)
- [Oracle Object Editors](#)
- [Sybase ASE Object Editors](#)

## OVERVIEW AND COMMON USAGE OF OBJECT EDITORS

An Object editor lets you view and modify settings and properties of existing object types and servers on a datasource. It also lets you add new resources and provides access to related, datasource management facilities. For example, the Tables editor lets you add or insert, edit, or drop columns, work with permissions to work with that table, access information on physical storage and the distribution of data and indexes across table spaces, and so on.

Each tab on the Object editor lets you perform a logical task or collection of logical tasks for that object type. The Object editor toolbar has a set of commands common to all object types and also includes a **Command** menu with commands specific to the object type you are currently working with. For example:



In order to work with object editors, you must be familiar with the following tasks:

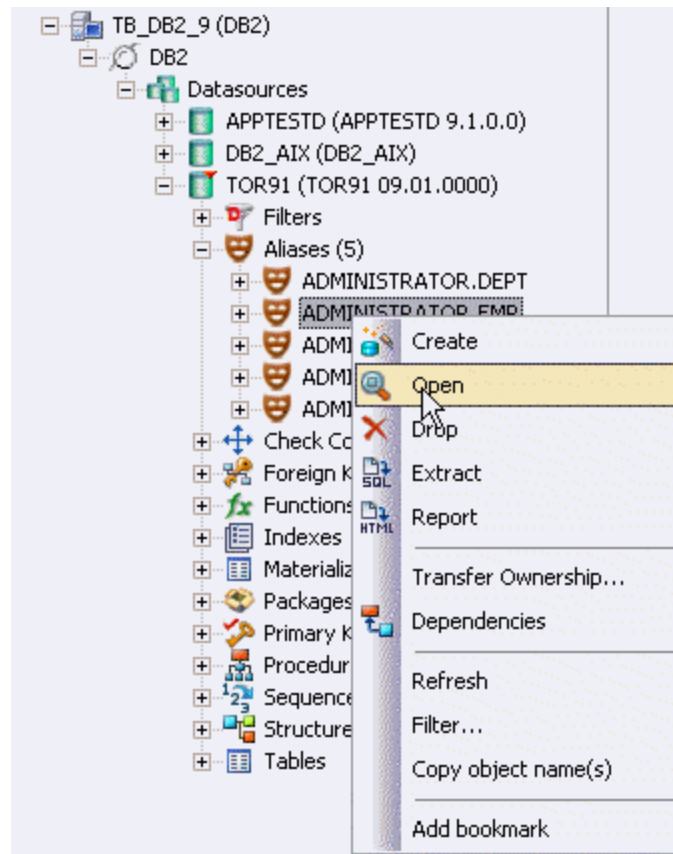
- [Opening an Object Editor](#) on a server or database object.
- [Viewing and Modifying Object Properties](#) using functionality common to all object editors as well as object-specific functionality available to specific object editors.
- [Previewing and Submitting Object Editor Changes](#) to effect your changes to the datasource.

## OPENING AN OBJECT EDITOR

Object editors are accessed from the Datasource Explorer pane of the Rapid SQL main window.

### To open an Object editor on a specific database object

- 1 Connect to the datasource that contains the object you want to modify. For more information, see [Datasource Management](#).
- 2 On the Datasource Explorer, expand the target datasource.
- 3 Continue to expand folders until objects of the target object type are displayed.
- 4 Right-click the object to be edited and select **Open** from the context menu.



Rapid SQL opens an editor for that object type on the object you selected.

For information on editing objects, see [Viewing and Modifying Object Properties](#).

## VIEWING AND MODIFYING OBJECT PROPERTIES

There are two categories of tasks you can perform using an Object editor:

- [Work with Tasks and Settings of Specific Object Types](#)
- [Work with Tasks and Settings Common to Object Editors](#)

### WORK WITH TASKS AND SETTINGS OF SPECIFIC OBJECT TYPES

Many of the tasks you perform using an Object editor are specific to the object type you are working with. For example, the Triggers editor lets you work directly with SQL code for the trigger while the Tables editor lets you work with columns of the table. Object-specific tasks are documented in the topics for specific Object editors provided later in this chapter.

### WORK WITH TASKS AND SETTINGS COMMON TO OBJECT EDITORS

The following tasks can be performed against many of the supported Object editors:

- [Using the object editor toolbar](#)
- [Opening Another Object in the Same Object Editor Explorer](#)
- [Viewing the SQL/DDDL for an Object](#)
- [Working with Privileges and Permissions](#)
- [Working with Object Dependencies](#)
- [Adding a Comment to an object](#)

#### USING THE OBJECT EDITOR TOOLBAR

The Object editor toolbar appears above the tabs of the Object editor.




It provides the following functions:

<b>Create</b>	Launches a Rapid SQL object creation wizard. For more information, see <a href="#">Creating objects</a> .
<b>Alter</b>	Enabled when a property of the current database has been modified. Opens a dialog that lets you preview the SQL code that will effect the modification and then submit the change. For more information, see <a href="#">Preview</a> .
<b>Drop</b>	Lets you drop one or more database objects and remove their definition from the system catalog. For more information, see <a href="#">Drop</a> .
<b>Extract</b>	Rapid SQL lets you extract data from one database to another database and extract the statements required to create objects into an Interactive SQL window. For more information, see <a href="#">Extract</a> .
<b>Report</b>	Lets you generate detailed or summary reports on database objects. For details, see <a href="#">Report</a> .
<b>Command Menu</b>	Provides menu commands specific to the object being viewed or modified. For a listing of commands available, see the topic for the specific Object editor, later in this chapter.
<b>Refresh</b>	Refreshes or clears Object editor contents.
<b>Close</b>	Closes the Object editor and if appropriate, prompts you to verify any changes.

Commands on the Object editor toolbar are disabled if they do not apply to the object type being viewed or modified.

#### OPENING ANOTHER OBJECT IN THE SAME OBJECT EDITOR EXPLORER

When an object editor is open, the area between the object editor toolbar and the tabs has one or more dropdown lists. For example:





These lists allow you to qualify any object of the same type as the object currently displayed in the Object editor, and display information for that object in the current Object editor. The number of, and specific dropdown lists, differ according to the type of Object editor but there are always sufficient controls to uniquely identify another object of the same type.

### To display the information for another object of the same type

- Use the dropdown lists to specify a different object.

### VIEWING THE SQL/DDDL FOR AN OBJECT

The object editors for most object types feature a **DDL** tab that lets you view an object's underlying SQL code.

### To view the underlying DDL/SQL for a database object

- 1 Open an editor on an object type for which DDL code can be displayed. For details, see [Opening an Object Editor](#).

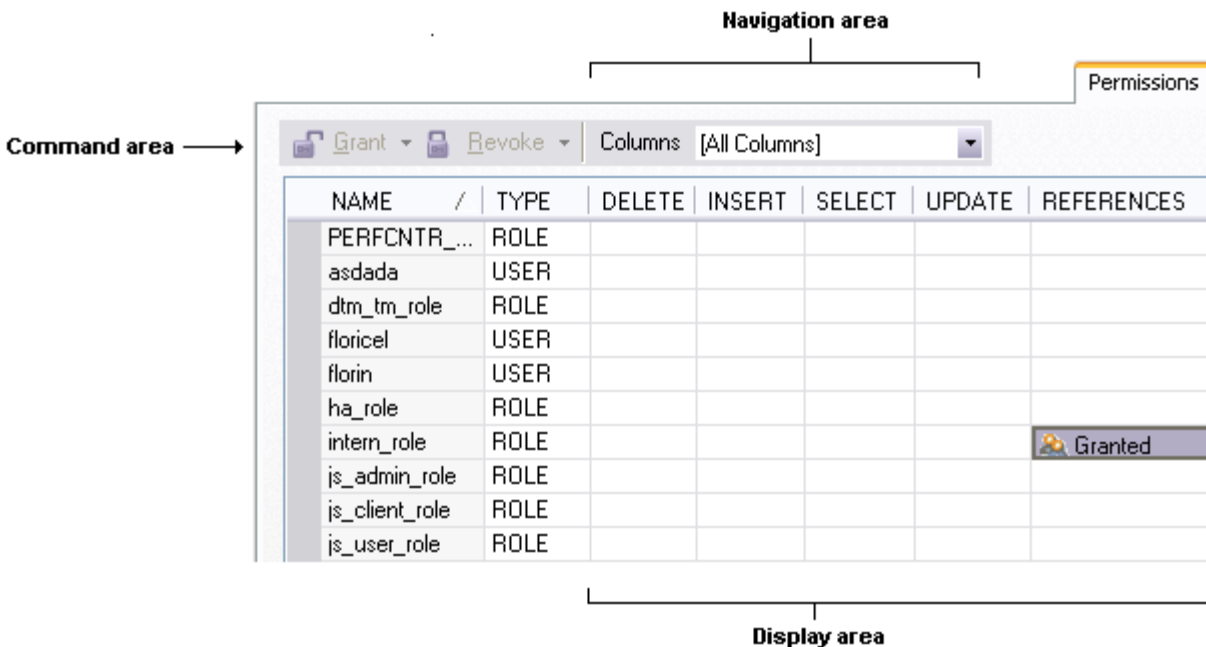
See the topics for specific Object editors later in this chapter for information on whether that object type supports display of underlying SQL code.

- 2 Click the **DDL** tab.


### WORKING WITH PRIVILEGES AND PERMISSIONS

When you open an Object editor on an object with associated privileges, the **Permissions** (or **Object Permissions** or **System Permissions**) tab for that editor displays the relevant privileges and lets you make changes accordingly:

For database objects such as tables or procedures, to which permissions or privileges are granted, a **Permissions** tab lets you manage permissions or privileges to that object from all grantees, such as users or groups, on that datasource. The tab has a navigation area, a command area, and a display area.



The Navigation area lets you change the content of the Display area to view more specific privilege details. For example, when viewing privileges for a database object such as a table, the dropdown lists in the navigation area, if present, let you drill down to populate the display area with privileges for a lower level component, such as a particular column. Similarly, when viewing privileges for a grantee such as a user, the **Object Permissions** tab's Navigation area lets you populate the display area with that user's permissions on specific object types, such as tables or procedures.



The Display area shows privilege details for an object or grantee. When viewing privileges for a grantee such as a user or role, the display area shows privileges for that recipient on objects currently qualified by the current selection in the Navigation area. When viewing privileges for an object such as a table, the display area shows the privileges for all grantees. Each cell in the Display area corresponds to a specific permissions and a cell representing a granted permission shows a distinctive icon.


The Command area lets you initiate granting or revocation of permissions.

### To view or modify privileges using an Object editor

- 1 Open an editor on a database object with associated privileges or permissions. For details, see [Opening an Object Editor](#).
- 2 Click the **Permissions** (or **Privileges** or **Object Privileges** or **System Privileges**) tab.
- 3 Use the dropdown lists in the navigation area, if present, to populate the Display area with details for more specific or different permissions.
- 4 In the Display area, select a cell corresponding to a specific permission.
- 5 Use the Command area controls to grant or revoke permissions.
- 6 Submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

### WORKING WITH OBJECT DEPENDENCIES

For objects such as views and procedures, whose definition references other objects, such as tables, Rapid SQL lets you view all dependencies. The Object editors for referencing or referenced objects have a **Dependencies** tab that shows all dependencies for an object.



In addition to letting you view dependencies, Object editors let you drop a referencing or referenced object, or open an Object editor on that object.

### To manage dependencies for an object

- 1 Open an editor on a database object with referencing or referenced objects. For details, see [Opening an Object Editor](#).

See the object editor descriptions in subsequent topics to verify that an Object editor supports display of dependencies.

- 2 Click the **Dependencies** tab.

The **Dependency Name** area lists all referenced or referencing objects. Objects are grouped according to their object type.


- 3 Optionally, select a referenced or referencing object in the right pane and either:
  - Use the **Edit** button to open an Object editor on that object
  - Use the [Drop](#) button to initiate dropping that object.

#### ADDING A COMMENT TO AN OBJECT

The object editors for certain object types feature a **Comment** tab that lets you add an explanatory note to specific object definitions. Comments are stored in the REMARKS column of the object's system-catalog.

#### To add a comment to an object

- 1 Open an editor on an object type that permits comments. For details, see [Opening an Object Editor](#).
- 2 Click the **Comment** tab.



- 3 In the **Comment** area, type an explanatory note of up to 254 characters long.
 

**NOTE:** Right-clicking in the Comment area opens a context menu that lets you perform edit text operations such as search, selection, copying and pasting, as well as coding-specific operations such as enabling/disabling line numbers and indenting lines.
- 4 Submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).


#### PREVIEWING AND SUBMITTING OBJECT EDITOR CHANGES

After you use an Object editor to modify the settings or properties of a database object, you can preview the SQL that will be executed to effect those changes on the datasource. You can then submit the SQL for execution on the server.

#### To preview and submit object editor changes to a database object

- 1 Click the **Alter** button on the Object editor toolbar. For details, see [Overview and common usage of object editors](#).
 

A **Preview: Alter** dialog opens.



2 Use one of the following options to submit and execute your changes:

- **Execute** button - executes the changes immediately.
- **Schedule** button - opens a dialog that lets you schedule the change. For more information, see [Scheduling](#).

## IBM DB2 FOR LINUX, UNIX, AND WINDOWS OBJECT EDITORS

Rapid SQL includes an Object Editor for all supported IBM DB2 or Linux, Unix, and Windows objects. To see an Editor for a specific object, click the corresponding link below:

- [Aliases Editor \(IBM DB2 LUW\)](#)
- [Check Constraints Editor \(IBM DB2 LUW\)](#)
- [Databases Editor \(IBM DB2 LUW\)](#)
- [Foreign Keys Editor \(IBM DB2 LUW\)](#)
- [Functions Editor \(IBM DB2 LUW\)](#)
- [Indexes Editor \(IBM DB2 LUW\)](#)
- [Materialized Query Tables Editor \(IBM DB2 LUW\)](#)
- [Packages Editor \(IBM DB2 LUW\)](#)
- [Primary Keys Editor \(IBM DB2 LUW\)](#)
- [Procedures Editor \(IBM DB2 LUW\)](#)
- [Sequences Editor \(IBM DB2 LUW\)](#)

- [Structured Types Editor \(IBM DB2 LUW\)](#)
- [Tables Editor \(IBM DB2 LUW\)](#)
- [Tablespaces Editor \(IBM DB2 LUW\)](#)
- [Triggers Editor \(IBM DB2 LUW\)](#)
- [Unique Keys Editor \(IBM DB2 LUW\)](#)
- [User Datatypes Editor \(IBM DB2 LUW\)](#)
- [Views Editor \(IBM DB2 LUW\)](#)

## ALIASES EDITOR (IBM DB2 LUW)

The Aliases Editor lets you view basic alias properties, change the comment associated with an alias, and view the DDL that will be issued to alter the alias definition.

### To edit an alias

- 1 Open an editor on the alias. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	In addition to the owning <b>Schema</b> , this tab lets you view the <b>Target Owner</b> , <b>Target Type</b> , and <b>Target Name</b> , of the referenced object.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## CHECK CONSTRAINTS EDITOR (IBM DB2 LUW)

The Check Constraints Editor lets you modify and enable/disable check constraints.

**To edit a check constraint:**

- 1 Open an editor on the check constraint. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Definition</b>	Use the <b>Enabled</b> control to enable/disable the check constraint.  You can also edit the condition in the <b>Check Condition</b> box. The <b>Table Column</b> button opens a dialog that lets you select and paste column names into the check condition expression.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**DATABASES EDITOR (IBM DB2 LUW)**

The Databases Editor lets you manage database tablespace placement details and modify configuration parameters for the database.

**To edit a database**

- 1 Open an editor on the database. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Placement</b>	Select a tablespace and click the <b>Edit</b> button to open an object editor on that tablespace. For more information, see <a href="#">Tablespaces Editor (IBM DB2 LUW)</a> .
<b>Options</b>	Lets you modify the DB2 Database Manager configuration file parameters for the database. Select a parameter and click the <b>Edit</b> button to open a dialog that lets you change the current value.  <b>NOTE:</b> To set database options for all future databases, set the database options on the model database.
<b>DDL</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**FOREIGN KEYS EDITOR (IBM DB2 LUW)**

The Foreign Keys Editor lets you manage column mapping for a foreign key, modify update and delete rule actions, and specify a NOT ENFORCED value.

**To edit a foreign key**

- 1 Open an editor on the foreign key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Column Mapping</b>	The existing column mapping for the foreign key is represented by selected columns in the <b>Main Table</b> and <b>Referenced Table</b> lists. Additional candidates in the <b>Referenced Table</b> list are indicated by enabled column check boxes.  If necessary, use the <b>Table</b> dropdown in the <b>Referenced Table</b> group to choose a new table for this foreign key. Select or deselect columns in the <b>Main Table</b> list and <b>Referenced Table</b> list to form the referential constraint between the two tables.
<b>Properties</b>	Lets you modify the <b>Delete Rule</b> (NO ACTION, RESTRICT, CASCADE, SET NULL) and <b>Update Rule</b> (NO ACTION, RESTRICT) actions. This tab also lets you deselect the <b>Enforced</b> check box to specify a NOT ENFORCED option.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**FUNCTIONS EDITOR (IBM DB2 LUW)**

The Functions Editor lets you manage the body, inputs and outputs, and properties for a function.

**To edit a function**

- 1 Open an editor on the function. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Displays the <b>Language</b> , <b>Return Type</b> , <b>External Name</b> , <b>SQL Access Level</b> , <b>Origin</b> , <b>Threadsafe</b> , <b>Fenced</b> , <b>Scratchpad</b> , <b>Scratchpad Length</b> , <b>Allow Parallel</b> , <b>Final Call</b> , <b>Parameter Style</b> , <b>Inherit Special Registers</b> , <b>DBInfo</b> , <b>Deterministic</b> , <b>External Action</b> , <b>Called on Null Input</b> , and <b>Parameter CCSID</b> properties.
<b>Parameters</b>	Lets you manage function parameters. On opening, this tab shows the existing parameters. Optionally you can:
	Select a parameter from the list, modify the <b>Type</b> and if appropriate, the <b>Precision</b> , <b>Size</b> , and <b>Scale</b> of the parameter.
	Click the New button, provide a name for the new parameter and modify its attributes.
	Select a parameter and click the Drop button to delete the parameter.
	Select a parameter and use the arrow buttons to reorder the parameter list.



Tab	Settings and tasks
<b>Return Scalar</b> (only available for functions created with a <b>Function Type</b> of EXTERNAL SCALAR, SOURCED or TEMPLATE and for SQL <b>Function Type</b> functions with an explicit <b>Return Type</b> of SCALAR)	Lets you change the <b>Type</b> of a return scalar and depending on the type specified, modify the <b>Precision</b> , <b>Scale</b> , and <b>Size</b> of the return scalar.
<b>Return Columns</b> (only available for functions created with a <b>Function Type</b> of OLEDB and EXTERNAL TABLE and for SQL <b>Function Type</b> functions with an explicit <b>Return Type</b> of TABLE or ROW )	Lets you manage return columns for a function. On opening, this tab shows the existing columns. Optionally you can:
	Select a column from the list, modify the <b>Type</b> and if appropriate, the <b>Precision</b> , <b>Size</b> , and <b>Scale</b> of the parameter. You can also specify that the return column is to be treated <b>As Locator</b> .
	Click the New button, provide a name for the new column and modify its attributes.
	Select a parameter and click the Drop button to delete the parameter.
	Select a parameter and use the arrow buttons to reorder the parameter list.
<b>Body</b>	Lets you edit the body of the function.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## INDEXES EDITOR (IBM DB2 LUW)

The Index Editor lets you manage columns and properties, and view space details for an index.

**NOTE:** IBM DB2 for Windows, Unix, and Linux, lets you segregate *Include columns*; columns that are to be part of the index but not part of the unique key.

**To edit an index**

- 1 Open an editor on the index. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Lets you manage columns that make up the index. On opening, this tab shows the existing columns. Optionally you can:	
		Change the <b>Sort</b> order of a column.
		Click the New button and select a column name from the dropdown, to add a column to the index.
		Select a column and click the Drop button to delete the column from the index.
		Select a column and use the arrow buttons to reorder the columns in the index.
<b>Properties</b>	Lets you work with properties in the following categories:	
	<b>Attributes</b>	Lets you view the <b>Defined By</b> property. Lets you set the <b>Index Type</b> (UNIQUE or NONUNIQUE), specify the index as <b>Clustered</b> , and <b>Allow Reverse Scans</b> .
	<b>Storage</b>	Lets you set <b>Percent Free</b> and <b>Minimum Percent Used</b> settings.
<b>Include Columns</b>	Lets you Include columns that are to be part of the index but not part of the unique key. On opening, the tab shows a listing of the Include columns currently defined as part of this index. Optionally, take one of the following actions:	
		Click the New button to add a new column to the index.
		Select an existing column and click the Delete button to delete that column.
<b>Space</b>	Lets you view the following property groups:	
	<b>Attributes</b>	<b>Make Unique, System Required, Total Keys, Page Fetch Pairs, and Distinct Keys</b>
	<b>Statistics</b>	<b>Index Level, Cluster Ratio, Cluster Factor, Leaf Pages, Sequential Pages, and Density</b>
	<b>Cards</b>	<b>First Key, First 2 Keys, First3 Keys, First 4 Keys, and Full Keys</b>
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .	
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## MATERIALIZED QUERY TABLES EDITOR (IBM DB2 LUW)

The Materialized Query Tables Editor lets you manage the columns, base query and options for a materialized query table, and work with storage and performance settings.

### To edit a materialized query table

- 1 Open an editor on the materialized query table. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Displays the currently defined columns for the materialized Query Table. Optionally, you can <b>Add</b> , <b>Insert</b> , or <b>Edit</b> columns for the materialized query table.	
<b>Definition</b>	<b>Base Query</b>	Lets you view the query that this materialized query table is based on.
	<b>Materialized Query Table Options</b>	Lets you specify an immediate or deferred <b>Refresh</b> , and explicitly set a Materialized Query Table as <b>Enabled</b> or <b>Disabled</b> or go with the <b>Default</b> . This tab also displays whether the Materialized Query Table is maintained by the system or by a user.
<b>Performance</b>	<b>Tablespace Placement</b>	Displays <b>Table Tablespace</b> , <b>Index Tablespace</b> , and <b>Long Tablespace</b> settings for the materialized query table.
	<b>Log Options</b>	Lets you select logging options: <b>Logged initially</b> <b>Data Capture</b> - Lets you select none or change to <b>Include Longvar Columns</b> . NOTE: If you select the <b>Data Capture</b> option, the table name / column name cannot be longer than 18 bytes.
	<b>Options</b>	Let you set the <b>Percent Free</b> , <b>Locksize</b> , <b>Append</b> , and <b>Volatile</b> settings for the materialized query table.
	<b>Partition Columns</b>	Use the <b>Add</b> and <b>Delete</b> buttons to manage the partition columns for the materialized query table.
<b>Space</b>	<b>Page Information</b>	Lets you view <b>Row Count</b> and <b>Num. of Overflow Rows</b> values.
	<b>Row Information</b>	Lets you view <b>Num. of Pages with Rows</b> , <b>Num. of Pages</b> , <b>Percent Free</b> values.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .	
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .	
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PACKAGES EDITOR (IBM DB2 LUW)

The Packages Editor lets you view contents and settings for a package.

### To edit a package

- 1 Open an editor on the package. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Definition</b>	Lets you view the following settings: <b>Binder, Definer, Default Schema, Degree, Function Path, Language Level, SQL Math Warn, SQL Warn, Buffered Inset, Status, Code Page, Total Sections, Multi-node Bound, Intra-partition, Query Optimization, Cursor Blocking, Isolation Level, Date Time Format, Last Bind Time, Explicit Bind Time, Explain Level, Explain Snapshot, and Explain Mode.</b>
<b>Statements</b>	Shows the contents of a package statement. Optionally, you can select a statement and click <b>Explain</b> to copy the statement to an ISQL Editor window. For more information, see <a href="#">ISQL Editor</a> .
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PRIMARY KEYS EDITOR (IBM DB2 LUW)

The Primary Keys Editor lets you manage columns and properties for a primary key, and view space usage/allocation details.

### To edit a primary key

- 1 Open an editor on the primary key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Columns</b>	Displays the columns that make up the primary key. Optionally, you can:
	Select a column and use the Delete key to remove the column from the primary key.
	Select a column and use the arrow keys to reorder the columns in the primary key.
<b>Properties</b>	Displays the owning <b>Table Schema</b> and <b>Table Name</b> , the <b>Name</b> of the primary key, a <b>Defined By</b> property, and <b>Percent Free</b> and <b>Minimum Percent Used</b> properties.

Tab	Settings and tasks	
<b>Space</b>	Lets you view the following property groups:	
	<b>Attributes</b>	<b>Make Unique, System Required, Total Keys, Page Fetch Pairs, and Distinct Keys</b>
	<b>Statistics</b>	<b>Index Level, Cluster Ratio, Cluster Factor, Leaf Pages, Sequential Pages, and Density</b>
	<b>Cards</b>	<b>First Key, First 2 Keys, First3 Keys, First 4 Keys, and Full Keys</b>
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object.</a>	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object.</a>	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes.](#)

## PROCEDURES EDITOR (IBM DB2 LUW)

The Procedures Editor lets you manage the body and parameters for a procedure.

### To edit a procedure

- Open an editor on the procedure. For details, see [Opening an Object Editor.](#)
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	<p>Lets you modify the following settings: <b>SQL Access Level</b> (MODIFIES SQL DATA, CONTAINS SQL, READS SQL DATA, NO SQL), number of <b>Result Sets</b>, <b>External Action</b>, New Save Point, Inherit Special Registers, <b>Deterministic</b>, and <b>Parameter CCSID</b>.</p> <p>Lets you view the following settings: <b>Language</b> (C, COBOL, JAVA, OLE, SQL CLR), <b>External Name</b>, Threadsafe, <b>Fenced</b>, <b>Parameter Type</b> (GENERAL, DB2, GENERAL WITH NULLS, SQL), <b>Program Type</b> (MAIN, SUB and not enabled for a <b>Parameter Style</b> of SQL), and <b>DBInfo</b>.</p>	
<b>Parameters</b>	Lets you manage procedure parameters. On opening, this tab shows the existing parameters. Optionally you can:	
		Select a parameter from the list, modify the <b>Type</b> , and if appropriate, the <b>Precision</b> , <b>Size</b> , and <b>Scale</b> of the parameter. You can also set the <b>Parameter Mode</b> to INPUT, OUTPUT, or INPUT_OUTPUT.
		Click the New button, provide a name for the new parameter and modify its attributes.
		Select a parameter and click the Drop button to delete the parameter.
		Select a parameter and use the arrow buttons to reorder the parameter list.
	<b>NOTE:</b> You cannot use host variables in the CALL statement for the name of the procedure.	
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object.</a>	
<b>Body</b>	Lets you modify the body of the procedure.	
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies.</a>	
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions.</a>	

Tab	Settings and tasks
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## SEQUENCES EDITOR (IBM DB2 LUW)

The Sequences Editor lets you manage cycle numbers, increments, the datatype, and other options for a sequence.

### To edit a sequence

- Open an editor on the sequence. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Definition</b>	Lets you work with settings in the following categories:	
	<b>Parameters</b>	<p><b>Increment By</b> - Lets you specify the interval between sequence numbers. This integer value can be any positive or negative integer, but it cannot be 0. This value can have 28 or fewer digits. The absolute of this value must be less than the difference of MAXVALUE and MINVALUE. If this value is negative, then the sequence descends. If the increment is positive, then the sequence ascends. If you omit this clause, the interval defaults to 1.</p> <p><b>Minimum Value</b> - Lets you specify the minimum value of the sequence. This integer value can have 28 or fewer digits.</p> <p><b>Maximum Value</b> - Lets you specify the maximum value the sequence can generate. This integer value can have 28 or fewer digits.</p>
	<b>Next Sequence Numbers</b>	Lets you work with sequence cycle numbers.
	<b>Sequence Datatype</b>	Lets you specify the datatype and width for the sequence.
	<b>Options</b>	Lets you specify <b>Cache Size</b> , <b>Cycle When Reach Max/Min</b> , and <b>Generate Numbers in Order</b> (useful when you are using the sequence number as a timestamp) values.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## STRUCTURED TYPES EDITOR (IBM DB2 LUW)

The Structured Types Editor lets you manage the attributes, methods and body for a structured type.

### To edit a structured type

- 1 Open an editor on the structured type. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	<p>Displays the initially-defined values for the following properties: <b>Instantiable</b>, <b>Final Type</b>, <b>With Function Access</b>, <b>Without Comparisons</b>, <b>Inline Length</b>, <b>Supertype Schema</b>, <b>Supertype Name</b>, <b>Cast (Source as Ref) With</b>, <b>Cast (Ref as Source) With</b>, <b>Reference Using</b>, <b>Size</b>, <b>Precision</b>, and <b>Scale</b>.</p> <p>For details on these properties, see <a href="#">Structured Type Wizard (DB2 LUW)</a>.</p>
<b>Attributes</b>	<p>Displays the currently defined attributes for the structured type. Optionally, take one of the following actions:</p>
	<p>Select an attribute and change the <b>Datatype</b> in the Property/Value list. Depending on the datatype you choose you can also provide additional datatype options. For details on option availability, see <a href="#">Structured Type Wizard (DB2 LUW)</a>.</p>
	<p>Add an attribute by clicking the <b>New</b> button and typing a name for the attribute.</p>
	<p>Select an attribute and click <b>Drop</b> to delete the attribute.</p>
<b>Methods</b>	<p>This tab lets you initiate creation of method specifications and prepopulate associated method bodies. On opening, this tab lists all method specifications associated with the structured type and for each method, includes name and language details. Optionally, take one of the following actions:</p>
	<p>Click <b>Add</b> to begin the process of adding a new method specification. Similarly, select a method and click <b>Edit</b> to modify the method specification. For more information, see <a href="#">Adding or editing structured type methods</a>.</p>
	<p>Select a method specification and click <b>Drop</b> to delete that method specification from the structured type.</p>
	<p>The <b>Create Body</b> and associated controls are available for selected methods for which no method body has yet been defined. They lets you prepopulate the method body definition:</p> <p><b>External name</b> - lets you provide a 'string' or SQL identifier for a method specified with a <b>Language</b> of C, JAVA, or OLE.</p> <p><b>Transform group</b> - lets you specify the transform group used when invoking the method. This setting is available for methods specified with a <b>Language</b> of C, JAVA, or OLE.</p> <p><b>As identifier</b> - This setting is available for methods specified with a <b>Language</b> of C. When this check box is checked, the method body is created with the external name as provided. Otherwise, the external named provided will appear within quotes in the CREATE METHOD statement on the <b>Body</b> tab.</p> <p><b>Inherit isolation level with lock request</b> - specifies whether the INHERIT clause is specified as INHERIT ISOLATION LEVEL WITHOUT LOCK REQUEST or INHERIT ISOLATION LEVEL WITH LOCK REQUEST. This setting is available for methods specified with a Language of SQL.</p> <p>After specifying options, click <b>Create Body</b> to work with the CREATE METHOD statement you generated on the editor's <b>Body</b> tab.</p>

Tab	Settings and tasks
<b>Body</b>	Lets you view and modify the CREATE METHOD statement generated with your choices on the <b>Methods</b> tab.
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

#### ADDING OR EDITING STRUCTURED TYPE METHODS

The **Add Method** and **Edit Method** wizards let you work with the methods of a structured type. They let you specify basic properties, parameter and return value details for methods.

#### To add or edit a structured type method:

- Open an editor on a structured type. For details, see [Opening an Object Editor](#).
- Navigate to the **Methods** panel, and then open a wizard using one of the following techniques:
  - Click **Add** to create a new method
  - Select an existing method and click **Edit**.
- Use the following table as a guide to understanding and modifying the settings on the tabs of this wizard:

Tab	Settings and tasks	
Properties	When adding a method this pane lets you specify a <b>Name</b> , <b>Specific Name</b> , <b>Language</b> (OLE, SQL, JAVA, C), and <b>SQL Access Level</b> (CONTAINS SQL, READS SQL DATA, NO SQL). When editing a method, these properties are for display only.	
<b>Advanced</b> (available after clicking the <b>Advanced</b> button on the <b>Properties</b> pane)	Lets you select or enable the following ADD METHOD options: <b>Deterministic</b> , <b>External Action</b> , <b>Called on NULL Input</b> , <b>DBINFO</b> , <b>Fenced</b> , <b>Allow Parallel</b> , <b>Scratchpad</b> , <b>Scratchpad Length</b> , <b>Parameter Style</b> (SQL, DB2GENERAL), and <b>Final Call</b> .	
<b>Parameters</b>	Lets you work with input parameters:	
		Add a parameter by clicking the New button and typing a name for the parameter.
		Select a parameter and change the <b>Datatype</b> in the Property/Value list. Depending on the datatype you choose you can also provide <b>Precision</b> , <b>Size</b> , and <b>Scale</b> options.
		Delete a selected parameter by clicking the DELETE button.
		Order the parameters using the arrow buttons.
<b>Return</b>	Lets you provide details of the return value of the method:	
	<b>Return Datatype</b>	Lets you specify the base <b>Type</b> of the return value. Depending on the datatype you choose you can also provide <b>Precision</b> , <b>Size</b> , and <b>Scale</b> options.
	<b>Cast Datatype</b>	This group lets you optionally use the CAST FROM form of the ADD METHOD RETURN clause. It lets you have the method return a different datatype, cast from the datatype specified in the <b>Return Datatype</b> group. This feature must be <b>Enabled</b> and provide the same type options as the <b>Return Datatype</b> group.



Tab	Settings and tasks
	An <b>As Locator</b> option specifies that the method return a LOB locator instead of the actual value. The option is only available for LOB and LOB-based datatypes. <b>As Locator</b> applies to a LOB type in the <b>Cast Datatype</b> group, if specified. Otherwise it applies to a LOB type in the <b>Return Datatype</b> group

- 4 When ready click **Finish**.

## TABLES EDITOR (IBM DB2 LUW)

The Tables Editor lets you manage basic properties, columns, dimension columns, distribution key columns, partitions, tablespaces, and constraints for a table.

**TIP:** Before modifying a table, familiarize yourself with the material in [Altering Tables for IBM DB2 LUW for Linux, Unix, and Windows](#).

### To edit a table

- 1 Open an editor on the table. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Displays the currently defined columns in the table. Optionally, you can:	
		Select a column, and in the <b>Property/Value</b> list modify property values for that column.
		Click <b>Add Column</b> , provide a name for the new column, and set property values for the column.
		Select a column and click <b>Delete</b> to remove the column from the table.
<b>Properties</b>	Displays the <b>Name</b> , <b>Created</b> , <b>Last RunStats</b> , <b>Invalidate Time</b> , and <b>Defined By</b> properties. Lets you set <b>Percent Free</b> , <b>Lock Size</b> , <b>Append Data</b> , <b>Volatile</b> , <b>Compress</b> , <b>Row Compression</b> , <b>Security Policy</b> , <b>RestrictDrop</b> , <b>LogIndexBuild</b> , <b>CCSID</b> , <b>Do not initially log</b> , and <b>Data Capture</b> properties.	
<b>Status</b>	Lets you view the following statistics: <b>Total Number of Rows</b> , <b>Number of Overflow Rows</b> , <b>Number of Pages With Rows</b> , <b>Number of Pages</b> , <b>Table Status</b> , <b>Row Organization</b> , <b>RowTypeSchema</b> , <b>RowTypeName</b> , <b>AccessMode</b> , <b>ActiveBlocks</b> , <b>AvgCompressedRowSize</b> , <b>AvgRowCompressionRatio</b> , <b>CodePage</b> , <b>PercentOfPagesSaved</b> , <b>LastRegeneratedTime</b> , and <b>ProtectionGranularity</b> .	
<b>Partition</b>	This tab provides details on the partition columns and data partitions for the table. Optionally you can:	
		Use the New or Delete buttons to create or drop a partition column or data partition.
		Use the Edit button to edit a data partition.
		Use the <b>Commands</b> menu to attach or detach a data partition.
<b>Tablespaces</b>	<b>Data Tablespace</b> , <b>Long Tablespace</b> , and <b>Index Tablespace</b>	Lets you view the <b>Database partition Group</b> , <b>Managed By</b> , <b>Page Size</b> , and <b>Extent Size</b> properties. Lets you choose the <b>Name</b> of a tablespace.

Tab	Settings and tasks
<b>Dimensions</b>	Lets you group columns to form a dimension:
	Click the New button to add a new column to the dimension for the table.
	Select a column and click the Edit button to modify the dimension column properties.
	Select a column and click the Delete button to drop a column from the dimension.
<b>Distribution Key Columns</b>	Lets you group one or more columns to form a distribution key:
	Click the New button to add a new column to the dimension for the table.
	Select a column and click the Edit button to modify the dimension column properties.
	Select a column and click the Delete button to drop a column from the dimension.
<b>Indexes</b>	Lets you manage indexes for a table:
	Click <b>Add</b> to open a dialog that lets you add a new index to the table.
	Select an index and click <b>Edit</b> to open a dialog that lets you edit index properties.
	Select an index and click <b>Drop</b> to open a dialog that lets you remove the index from the table.
<b>Constraints</b>	Lets you manage primary key, unique key, foreign key, and check constraints for a table. Constraints are grouped by type, under folders:
	Select a constraint type folder and click <b>Add</b> to open a dialog that lets you add a constraint of that type.
	Select a constraint and click <b>Edit</b> to open a dialog that lets you modify the constraint details.
	Select a constraint and click <b>Drop</b> to remove the constraint.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

#### ALTERING TABLES FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

The ALTER TABLE command of Transact-SQL is limited to adding NULL columns to the end of a table and to adding or dropping constraints. Unfortunately, this scenario does not address many requirements of administrators and developers who need to add, delete or modify columns more broadly:

- Add columns anywhere in a table, not just the end
- Add columns that do not permit a NULL value.
- Change the NULL/NOT NULL status of table columns
- Change column datatypes to other compatible types
- Change the length of datatypes

- Delete a column

Due to the limitations of the ALTER TABLE command, the only way to make broader modifications is to write SQL scripts that step through all desired changes. To perform an enhanced table alter, Rapid SQL constructs an SQL script that completes the following steps:

- 1 Renames the existing table so that the original and its data remain intact
- 2 Builds a CREATE TABLE statement with the new table definition, including declared defaults, primary key and check constraints
- 3 Builds an INSERT statement to copy data from the original, renamed table to the new one
- 4 Builds foreign keys on the new table
- 5 Reapplies any privileges granted on the table
- 6 Rebuilds all dependencies on the new table, including indexes, triggers, procedures, packages, functions and views. When rebuilding procedures, functions, packages and views, Rapid SQL also rebuilds any permissions on them.

## TABLESPACES EDITOR (IBM DB2 LUW)

The Tablespaces Editor lets you manage containers, basic properties, performance, space usage, and associated objects for a tablespace.

### To edit a tablespace

- 1 Open an editor on the tablespace. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	Lets you work with settings in the following categories:	
	<b>Tablespace</b>	Lets you view the <b>Type</b> (REGULAR, LARGE, TEMPORARY, or USER TEMPORARY), <b>Use Automatic Storage</b> , <b>Managed By</b> , and <b>Database Partition Group</b> properties. Lets you modify the <b>Buffer Pool</b> and <b>Drop Recovery</b> properties.
	<b>Performance</b>	Lets you view the <b>Page Size</b> and <b>Extent Size</b> properties. Lets you specify or select the <b>Prefetch Automatic</b> , <b>Overhead</b> , <b>Transfer Rate</b> , and <b>File System Caching</b> properties.
	<b>Automatic Storage</b>	Lets you work with the <b>AutoResize</b> (enables following size settings), <b>Initial Size</b> , <b>Increase Size</b> , <b>Max Size Unlimited</b> , and <b>Max Size</b> (enabled if <b>Max Size Unlimited</b> is not set) attributes. The <b>Initial Size</b> and <b>Max Size</b> values can be provided as kB, mB, or gB values.

Tab	Settings and tasks	
<b>Container</b>	This option is only available for database managed tablespaces. This tab lets you add or delete containers of a tablespace.	
		Select an existing container and in the Property/Value list, use the <b>Size</b> controls to RESIZE, EXTEND, or REDUCE the container.
		Add a container by clicking <b>New</b> , and then in the Property/Value list, provide <b>Name</b> , number of <b>Database Partitions</b> , <b>Type</b> (DEVICE, FILE) and original <b>Size</b> .
		Click <b>Delete</b> to remove a container.
<b>Performance</b>	Lets you manage settings in the following categories:	
	<b>Page Setting</b>	Lets you view <b>Page Size</b> and <b>Extent Size</b> settings and modify the <b>Prefetch Size</b> setting.
	<b>I/O Setting</b>	Let you modify the <b>Overhead</b> and <b>Transfer Rate</b> settings.
	<b>Dropped Table Settings</b>	Lets you view the <b>Recovery Status</b> .
	<b>Defaults</b>	Lets you view the <b>Nodegroup</b> and specify a <b>Bufferpool</b> .
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .	
<b>Space</b>	Lets you view the table usage and the distribution of space for a tablespace. Specific statistics include <b>Free Pages</b> , <b>Used Pages</b> , <b>Reserved Pages</b> , and <b>Total Pages</b> .	
<b>Objects</b>	Lets you manage database objects associated with the tablespace. Objects are organized in a tree structure with folders containing the objects. Optionally, take one of the following actions:	
		Select an object and click <b>Edit</b> to open an object editor on the selected object.
		Select an object and click <b>Drop</b> to initiate dropping the selected object.
<b>Privileges</b>	For details on using this tab, see <a href="#">Previewing and Submitting Object Editor Changes</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## TRIGGERS EDITOR (IBM DB2 LUW)

The Triggers Editor lets you view properties for a trigger.

### To edit a trigger

- Open an editor on the trigger. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you view the <b>Trigger Timing</b> , <b>Trigger Events</b> , <b>Trigger Type</b> , <b>Object Status</b> , <b>Definer</b> , and <b>Function Path</b> properties.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .

Tab	Settings and tasks
<b>Definition</b>	Lets you view and modify the CREATE TRIGGER DDL that will implement your changes. To modify a trigger, Rapid SQL must drop then create the trigger.

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## UNIQUE KEYS EDITOR (IBM DB2 LUW)

The Unique Keys Editor lets you manage columns and basic properties for a unique key, as well as view space details.

### To edit a unique key

- Open an editor on the unique key. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Lets you manage columns that make up the unique key. On opening, this tab shows the existing columns. For each column, the datatype (and if applicable the precision in brackets) and whether the table definition permits nulls in the target table column.  Optionally you can:	
		Click the New button and select a column name from the dropdown, to add a column to the index.
		Select a column and click the Drop button to delete the column from the index.
<b>Properties</b>	Displays the owning <b>Table Schema</b> and <b>Table Name</b> , the Name of the primary key, a <b>Defined By</b> property, and <b>Percent Free</b> and <b>Minimun Percent Used</b> properties.	
<b>Space</b>	Lets you view settings in the following categories:	
	<b>Attributes</b>	Lets you view <b>Make Unique</b> , <b>System Required</b> , <b>Total Keys</b> , <b>Distinct Keys</b> , and <b>Page Fetch Pairs</b> settings.
	<b>Statistics</b>	Lets you view <b>Index Level</b> , <b>Cluster Ration</b> , <b>Cluster Factor</b> , <b>Leaf Pages</b> , <b>Sequential Pages</b> , and <b>Density</b> settings.
	<b>Cards</b>	Lets you view <b>First Key</b> , <b>First 2 Keys</b> , <b>First 3 Keys</b> , <b>First 4 Keys</b> and <b>Full Keys</b> settings.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## USER DATATYPES EDITOR (IBM DB2 LUW)

The User Datatypes Editor lets you manage the basic properties of a user datatype.

**To edit a user datatype**

- 1 Open an editor on the user datatype. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Displays the user <b>Datatype</b> name. Lets you select a new <b>Owner</b> or <b>Type</b> , and depending on the type specified, offers additional type-specific properties such as <b>Size</b> . Also lets you modify the <b>Allow Bit Data</b> property, if appropriate.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**USERS EDITOR (IBM DB2 LUW)**

The Users editor lets you manage object dependencies and permissions for a user.

**To edit a user**

- 1 Open an editor on the user. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Displays the user <b>Name</b> .
<b>Objects</b>	Lets you manage database objects associated with the user. Objects are organized in a tree structure with folders containing the objects. Optionally, take one of the following actions:
	Select an object and click <b>Edit</b> to open an object editor on the selected object.
	Select an object and click <b>Drop</b> to initiate dropping the selected object.
<b>Object Permissions and System Permissions</b>	For details on using these tabs, see <a href="#">Working with Privileges and Permissions</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**VIEWS EDITOR (IBM DB2 LUW)**

The Views Editor lets you manage the columns as well as view and modify properties for a view.

## To edit a view

- 1 Open an editor on the view. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
Properties	Displays all columns for the view. Details for each column include the <b>Column Name</b> , the <b>Datatype</b> (and if applicable, with the precision in parentheses), and whether or not <b>Nulls</b> are allowed for that column.  You can also set the <b>EnableQueryOptimization</b> property for the view.
Definition	Lets you view and modify the CREATE VIEW DDL that will implement any changes you make in this editor. To modify a view, Rapid SQL must drop then create the view.
Comment	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
Dependencies	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
Permissions	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## IBM DB2 FOR Z/OS OBJECT EDITORS

Rapid SQL includes an Object Editor for all supported IBM DB2 for z/OS objects:

- [Aliases Editor \(IBM DB2 Z/OS\)](#)
- [Check Constraints Editor \(IBM DB2 Z/OS\)](#)
- [Databases Editor \(IBM DB2 Z/OS\)](#)
- [DBRMs Editor \(IBM DB2 Z/OS\)](#)
- [Foreign Keys Editor \(IBM DB2 Z/OS\)](#)
- [Functions Editor \(IBM DB2 Z/OS\)](#)
- [Indexes Editor \(IBM DB2 Z/OS\)](#)
- [Packages Editor \(IBM DB2 Z/OS\)](#)
- [Plans Editor \(IBM DB2 Z/OS\)](#)
- [Primary Keys Editor \(IBM DB2 Z/OS\)](#)
- [Procedures Editor \(IBM DB2 Z/OS\)](#)
- [Stogroups Editor \(IBM DB2 Z/OS\)](#)
- [Synonyms Editor \(IBM DB2 Z/OS\)](#)
- [Tables Editor \(IBM DB2 Z/OS\)](#)

- [Tablespaces Editor \(IBM DB2 Z/OS\)](#)
- [Triggers Editor \(IBM DB2 Z/OS\)](#)
- [Unique Keys Editor \(IBM DB2 Z/OS\)](#)
- [User Datatypes Editor \(IBM DB2 Z/OS\)](#)
- [Users Editor \(IBM DB2 Z/OS\)](#)
- [Views Editor \(IBM DB2 Z/OS\)](#)

## ALIASES EDITOR (IBM DB2 Z/OS)

The Aliases Editor lets you view details of an alias definition.

### To edit an alias

- 1 Open an editor on the alias. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	In addition to the owning <b>Schema</b> and <b>Name</b> of the alias, this tab lets you view the <b>Target Owner</b> , <b>Target Type</b> , and <b>Target Name</b> , of the referenced object.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## CHECK CONSTRAINTS EDITOR (IBM DB2 Z/OS)

The Check Constraints Editor lets you view definition details and edit a check condition expression.



**To edit a check constraint**

- 1 Open an editor on the check constraint. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Definition</b>	Lets you view basic identification information on the check constraint: <b>Table Schema</b> and <b>Table Name</b> , the <b>Name</b> of the constraint and the date it was <b>Created</b> .  You can also edit the condition in the <b>Check Condition</b> box. The <b>Table Columns</b> button acts a time saver in editing the condition. It opens a dialog that lets you select and paste column names into the check condition expression.
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**DATABASES EDITOR (IBM DB2 Z/OS)**

The Databases Editor lets you manage basic properties, permissions, and object dependencies for a database.

**To edit a database**

- 1 Open an editor on the database. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	The <b>Database Creation</b> group lets you view the <b>Name</b> and <b>Group Member</b> properties and set the <b>Type</b> property. The <b>Database Attributes</b> group lets you view the <b>Encoding Scheme</b> , <b>Create Date</b> , and <b>Last Altered</b> properties and lets you set the <b>Tablespace Buffer Pool</b> , <b>Index Buffer Pool</b> , <b>Storage Group</b> , and <b>CCSID</b> properties.
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**FOREIGN KEYS EDITOR (IBM DB2 Z/OS)**

The Foreign Keys Editor lets you manage column mapping and specify a delete rule for a foreign key.

**To edit a foreign key**

- 1 Open an editor on the foreign key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Column Mapping</b>	The existing column mapping for the foreign key is represented by selected columns in the <b>Main Table</b> and <b>Referenced Table</b> lists. Additional candidates in the <b>Referenced Table</b> list are indicated by enabled column check boxes.  If necessary, use the <b>Table</b> dropdown in the <b>Referenced Table</b> group to choose a new table for this foreign key. Select or deselect columns in the <b>Main Table</b> list and <b>Referenced Table</b> list to form the referential constraint between the two tables.
<b>Properties</b>	Lets you specify a <b>Delete Rule</b> (CASCADE, NO ACTION, RESTRICT, SET NULL).
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**FUNCTIONS EDITOR (IBM DB2 Z/OS)**

The Functions Editor lets you view and modify properties for a function, manage its inputs and outputs, and modify the code in the body of the function.

**To edit a function**

- 1 Open an editor on the function. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you work with properties in the following categories:
	<b>Identification</b> Lets you view the <b>Create Timestamp</b> , <b>Schema</b> , <b>Routine ID</b> , and <b>Origin</b> properties. Lets you set the <b>External Name</b> and <b>Collection ID</b> properties.
	<b>Run Time</b> Lets you view the <b>Result Sets</b> property. Lets you set the <b>WLM Environment</b> , <b>WLM For Nested</b> , <b>ASU Time</b> , and <b>Run Options</b> properties.
	<b>Structure</b> Lets you view the <b>Language</b> , <b>Parameter Style</b> , <b>Number of LOB Columns</b> , <b>Number of Parameters</b> , and <b>Allow Parallel</b> properties. Lets you set the <b>Program Type</b> , <b>Security Type</b> , <b>SQL Access Level</b> , <b>Inherit Special Registers</b> , <b>Fenced</b> , <b>DBINFO</b> , <b>Deterministic</b> , <b>Called On Null Input</b> , <b>External Action</b> , <b>Stay Resident</b> , <b>Final Call</b> , <b>Scratchpad</b> , and <b>Scratchpad Length</b> properties.
	<b>Run Estimates</b> Lets you view the <b>Initial I/Os</b> , <b>I/Os Per Invocation</b> , <b>Initial Instructions</b> , and <b>Instructions Per Invocation</b> properties.
	<b>Java Structure</b> Lets you view the <b>Java Class</b> , <b>Jar ID</b> , <b>Package ID</b> , <b>Method ID</b> , <b>Jar Schema</b> , and <b>Java Signature</b> properties.

Tab	Settings and tasks	
<b>Parameters, Return Columns, and Return Scalar</b>	On opening, these tabs displays the current function parameters, return columns or return scalar. Optionally, take one of the following actions:	
		Select a parameter, return column, or return scalar in the <b>Attributes</b> area, modify values, as permissible. Attributes differ by the tab you chose but typically include items such as <b>Type</b> , <b>Precision</b> , <b>Scale</b> , <b>As Locator</b> , and <b>Size</b> .
		Click the New button to provide datatype and size details for a new parameter, return column, or return scalar.
		Select a parameter, return column, or return scalar and click <b>Drop</b> to delete that parameter.
<b>Body</b>	Lets you modify the code in the text area of the tab.	
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .	
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## INDEXES EDITOR (IBM DB2 Z/OS)

The Indexes Editor lets you manage basic, storage, and space properties for an index, as well as work with its columns and partitions.

### To edit an index

- Open an editor on the index. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Lets you manage the columns or key-expressions that make up the index. On opening, this tab shows the existing columns/key-expressions.  <b>NOTE:</b> Key-expression functionality is only enabled if the <b>Index on Expression</b> setting is enabled on the <b>Properties</b> tab. Optionally you can:	
		Change the <b>Sort</b> order of a column or key-expression.
		If <b>Index on Expression</b> is disabled, click the New button and select a column name from the dropdown to add a column to the index.
		If <b>Index on Expression</b> is enabled, click the New button and type a valid key-expression in the <b>Expression</b> field, to add an expression to the index.
		Select a column or key-expression and click the Drop button to delete the column from the index.

Tab	Settings and tasks	
<b>Properties</b>	Lets you work with the <b>Buffer Pool, Piece Size, Close, Copy, Compress, Index on Expression, and Padded</b> properties. For more information on these properties, see <a href="#">Indexes (DB2 z/OS) - Properties</a> .	
<b>Storage</b>	Lets you select a dataset management scheme and provide associated attribute values. For details, see <a href="#">Indexes (DB2 z/OS) - Storage</a> .	
<b>Partitions</b>	Lets you work with partitions for the index. For each data partition, the listing shows the storage group, VCAT Catalog, primary and secondary space allocations, and if appropriate, the free space percentage, GBP Cache, limit key value, and erase on delete settings. Optionally, you can select a partition and click the Edit button to open an editor that lets you edit that partition.	
<b>Space</b>	Lets you view the following property groups:	
	<b>Attributes</b>	<b>Make Unique, System Required, Total Keys, Page Fetch Pairs, and Distinct Keys</b>
	<b>Statistics</b>	<b>Index Level, Cluster Ratio, Cluster Factor, Leaf Pages, Sequential Pages, Density, DASD Storage, and Data Blocks/Key.</b>
	<b>Cards</b>	<b>First Key and Full Keys</b>
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDL for an Object</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PACKAGES EDITOR (IBM DB2 Z/OS)

The Packages Editor lets you manage properties and bind parameters for a package, as well as view environment information, package contents and associated plans.

### To edit a package

- Open an editor on the package. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you work with the following properties: <b>Average Size, Bind Time, Consistency Token, Dec31, Decimal Point, Group Member, Katakana Charset, Language, Mixed Character Set, Operative, Package Size, Package Source, PDS Name, Precompile Timestamp, Release Bound, String Delimiter, SysEntries, Timestamp, and Valid.</b>
<b>Bind Parameters</b>	Lets you set the <b>Schema Path, Creator, Qualifier, Explain, Validate, Degree, Dynamic Rules, Lock Isolation, Encoding, Page Writes, Resource Release, CurrentData, DB Protocol, Reoptvar, Keep Dynamic, Defer Prepare, SQL Error, and Optimization Hint</b> properties.

Tab	Settings and tasks
<b>Plan/Packlists</b>	Displays the plans contained in the package if the package was bound individually, or as part of a complete collection ID. Optionally, you can:
	Select a plan and click <b>Edit</b> to open the Plans editor on the plan. For more information, see <a href="#">Plans Editor (IBM DB2 Z/OS)</a> .
	Select a plan and click <b>Rebind</b> to initiate rebinding the plan. For more information, see <a href="#">Rebind Plans</a> .
	Select a plan and click <b>Free</b> to initiate deleting the plan. For more information, see <a href="#">Drop</a> .
<b>Statements</b>	Shows the contents of any package statement on the datasource. Optionally, select a statement and click <b>Explain</b> to copy the statement to an ISQL Editor window. NOTE: For more information, see <a href="#">ISQL Editor</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Privileges</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>Environments</b>	Displays run-time environments information for a package. Use the arrow buttons to move environments between the <b>Enabled Environments</b> and <b>Disabled Environments</b> lists. To modify a connection, select the environment and click Edit Connections. For more information, see <a href="#">Connection Editor</a> .
<b>Command</b>	Displays the command that originally built the package.

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## CONNECTION EDITOR

The Connection Editor lets you modify plan and package connections. It is opened from the Packages editor. For more information, see [Packages Editor \(IBM DB2 Z/OS\)](#).

The table below describes the options and functionality of the Connection Editor:

Option	Description
<b>Connections</b>	Displays the connections for the plan or package.
<b>Add</b>	Click to add the connection.

## PLANS EDITOR (IBM DB2 Z/OS)

The Plans Editor lets you manage plan properties, view DBRM information, manage packages and contents, and manage run-time environments for a plan.

**To edit a plan**

- 1 Open an editor on the plan. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Definition</b>	Lets you work with the following attribute groups:	
	<b>Properties</b>	Lets you view the <b>Timestamp</b> , <b>Group Member</b> , <b>Plan Size</b> , <b>Average Size</b> , <b>Valid</b> , <b>Operative</b> , <b>Pkg. List Entries</b> , <b>SysEntries</b> , and <b>Release Bound</b> properties. Lets you set the <b>Owner</b> , <b>Qualifier</b> , and <b>Current Server</b> properties.
	<b>Bind Parameters</b>	Lets you set the <b>Schema Path</b> , <b>Explain</b> , <b>Validate</b> , <b>Degree</b> , <b>Dynamic Rules</b> , <b>Lock Isolation</b> , <b>Resource Acquire</b> , <b>Resource Release</b> , <b>Disconnect</b> , <b>Current Data</b> , <b>DB Protocol</b> , <b>Reoptvar</b> , <b>SQL Rules</b> , <b>Keep Dynamic</b> , <b>Defer Prepare</b> , <b>Encoding</b> , <b>Page Writes</b> , <b>Flag</b> , <b>Optimization Hint</b> , and <b>Cache Size</b> properties.
<b>DBRMs</b>	Lists details for each DBRM associated with a plan. Optionally, select a DBRM from the list and click <b>Edit</b> to open the DBRM editor on that DBRM. For more information, see <a href="#">DBRMs Editor (IBM DB2 Z/OS)</a> .	
<b>Packages</b>	Displays details for each package associated with a plan. Optionally, you can:	
		Select a package and click <b>Edit</b> to open the Package editor on that package. For details, see <a href="#">Packages Editor (IBM DB2 Z/OS)</a> .
		Select a package and click <b>Rebind</b> to open the <a href="#">Rebind Packages</a> dialog.
		Select a package and click <b>Free</b> to open the <a href="#">Free Packages</a> dialog.
<b>DBRM/Packages</b>	Displays the entire contents of the plan, DBRMs and packages, in a single display. The first column contains either a 'D' for DBRM or 'P' for packages. Optionally, take one of the following actions:	
		Select a package list entry and click <b>Edit</b> to open the Package editor on that package. For details, see <a href="#">Packages Editor (Oracle)</a> .
		Select a package and click <b>Rebind</b> to open the <a href="#">Rebind Packages</a> dialog.
		Select a package and click <b>Free</b> to open the <a href="#">Free Packages</a> dialog.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .	
<b>Privileges</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .	
<b>Environments</b>	Displays run-time environments information for a plan. Optionally, you can:	
		Use the arrow buttons to move environments between the <b>Enabled Environments</b> and <b>Disabled Environments</b> lists.
		Select an environment and click <b>Edit Connections</b> to open the <a href="#">Connection Editor</a> .
<b>Command</b>	Displays the command that originally built the plan.	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PRIMARY KEYS EDITOR (IBM DB2 Z/OS)

The Primary Keys Editor lets you manage primary key columns as well as work with storage, space, and partitions for a primary key.

**TIP:** The [refresh](#) button lets you refresh or clear the editor's contents, and log SQL.

### To edit a primary key

- 1 Open an editor on the primary key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Column</b>	Lets you manage columns that make up the primary key. On opening, this tab shows the existing columns. For each column, the listing shows the datatype (and if applicable the precision in brackets) and whether the table definition permits nulls in the target table column. Optionally you can:	
		Change the <b>Sort</b> order of a column.
		Click the New button and select a column name from the dropdown, to add a column to the primary key.
		Select a column and click the Drop button to delete the column from the primary key.
		Select a column and use the arrow buttons to reorder the columns in the primary key.
<b>Properties</b>	Lets you set <b>Buffer Pool</b> , <b>Close</b> , <b>Copy</b> and <b>Piece Size</b> properties.	
<b>Storage</b>	Lets you view <b>Storage Group</b> and <b>VCAT Catalog</b> properties. Lets you set <b>Primary Space Allocation</b> , <b>Secondary Space Allocation</b> , <b>Erase</b> , <b>Free Page</b> , <b>Percent Free</b> , and <b>GBP Cache</b> properties.	
<b>Partitions</b>	Lets you work with partitions for the primary key.	
<b>Space</b>	Lets you view values in the following property groups:	
	<b>Attributes</b>	<b>Make Unique</b> , <b>System Required</b> , <b>Total Keys</b> , <b>Page Fetch Pairs</b> , and <b>Distinct Keys</b> .
	<b>Statistics</b>	<b>Index Level</b> , <b>Cluster Ratio</b> , <b>Cluster Factor</b> , <b>Leaf Pages</b> , <b>Sequential Pages</b> , and <b>Density</b> .
	<b>Cards</b>	<b>First Key</b> and <b>Full Keys</b> .
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PROCEDURES EDITOR (IBM DB2 Z/OS)

The Procedures Editor lets you manage the properties and input/output parameters of a procedure.

### To edit a procedure

- 1 Open an editor on the procedure. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	Lets you work with properties in the following categories:	
	<b>Identification</b>	Lets you view the <b>Schema</b> and <b>Routine ID</b> properties. Lets you set the <b>External Name</b> and <b>Collection ID</b> properties.
	<b>Run Time</b>	Lets you set the <b>Result Sets</b> , <b>WLM Environment</b> , <b>WLM For Nested</b> , <b>ASU Time</b> , and <b>Run Options</b> properties.
	<b>Structure</b>	Lets you view the <b>Language</b> , <b>Number of LOB Columns</b> , <b>Number of Parameters</b> , and <b>DBINFO</b> properties. Lets you set the <b>Program Type</b> , <b>Security Type</b> , <b>SQL Access Level</b> , <b>Parameter Style</b> , <b>Inherit Special Registers</b> , <b>Fenced</b> , <b>Commit on Return</b> , <b>Deterministic</b> , and <b>Stay Resident</b> properties.
	<b>Run Estimates</b>	Lets you view the <b>Initial I/Os</b> , <b>I/Os Per Invocation</b> , <b>Initial Instructions</b> , and <b>Instructions Per Invocation</b> properties.
	<b>Java Structure</b>	Lets you view the <b>Java Class</b> , <b>Jar ID</b> , <b>Package ID</b> , <b>Method ID</b> , <b>Jar Schema</b> , and <b>Java Signature</b> properties.
<b>Parameters</b>	Displays a listing of the existing input/output parameters for the procedure. For the selected parameter, the <b>Datatype</b> list shows details for that parameter, including the <b>Type</b> and <b>Parameter Mode</b> (INPUT, OUTPUT, INPUT_OUTPUT). Depending on the type other parameters such as <b>Size</b> , <b>Precision</b> , or <b>Scale</b> may be available for viewing or modification.  Optionally you can:	
		Select a parameter and in the <b>Datatype</b> list, modify details for that parameter.
		Click the Add button to add a new parameter, provide a name for the new parameter, and edit the parameter values in the <b>Datatype</b> list.
		Click the Delete button to drop a selected parameter.
		Use the arrow buttons to change the order location of a selected parameter.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .	
<b>Body</b>	Lets you modify the SQL code for procedures on the current datasource	
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).



## STOGROUPS EDITOR (IBM DB2 Z/OS)

The Stogroups Editor lets you view and modify volumes, manage function privileges, and view DDL for a stogroup.

### To edit a stogroup

- 1 Open an editor on the stogroup. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Volume Devices</b>	Shows details on volumes associated with the stogroup. Optionally you can:
	Click <b>Add</b> to open a dialog that lets you add volumes.
	Select a volume and click <b>Remove</b> to delete the volume.
<b>Privileges</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## SYNONYMS EDITOR (IBM DB2 Z/OS)

The Synonyms Editor lets you view base object information and manage database object dependencies for a synonym.

### To edit a synonym

- 1 Open an editor on the synonym. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you view the owning <b>Schema</b> , the <b>Name</b> of the synonym, and the synonym's <b>Referenced Object Owner</b> , <b>Referenced Object Type</b> , and the <b>Referenced Object Name</b> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## TABLES EDITOR (IBM DB2 Z/OS)

The Tables Editor lets you: manage columns, basic properties, partitions, indexes, and constraints for a table.

**NOTE:** Before editing tables, refer to the material in [Altering Tables for IBM DB2 z/OS](#).

### To edit a table

- 1 Open an editor on the table. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Displays the currently defined columns in the table. For each column, the Property/Value list displays the <b>Name</b> and <b>Type</b> for the column. In addition, depending on the <b>Type</b> selected, the list also displays <b>Size</b> , <b>Scale</b> , <b>Identity Column</b> , <b>Allow Nulls</b> , <b>Default Value</b> , <b>Comment</b> , and <b>For Data</b> property values, as appropriate. The <b>Statistics</b> group displays a <b>Number of Distinct Values in the Column</b> value. Optionally, you can:	
		Select a column, and in the <b>Property/Value</b> list modify property values for that column.
		Click <b>Add Column</b> , provide a name for the new column, and set property values for the column.
		Select a column and click <b>Delete</b> to remove the column from the table.
<b>Properties</b>	Lets you work with settings in the following categories:	
	<b>Table</b>	Lets you view <b>EditProcedure</b> , <b>Table Type</b> , <b>Table Status</b> , <b>Check Flag</b> , <b>DBID</b> , and <b>OBID</b> properties. Lets you set <b>Volatile</b> , <b>Audit</b> , <b>RestrictDrop</b> , <b>Label</b> , and <b>ValidProc</b> properties.
	<b>Tablespace Placement</b>	Lets you select a <b>Tablespace</b> .
	<b>Log Options</b>	Lets you select a <b>Data Capture</b> option of DATA CAPTURE NONE or DATA CAPTURE CHANGES.
	<b>Statistics</b>	Lets you view the following statistics: <b>Last Runstats</b> , <b>Total number of rows</b> , <b>Average Row Length</b> , <b>Number of Pages</b> , <b>Percent Compressed Rows</b> , <b>Max Record Length</b> , and <b>DASD storage</b> .
<b>Partitions</b>	Displays existing partition columns and data partitions. Optionally you can add, edit, or delete partition columns and data partitions.	
<b>Indexes</b>	Displays the list of indexes for the table. Optionally, take one of the following actions:	
		Click <b>Add</b> to open a dialog that lets you add a new index to the table.
		Select an index and click <b>Edit</b> to open a dialog that lets you edit index properties.
		Select an index and click <b>Drop</b> to open a dialog that lets you remove the index from the table.

Tab	Settings and tasks
<b>Constraints</b>	Displays constraints in a tree structure. The tree contains folders which contain all constraints associated with the target table. The objects are organized in folders based on the type of constraint. Optionally take one of the following actions:
	Select a constraint type folder and click <b>Add</b> to open a dialog that lets you add a constraint of that type.
	Select a constraint and click <b>Edit</b> to open a dialog that lets you modify the constraint details.
	Select a constraint and click <b>Drop</b> to remove the constraint.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object.</a>
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies.</a>
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions.</a>
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object.</a>

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes.](#)

#### ALTERING TABLES FOR IBM DB2 Z/OS

The ALTER TABLE command of Transact-SQL is limited to adding NULL columns to the end of a table and to adding or dropping constraints. Unfortunately, this scenario does not address many requirements of administrators and developers who need to add, delete or modify columns more broadly:

- Add columns anywhere in a table, not just the end
- Add columns that do not permit a NULL value.
- Change the NULL/NOT NULL status of table columns
- Change column datatypes to other compatible types
- Change the length of datatypes
- Delete a column

Due to the limitations of the ALTER TABLE command, the only way to make broader modifications is to write SQL scripts that step through all desired changes. To perform an enhanced table alter, Rapid SQL constructs an SQL script that completes the following steps:

- 1 Renames the existing table so that the original and its data remain intact
- 2 Builds a CREATE TABLE statement with the new table definition, including declared defaults, primary key and check constraints
- 3 Builds an INSERT statement to copy data from the original, renamed table to the new one
- 4 Builds foreign keys on the new table
- 5 Reapplies any privileges granted on the table

- Rebuilds all dependencies on the new table, including indexes, triggers, procedures, packages, functions and views. When rebuilding procedures, functions, packages and views, Rapid SQL also rebuilds any permissions on them

## TABLESPACES EDITOR (IBM DB2 Z/OS)

The Tablespaces Editor lets you work with the basic properties and partitions for a tablespace as well as view space details, status, and objects stored on the tablespace.

### To edit a tablespace

- Open an editor on the tablespace. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	Lets you work with properties in the following categories:	
	<b>Bufferpool</b>	<b>Buffer Pool</b>
	<b>Partitions and size</b>	<b>Number of partitions</b> , Partition size (DSSIZE), Segment size, and Max rows per page
	<b>Other parameters</b>	<b>GBPCACHE</b> , <b>Compress</b> , <b>Track modified pages</b> , <b>Encoding scheme</b> , <b>Log</b> , <b>Define</b> , <b>Member Cluster</b> , <b>Close rule</b> , <b>Lock Size</b> , and <b>Maximum locks</b>
	For more information on these properties, see <a href="#">Tablespaces (DB2 z/OS) - Properties</a> .	
<b>Partitions</b>	Displays a list of partitions for the tablespace. Details for each partition include the storage group, VCAT, primary and secondary space allocations, the free space portion of each page and free space frequency, the group buffer cache scheme, and whether modifications are tracked.  Optionally, you can:	
		Select a partition and click <b>Edit</b> to open a dialog that lets you modify properties for that partition.
		Select a partition and click <b>Clone</b> to open a dialog that lets you apply the attributes of the selected partition to another partition.
<b>Status</b>	Lets you display CLAIMERS, LOCKS, LPL, USE, or WEPR status details	
<b>Space</b>	Lets you view space usage and allocation details for the tablespace.	
<b>Objects</b>	Displays the objects stored on the tablespace. Objects are organized in a tree structure with folders containing the objects.  Optionally, you can:	
		Select an object and click <b>Edit</b> to open an object editor on that object.
		Select an object and click <b>Drop</b> to initiate dropping that object.
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .	
<b>DDL</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## TRIGGERS EDITOR (IBM DB2 Z/OS)

The Triggers Editor lets you modify the CREATE TRIGGER statement and manage properties for a trigger.

### To edit a trigger

- 1 Open an editor on the trigger. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you view the <b>Trigger Timing</b> , <b>Trigger Events</b> , <b>Trigger Type</b> , <b>Object Status</b> , <b>Definer</b> , and <b>Function Path</b> properties.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Definition</b>	Lets you modify the CREATE TRIGGER body for a trigger. To modify a trigger, Rapid SQL must drop then create the trigger.

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## UNIQUE KEYS EDITOR (IBM DB2 Z/OS)

The Unique Keys Editor lets you manage columns, basic properties, and partitions for a unique key, as well as view storage details.

**To edit a unique key**

- 1 Open an editor on the unique key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Column</b>	Lets you manage columns that make up the primary key. On opening, this tab shows the existing columns. For each column, the listing shows the datatype (and if applicable the precision in brackets) and whether the table definition permits nulls in the target table column. Optionally you can:	
		Change the <b>Sort</b> order of a column.
		Click the New button and select a column name from the dropdown, to add a column to the primary key.
		Select a column and click the Drop button to delete the column from the primary key.
		Select a column and use the arrow buttons to reorder the columns in the primary key.
<b>Properties</b>	Lets you set the <b>Buffer Pool</b> , <b>Piece Size</b> , <b>Close</b> , and <b>Copy</b> properties.	
<b>Storage</b>	Lets you view the <b>Storage Group</b> and <b>VCAT Catalog</b> properties. Lets you set the <b>Primary Space Allocation</b> , <b>Secondary Space Allocation</b> , <b>Erase</b> , <b>Free Page</b> , <b>Percent Free</b> , and <b>GBP Cache</b> properties.	
<b>Partition</b>	Lets you work with partitions for the index.	
<b>Space</b>	Lets you view values in the following property groups:	
	<b>Attributes</b>	<b>Make Unique</b> , <b>System Required</b> , <b>Total Keys</b> , <b>Page Fetch Pairs</b> , and <b>Distinct Keys</b>
	<b>Statistics</b>	<b>Index Level</b> , <b>Cluster Ratio</b> , <b>Cluster Factor</b> , <b>Leaf Pages</b> , <b>Sequential Pages</b> , and <b>Density</b>
	<b>Cards</b>	<b>First Key</b> and <b>Full Keys</b>
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**USER DATATYPES EDITOR (IBM DB2 Z/OS)**

The User Datatypes Editor lets you manage basic properties of a user datatype.

**To edit a user datatype**

- 1 Open an editor on the user datatype. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	Lets you work with properties in the following categories:	
	<b>Base Datatype</b>	Lets you set <b>Type</b> and <b>Size</b> settings.
	<b>Character Options</b>	Lets you set the <b>For Data</b> and <b>CCSID</b> properties.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .	
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**USERS EDITOR (IBM DB2 Z/OS)**

The Users Editor lets you manage permissions for a user and the objects owned by that user.

**To edit a user**

- 1 Open an editor on the user. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	Displays the user <b>Name</b> .	
<b>Objects</b>	Lets you manage database objects associated with the user. Objects are organized in a tree structure with folders containing the objects. Optionally, take one of the following actions:	
		Select an object and click <b>Edit</b> to open an object editor on the selected object.
		Select an object and click <b>Drop</b> to initiate dropping the selected object.
<b>Object Permissions and System Permissions</b>	For details on using these tabs, see <a href="#">Working with Privileges and Permissions</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## VIEWS EDITOR (IBM DB2 Z/OS)

The Views Editor lets you view columns for a view and work with the dependencies and permissions for the view.

### To edit a view

- 1 Open an editor on the view. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Displays all columns for the view. Details for each column include the <b>Column Name</b> , the <b>Datatype</b> (and if applicable, with the precision in parentheses), and whether or not <b>Nulls</b> are allowed for that column.
<b>Definition</b>	Lets you view and modify the CREATE VIEW DDL that will implement any changes you make in this editor. To modify a view, Rapid SQL must drop then create the view.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## INTERBASE/FIREBIRD OBJECT EDITORS

Rapid SQL includes an Object Editor for all supported InterBase/Firebird objects. To see an Editor for a specific object, click the corresponding link below:

- [Blob Filters Editor \(InterBase/Firebird\)](#)
- [Domains Editor \(InterBase/Firebird\)](#)
- [Encryption Keys editor \(InterBase/Firebird\)](#)
- [Exceptions Editor \(InterBase/Firebird\)](#)
- [External Functions editor \(InterBase/Firebird\)](#)
- [Foreign Keys editor \(InterBase/Firebird\)](#)
- [Generators editor \(InterBase/Firebird\)](#)
- [Indexes editor \(InterBase/Firebird\)](#)
- [Primary Keys editor \(InterBase/Firebird\)](#)
- [Procedures editor \(InterBase/Firebird\)](#)
- [Roles editor \(InterBase/Firebird\)](#)



- [Shadows editor \(InterBase/Firebird\)](#)
- [Tables editor \(InterBase/Firebird\)](#)
- [Triggers editor \(InterBase/Firebird\)](#)
- [Unique Keys editor \(InterBase/Firebird\)](#)
- [Users editor \(InterBase/Firebird\)](#)
- [Views editor \(InterBase/Firebird\)](#)

## BLOB FILTERS EDITOR (INTERBASE/FIREBIRD)

The Blob Filters editor lets you modify the input and output types, entry point, and the module name of a blob filter declaration.

### To edit a blob filter

- 1 Open an editor on the blob filter. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	In addition to displaying the <b>Name</b> property, this tab also lets you modify the <b>Input subtype</b> , <b>Output subtype</b> , <b>Entrypoint</b> , and <b>Module name</b> properties. For details on these properties, see <a href="#">Blob Filters Editor (InterBase/Firebird)</a> .
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## DOMAINS EDITOR (INTERBASE/FIREBIRD)

The Domains editor lets you modify datatype details and other basic clause/option values of a domain.

**To edit a domain**

- 1 Open an editor on the domain. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	In addition to displaying the <b>Datatype</b> name, this tab lets you work with the <b>Type</b> , <b>Allow Nulls</b> , <b>Size</b> , Precision, <b>Scale</b> , <b>Array</b> , Character Set, Collation, <b>LOB Segment Length</b> , <b>Default</b> , and <b>Check</b> settings. For details on these settings and their availability, see <a href="#">Domains (InterBase/Firebird) - Properties</a> .
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**ENCRYPTION KEYS EDITOR (INTERBASE/FIREBIRD)**

The Encryption Keys editor lets you modify the encryption algorithm and key length options for a key, as well as the padding and Cipher Block Chaining versus Electronic Cookbook details.

**To edit an encryption key**

- 1 Open an editor on the encryption key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	In addition to displaying the <b>Name</b> of the encryption key, this tab lets you work with the <b>Algorithm</b> , <b>Pad random</b> , <b>Init Vector random</b> , and <b>IsDefault</b> settings. This tab also lets you modify the associated <b>Password</b> . For details on these settings, see <a href="#">Encryption Keys (InterBase/Firebird) - Properties</a> .
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## EXCEPTIONS EDITOR (INTERBASE/FIREBIRD)

The Exceptions Editor lets you modify the text of the message associated with an exception.

### To edit an exception

- 1 Open an editor on the exception. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	In addition to displaying the <b>Name</b> of the exception, this tab lets you work with the <b>Text</b> setting. For details on these settings, see <a href="#">Exceptions (InterBase/Firebird) - Properties</a> .
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## EXTERNAL FUNCTIONS EDITOR (INTERBASE/FIREBIRD)

The External Functions editor lets you modify basic properties, input parameters, and return datatype of an external function.

### To edit an external function

- 1 Open an editor on the external function. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	In addition to displaying the <b>Name</b> of the external function, this tab lets you work with the <b>EntryPoint</b> and <b>Module Name</b> settings. For details on these settings, see <a href="#">External Functions (InterBase/Firebird) - Properties</a> .
<b>Parameters</b>	Displays the input parameters to the external function. Optionally, you can:
	Use the New button to add a new parameter and type a name for the parameter in the space provided. With the parameter selected, in the <b>Attributes</b> area select a <b>Type</b> , and if appropriate, provide or select <b>Precision</b> , <b>Scale</b> , and <b>Size</b> options.
	Modify a parameter by selecting it, and in the <b>Attributes</b> area editing the <b>Type</b> , <b>Precision</b> , <b>Scale</b> , and <b>Size</b> options.
	Use the Delete button to drop a selected parameter and use the arrow keys to reorder the parameter list.
	Use the arrow keys to change the list ordering of a selected parameter.

Tab	Settings and tasks
Return Type	Displays datatype details of the value returned by the external function.
Dependencies	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## FOREIGN KEYS EDITOR (INTERBASE/FIREBIRD)

The Foreign Keys Editor lets you manage column mapping for a foreign key as well as modify the update and delete rule actions.

### To edit a foreign key

- 1 Open an editor on the foreign key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Column Mapping</b>	The existing column mapping for the foreign key is represented by selected columns in the <b>Main Table</b> and <b>Referenced Table</b> lists. Additional candidates in the <b>Referenced Table</b> list are indicated by enabled column check boxes.  If necessary, use the <b>Table</b> dropdown in the <b>Referenced Table</b> group to choose a new table for this foreign key. Select or deselect columns in the <b>Main Table</b> list and <b>Referenced Table</b> list to form the referential constraint between the two tables.
<b>Properties</b>	Lets you modify the <b>Delete rule</b> and <b>Update Rule</b> actions. For details on these properties, see <a href="#">Foreign keys (InterBase/Firebird) - Properties</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## GENERATORS EDITOR (INTERBASE/FIREBIRD)

The Generators Editor lets you view the name and work with the object dependencies for a generator.

**To edit a generator**

- 1 Open an editor on the generator. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you view the <b>Name</b> of the generator and modify the <b>Current Value</b> . For details on these properties, see <a href="#">Generators (InterBase/Firebird) - Properties</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**INDEXES EDITOR (INTERBASE/FIREBIRD)**

The Indexes Editor lets you modify the columns and basic properties for an index.

**To edit an index**

- 1 Open an editor on the index. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Columns</b>	Lets you manage columns that make up the index. On opening, this tab shows the existing columns. For each column, the listing displays the datatype and whether the table definition permits nulls.  Optionally you can:
	Click the New button and select a column name from the dropdown, to add a column to the index.
	Select a column and click the Drop button to delete the column from the index.
	Select a column and use the arrow buttons to reorder the columns in the index.
<b>Properties</b>	In addition to displaying the <b>Name</b> and associated <b>Table Name</b> for the index, this tab lets you modify the <b>Unique</b> , <b>Enabled</b> , and <b>Descending</b> properties. For details on these properties, see <a href="#">Indexes (InterBase/Firebird) - Properties</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PRIMARY KEYS EDITOR (INTERBASE/FIREBIRD)

The Primary Keys Editor lets you modify the columns that make up a primary key.

### To edit a primary key

- 1 Open an editor on the primary key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Lets you manage columns that make up the primary key. On opening, this tab shows the existing columns. For each column, the listing displays the datatype and whether the table definition permits nulls.  Note that only columns defined with the <b>Allow Nulls</b> property disabled can be added to a primary key.  Optionally you can:	
		Click the New button and select a column name from the dropdown, to add a column to the index.
		Select a column and click the Drop button to delete the column from the index.
		Select a column and use the arrow buttons to reorder the columns in the index.
<b>Properties</b>	Displays the <b>Name</b> of the index and the associated <b>Table Name</b> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PROCEDURES EDITOR (INTERBASE/FIREBIRD)

The Procedures Editor lets you modify the body and parameters for a procedure.

### To edit a procedure

- 1 Open an editor on the procedure. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Displays the <b>Name</b> of the procedure.

Tab	Settings and tasks
<b>Parameters</b>	Lets you manage procedure parameters. On opening, this tab shows the existing parameters. Optionally you can:
	Select a parameter from the list, modify the <b>Type</b> , and if appropriate, the <b>Precision</b> , <b>Size</b> , and <b>Scale</b> of the parameter. You can also set the <b>Parameter Mode</b> to INPUT or OUTPUT.
	Click the New button, provide a name for the new parameter and modify its attributes.
	Select a parameter and click the Delete button to delete the parameter.
	Select a parameter and use the arrow buttons to reorder the parameter list.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>Body</b>	Lets you modify the body of the procedure.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## ROLES EDITOR (INTERBASE/FIREBIRD)

The Roles Editor lets you manage the users assigned to a role and the permissions associated with the role.

### To edit a role

- Open an editor on the role. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Displays the <b>Name</b> and <b>Authorization Owner</b> for the role.
<b>Users</b>	Let you manage users for the role. Use the <b>Join Role</b> and <b>Leave Role</b> buttons to move users between the <b>Users In Role</b> and <b>Users Not In Role</b> lists.
<b>Object Permissions</b>	For details on using these tabs, see <a href="#">Working with Privileges and Permissions</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## SHADOWS EDITOR (INTERBASE/FIREBIRD)

The Shadows Editor lets you modify CONDITIONAL, AUTO, and MANUAL argument usage and the file specification for a shadow.

**To edit a shadow**

- 1 Open an editor on the shadow. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Displays the <b>Shadow Set</b> number and lets you modify the <b>Conditional</b> and <b>Behavior</b> properties. For details on these properties, see <a href="#">Shadows (InterBase/Firebird) - Properties</a> .
<b>Storage</b>	Use the <b>Delete</b> button to drop a selected file. Modify an existing file by selecting the file in the primary and secondary files list, and modify the <b>Physical File Name</b> , <b>Starting Page</b> (for secondary files), and <b>Size</b> properties. Add a new file by clicking <b>New</b> , and providing file specification values in the <b>Property/Value</b> list.
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**TABLES EDITOR (INTERBASE/FIREBIRD)**

The Tables Editor lets you manage basic properties, columns, indexes, and constraints for a table.

**To edit a table**

- 1 Open an editor on the table. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Columns</b>	Displays the currently defined columns in the table. Optionally, you can:
	Select a column, and in the <b>Property/Value</b> list modify property values for that column. For details on the properties, see <a href="#">Tables (InterBase/Firebird) - Columns</a> .
	Use the <b>Add Column</b> dropdown to add a new column to the bottom of the column list or to insert a new column above the currently selected column in the column list. Provide a <b>Name</b> for the column, and then set property values for the column.
	Select a column and click <b>Delete</b> to remove the column from the table.
<b>Properties</b>	Displays the <b>Name</b> of the table and the date and time <b>Created</b> .



Tab	Settings and tasks
<b>Indexes</b>	Lets you manage indexes for a table:
	Click <b>Add</b> to open a dialog that lets you add a new index to the table. For more information, see <a href="#">Indexes Wizard (InterBase/Firebird)</a> .
	Select an index and click <b>Edit</b> to open a dialog that lets you edit index properties. For more information, see <a href="#">Indexes editor (InterBase/Firebird)</a> .
	Select an index and click <b>Drop</b> to open a dialog that lets you remove the index from the table.
<b>Constraints</b>	Lets you manage primary key, unique key, foreign key, and check constraints for a table. Constraints are grouped by type, under folders:
	Select a constraint type folder and click <b>Add</b> to open a dialog that lets you add a constraint of that type.
	Select a constraint and click <b>Edit</b> to open a dialog that lets you modify the constraint details.
	Select a constraint and click <b>Drop</b> to remove the constraint.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## TRIGGERS EDITOR (INTERBASE/FIREBIRD)

The Triggers Editor lets you modify the basic properties, trigger body, and dependencies for a trigger.

### To edit a trigger

- Open an editor on the trigger. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Displays the <b>Parent Type</b> , <b>Parent Name</b> , and <b>Name</b> properties for the trigger, and lets you modify the <b>Trigger Timing</b> , <b>Trigger Events</b> , <b>Position</b> , and <b>Enabled</b> properties. For details on these properties, see <a href="#">Triggers (InterBase/Firebird) - Properties</a> .
<b>Body</b>	Lets you work with the trigger body.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## UNIQUE KEYS EDITOR (INTERBASE/FIREBIRD)

The Unique Keys editor lets you modify the columns that make up a unique key.

### To edit a unique key

- 1 Open an editor on the unique key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Lets you manage columns that make up the unique key. On opening, this tab shows the existing columns. For each column, the listing includes the datatype and whether the table definition permits nulls.  Note that only columns defined with the <b>Allow Nulls</b> property disabled can be added to a unique key.  Optionally you can:	
		Click the New button and select a column name from the dropdown, to add a column to the unique key.
		Select a column and click the Drop button to delete the column from the unique key.
		Select a column and use the arrow buttons to reorder the columns in the unique key.
<b>Properties</b>	Displays the <b>Name</b> of the unique key and the associated <b>Table Name</b> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## USERS EDITOR (INTERBASE/FIREBIRD)

The Users Editor lets you modify basic properties, assigned roles, and object permissions for a user.

**To edit a user**

- 1 Open an editor on the user. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	In addition to displaying the login <b>Name</b> for the user, this tab lets you work with the <b>Password, Group, Default Role, First Name, Middle Name, Last Name, Active, System User Name, Description, GID, and UID</b> properties. For details on these properties, see <a href="#">Users (InterBase/Firebird) - Properties</a> .
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>Roles</b>	Lets you assign the valid, non-default roles that a user can specify when logging in to the database. A default login role can be assigned for this user on the Properties tab/panel. For details, see <a href="#">Users (InterBase/Firebird) - Roles</a> . To assign roles to the user, select the associated check box.
<b>Object Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**VIEWS EDITOR (INTERBASE/FIREBIRD)**

The Views editor lets you work with the CREATE VIEW DDL for a view, grant and revoke associated permissions, and work with the view's dependencies.

**To edit a view**

- 1 Open an editor on the view. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Displays the <b>Name</b> of the view and a listing of the associated columns. Details for each column include the column <b>Name</b> , the <b>Datatype</b> , whether or not <b>Nulls</b> are allowed for that column, and an associated <b>Comment</b> .
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>Definition</b>	Lets you view and modify the CREATE VIEW DDL for the view.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## MICROSOFT SQL SERVER OBJECT EDITORS

Rapid SQL includes an Object Editor for all supported Microsoft SQL Server objects. To see an Editor for a specific object, click the corresponding link below:

- [Check Constraints Editor \(SQL Server\)](#)
- [Databases Editor \(SQL Server\)](#)
- [Defaults Editor \(SQL Server\)](#)
- [Extended Procedures Editor \(SQL Server\)](#)
- [When finished, you can submit your changes. For details, see Previewing and Submitting Object Editor Changes. Foreign Keys Editor \(SQL Server\)](#)
- [Functions Editor \(SQL Server\)](#)
- [Indexes Editor \(SQL Server\)](#)
- [Logins Editor \(SQL Server\)](#)
- [Primary Keys Editor \(SQL Server\)](#)
- [Procedures Editor \(SQL Server\)](#)
- [Rules Editor \(SQL Server\)](#)
- [Schemas Editor \(SQL Server\)](#)
- [Tables Editor \(SQL Server\)](#)
- [Triggers Editor \(SQL Server\)](#)
- [Unique Keys Editor \(SQL Server\)](#)
- [Users Editor \(SQL Server\)](#)
- [User Datatypes Editor \(SQL Server\)](#)
- [User Messages Editor \(SQL Server\)](#)
- [Views Editor \(SQL Server\)](#)

### CHECK CONSTRAINTS EDITOR (SQL SERVER)

The Check Constraints Editor lets you view and modify check constraints properties and edit the check constraint expression.

**To edit a check constraint**

- 1 Open an editor on the check constraint. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Definition</b>	Use the <b>Enabled</b> check box to enable or disable the check constraint. Use the <b>Not For Replication</b> check box to enable or disable the check constraint for replication.  You can also edit the condition in the <b>Check Condition</b> box. The <b>Table Column</b> button opens a dialog that lets you select and paste column names into the check condition expression.
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**DATABASES EDITOR (SQL SERVER)**

The Databases Editor lets you manage basic properties, log and data files for a database.

**To edit a database**

- 1 Open an editor on the database. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
Options	Lets you work with options in the following categories:	
	Creation Properties	<p>Lets you choose values for the <b>Compatible Level</b> and <b>Owner</b> settings.</p> <p><b>NOTE:</b> The Compatible level option sets certain database behaviors to be compatible with the specified earlier version of Microsoft® SQL Server. The compatibility level affects the behaviors in the specified database, not the entire server.To set database options for all future databases, set the database options on the model database.</p> <p><b>NOTE:</b> When changing the database owner (dbo) select the check box to transfer the existing aliases of users who could act as the old dbo (including their permissions) to the new dbo.</p>
	Properties	<p>Lets you set the following properties: <b>ANSI null default, ANSI nulls, ANSI padding, ANSI warnings, arithabort, auto create statistics, auto update statistics, auto close, auto shrink, concat null yields nul, cursor close on commit, db chaining, dbo use only, default to local cursor, merge publish, numeric roundabout, offline, published, quoted identifier, read only, recursive triggers, select into/bulkcopy/pll sort, single user, subscribed, torn page detection, and trunc log on chkpt.</b></p>

Tab	Settings and tasks	
<b>Placement</b>	Displays the currently defined data files for the database. Optionally, you can:	
		Select a file from the list on the left and in the <b>Device File Properties</b> group, modify the <b>Size</b> , <b>File Growth Rate</b> , <b>Max Size</b> , and <b>Unlimited Max Size</b> settings.
		Click the <b>New</b> button, provide a <b>Device File Name</b> for the new file, and use the other settings in the <b>Device File Properties</b> group to provide additional details for the file.
		Select a file from the list on the left and click <b>Delete</b> to remove the file.
<b>Transaction Log</b>	Displays the currently defined transaction logs for the database. Optionally, you can:	
		Select a transaction log from the list on the left and in the <b>Log Device Properties</b> group, modify the <b>Size</b> , <b>File Growth Rate</b> , <b>Max Size</b> , and <b>Unlimited Max Size</b> settings.
		Click the <b>New</b> button, provide a <b>Device File Name</b> for the new transaction log, and use the other settings in the <b>Log Device Properties</b> group to provide additional details for the transaction log.
		Select a transaction log from the list on the left and click <b>Delete</b> to remove the transaction log.
<b>Space</b>	Lets you view pie charts showing the data space usage and the transaction log (if available) space usage for the database.	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## DEFAULTS EDITOR (SQL SERVER)

The Defaults Editor lets you change the owner and value of a default.

### To edit a default

- Open an editor on the default. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you modify the <b>Value</b> of the Default and the owning <b>Schema</b> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## EXTENDED PROCEDURES EDITOR (SQL SERVER)

The Procedures Editor lets you modify the library name for an extended procedure. It also lets you work with dependencies and permissions.

**NOTE:** Extended Procedures are only available on the master database.

### To edit an extended procedure

- 1 Open an editor on the extended procedure. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you provide a dynamic-link library (DLL) <b>Library Name</b> for the extended procedure.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## FOREIGN KEYS EDITOR (SQL SERVER)

The Foreign Keys Editor lets you manage the columns and basic properties of a foreign key.

### To edit a foreign key

- 1 Open an editor on the foreign key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Column Mapping</b>	The existing column mapping for the foreign key is represented by selected columns in the <b>Main Table</b> and <b>Referenced Table</b> lists. Additional candidates in the <b>Referenced Table</b> list are indicated by enabled column check boxes.  If necessary, use the <b>Table</b> dropdown in the <b>Referenced Table</b> group to choose a new table for this foreign key. Select or deselect columns in the <b>Main Table</b> list and <b>Referenced Table</b> list to form the referential constraint between the two tables.
<b>Properties</b>	Lets you specify <b>Enabled</b> and <b>Not For Replication</b> properties. This tab also lets you select a <b>Delete rule</b> and an <b>Update rule</b> (NONE, SET NULL, SET DEFAULT, CASCADE).
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## FUNCTIONS EDITOR (SQL SERVER)

The Functions Editor lets you view and modify function definitions and dependencies.

### To edit a function

- 1 Open an editor on the function. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you modify the owning <b>Schema</b> of the function.
<b>Definition</b>	Lets you view and modify the dynamic-link library (DLL) or view the data definition language (DDL).
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## INDEXES EDITOR (SQL SERVER)

The Indexes Editor lets you manage columns and basic properties of an index and view index statistics.

### To edit an index

- 1 Open an editor on the index. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Columns</b>	Lets you manage columns that make up the index. On opening, this tab shows the existing columns. For each column, the datatype (and if applicable the precision in brackets) and whether the table definition permits nulls in the target table column. Optionally you can:
	Change the <b>Sort</b> order of a column.
	Click the New button and select a column name from the dropdown, to add a column to the index.
	Select a column and click the Drop button to delete the column from the index.
	Select a column and use the arrow buttons to reorder the columns in the index.



Tab	Settings and tasks	
<b>Properties</b>	Lets you work with settings in the following categories:	
	<b>Creation</b>	In addition to displaying identifying information, you can modify the <b>Build Online</b> and <b>Max degree of parallelism</b> properties. For details on these properties, see <a href="#">Index Wizard (SQL Server)</a> .
	<b>Attributes</b>	Lets you set <b>Index Type</b> , <b>Clustered</b> , <b>Ignore Duplicate Key</b> (for <b>Index Type</b> of UNIQUE), <b>Statistics Recompute</b> , <b>Allow Row Locks</b> , and <b>Allow Page Locks</b> properties. For details on these properties, see <a href="#">Index Wizard (SQL Server)</a> .  <b>NOTE:</b> You cannot reorganize an index (primary key, or unique key) that has an <b>Allow Page Locks</b> property set to FALSE. For information on reorganizing indexes, see <a href="#">Reorganizing SQL Server indexes, primary keys, and unique keys</a> .
	<b>Storage</b>	Lets you set <b>Fill Factor</b> , <b>File Group</b> , <b>Pad Index</b> , and <b>Sort in Tempdb</b> settings.
<b>Statistics</b>	Lets you view statistics in the following categories:	
	<b>Page Statistics</b>	<b>Data Pages</b> , <b>Reserved Pages</b> , <b>Used Pages</b> , and <b>Total Pages Modified</b> .
	<b>Row Statistics</b>	<b>Maximum Row Size</b> , <b>Minimum Row Size</b> , <b>Max Size of Non-Leaf Index Row</b> , and <b>Total Rows Modified</b> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## LOGINS EDITOR (SQL SERVER)

The Logins Editor lets you manage basic properties for a login and manage associated users and roles.

### To edit a login

- Open an editor on the login. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	Lets you work with settings in the following categories:	
	<b>General</b>	Select a new <b>Default Database</b> or <b>Default Language</b> . Master should not be specified as the default database in order to prevent users from creating objects there by mistake.
	<b>Authentication</b>	<b>Password</b> - For STANDARD account types, you can modify the password. <b>Check Policy</b> , <b>Check expiration</b> , and <b>Must change</b> - For STANDARD account types, use these properties to modify the corresponding SQL Server Enforce password policy and Enforce password expiration properties. <b>Asymmetric Key</b> - For ASYMMETRIC KEY account types, you can select a new asymmetric key.

Tab	Settings and tasks
<b>Users</b>	Lets you manage each login's user account(s). Before a login name can access a user database, it must be added as a user or alias to that database by the system administrator or database owner. Use the <b>Add</b> button to add a selected database from the <b>Databases where the login does NOT have a user account</b> to the <b>Databases where the login HAS a user account</b> list. Conversely, remove a login from a database by using the <b>Remove</b> button to move a selected database from the <b>Databases where the login HAS a user account</b> list.
<b>Server Roles</b>	Lets you select the defined roles that are to be assigned to this login.
<b>DDL</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PRIMARY KEYS EDITOR (SQL SERVER)

The Primary Keys Editor lets you manage the columns and basic properties for a primary key and lets you view page and row statistics.

### To edit a primary key

- Open an editor on the primary key. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Columns</b>	Lets you manage columns that make up the primary key. On opening, this tab shows the existing columns. For each column, the listing shows the datatype (and if applicable the precision in brackets) and whether the table definition permits nulls in the target table column. Optionally you can:
	Change the <b>Sort</b> order of a column.
	Click the New button and select a column name from the dropdown, to add a column to the primary key.
	Select a column and click the Drop button to delete the column from the primary key.
	Select a column and use the arrow buttons to reorder the columns in the primary key.

Tab	Settings and tasks	
<b>Properties</b>	Lets you work with settings in the following categories:	
	<b>Creation</b>	In addition to displaying identifying information, you can modify the <b>Build Online</b> and <b>Max degree of parallelism</b> properties. For details on these properties, see <a href="#">Primary Key Wizard (SQL Server)</a> .
	<b>Attributes</b>	<b>Clustered</b> - Indicates whether the target primary key is clustered. <b>Ignore Duplicate Key</b> - Indicates whether the target primary key ignores duplicate key values. If you select this option, the transaction that generated the duplicate key values can continue. <b>Statistics Recompute</b> - Indicates that index statistics are automatically recomputed as the index is updated. Microsoft does not recommend this.
	<b>Storage</b>	<b>File Group</b> - Lets you select a file group. <b>Fill Factor</b> - Lets you specify the fill factor that specifies how full each index page can be. If no fill factor is specified, Microsoft SQL Server uses the database's default fill factor. <b>Pad Index</b> - If you specified a Fill factor of more than 0 percent, and you selected the option to create a unique index, you can specify to use the same percentage you specified in Fill Factor as the space to leave open on each interior node. By default, Microsoft SQL Server sets a 2 row index size.
<b>Statistics</b>	Lets you view statistics in the following categories:	
	<b>Page Statistics</b>	<b>Data Pages, Reserved Pages, Used Pages, and Total Pages Modified.</b>
	<b>Row Statistics</b>	<b>Maximum Row Size, Minimum Row Size, Max Size of Non-Leaf Index Row, and Total Rows Modified.</b>
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PROCEDURES EDITOR (SQL SERVER)

The Procedures Editor lets you manage properties, the definition, dependencies, and permissions for a procedure.

### To edit a procedure

- Open an editor on the procedure. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you change the owning <b>Schema</b> .
<b>Definition</b>	Lets you view the SQL code for the procedure.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## RULES EDITOR (SQL SERVER)

The Rules Editor lets you view or modify basic properties of a rule.

### To edit a rule

- 1 Open an editor on the rule. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you specify whether the rule is <b>Enabled</b> and view the <b>Fire On Insert</b> , <b>Fire On Update</b> , <b>Fire On Delete</b> , and <b>Encrypted</b> settings.
<b>Definition</b>	Lets you view and modify the SQL code that will implement any changes you make using this editor.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## SCHEMAS EDITOR (SQL SERVER)

The Schemas Editor lets you change the owner of a group and manage the objects contained in the schema.

### To edit a user

- 1 Open an editor on the user. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you modify the <b>Owner</b> of the schema.
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>Objects</b>	Displays the specific objects owned by the schema, grouped under object type folders. Optionally, you can:
	Select an object and click <b>Edit</b> to open an object editor on that object.
	Select an object and click <b>Drop</b> to initiate dropping the object.
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## TABLES EDITOR (SQL SERVER)

The Tables Editor lets you manage columns, basic properties, indexes, and constraints for a table and view space usage details.

### To edit a table

- 1 Open an editor on the table. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Displays the currently defined columns in the table. Optionally, you can:	
		Select a column, and in the <b>Property/Value</b> list modify property values for that column.
		Click <b>Add Column</b> , provide a name for the new column, and set property values for the column.
		Select a column and click <b>Delete</b> to remove the column from the table.
<b>Properties</b>	Lets you work with settings in the following categories:	
	<b>Physical Storage</b>	Lets you specify a <b>Filegroup</b> and <b>Text Image Filegroup</b> . This is only available for Microsoft SQL Server 7.0 and later.
	<b>Full-Text Index</b>	Lets you view whether Full-Text Indexing is installed and active.  The full-text index feature provides support for sophisticated word searches in character string data. A full-text index stores information about significant words and their location within a given column. This information is used to quickly complete full-text queries that search for rows with particular words or combinations of words.  This feature is available for Microsoft SQL Server 8.0 or later.
	<b>Text In Row</b>	Lets you enable Text In Row Data and specify a limit.
<b>Indexes</b>	Displays any indexes currently defined for the table. Optionally, you can:	
		Click <b>Add</b> to open an Index editor, letting you create a new index for this table. For more information, see <a href="#">Index Wizard (SQL Server)</a> .
		Select an index and click <b>Edit</b> to open an wizard that lets you modify that index. For more information, see <a href="#">Indexes Editor (SQL Server)</a> .
		Select an index and click <b>Drop</b> to to open a dialog that lets you confirm that you want to drop the index. When dropping an index with the Clustered option enabled, the confirmation dialog includes an <b>Online</b> option that lets you specify an online drop (ONLINE=ON clause).

Tab	Settings and tasks	
<b>Constraints</b>	Lets you manage constraints for the table. Constraints are grouped by type, under folders. Optionally take one of the following actions:	
		Select a constraint type folder and click <b>Add</b> to open a dialog that lets you add a constraint of that type.
		Select a constraint and click <b>Edit</b> to open a dialog that lets you modify the constraint details.
		Select a constraint and click <b>Drop</b> to open a dialog that lets you confirm that you want to drop the index. When dropping a clustered unique key or clustered primary key, the confirmation dialog includes an <b>Online</b> option that lets you specify an online drop (ONLINE-ON clause).
<b>Space</b>	Lets you view pie charts showing the space usage for the table. Optionally you can double-click a slice in the pie chart for detailed statistics.	
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .	
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## TRIGGERS EDITOR (SQL SERVER)

The Triggers Editor lets you enable and disable a trigger, view and modify the CREATE TRIGGER statement that will implement changes, and manage dependencies.

### To edit a trigger

- 1 Open an editor on the trigger. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you set the <b>Enabled</b> property. It also lets you view the <b>Trigger Timing</b> (AFTER, INSTEAD OF), <b>Fire On Insert</b> , <b>Fire On Update</b> , <b>Fire On Delete</b> , and <b>Encrypted</b> properties.
<b>Definition</b>	Lets you modify the CREATE TRIGGER body for a trigger. To modify a trigger, edit the text of the trigger body in the Trigger Text area. Rapid SQL must drop then create the trigger.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## UNIQUE KEYS EDITOR (SQL SERVER)

The Unique Keys Editor lets you manage columns and properties for a unique key and view associated statistics.

### To edit a unique key

- 1 Open an editor on the unique key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Lets you manage columns that make up the unique key. On opening, this tab shows the existing columns. For each column, the listing displays the datatype (and if applicable the precision in brackets) and whether the table definition permits nulls in the target table column. Optionally you can:	
		Change the <b>Sort</b> order of a column.
		Click the New button and select a column name from the dropdown, to add a column to the index.
		Select a column and click the Drop button to delete the column from the index.
		Select a column and use the arrow buttons to reorder the columns in the index.
<b>Properties</b>	Lets you work with settings in the following categories:	
	<b>Creation</b>	In addition to displaying identifying information, you can modify the <b>Build Online</b> and <b>Max degree of parallelism</b> properties. For details on these properties, see <a href="#">Unique Key Wizard (SQL Server)</a> .
	<b>Attributes</b>	<b>Clustered</b> - Indicates whether the target index is clustered. <b>Ignore Duplicate Key</b> - Indicates whether the target primary key ignores duplicate key values. If you select this option, the transaction that generated the duplicate key values can continue. <b>Statistics Recompute</b> - Indicates that index statistics are automatically recomputed as the index is updated. Microsoft does not recommend this.
	<b>Storage</b>	<b>File Group</b> - Lets you specify the filegroup on which to place the index. This is for Microsoft SQL Server 7.0 or later. <b>Fill Factor</b> - Lets you specify a percentage that indicates how full Microsoft SQL Server should make the leaf level of each index page during index creation. When an index page fills up, Microsoft SQL Server must take time to split the index page to make room for new rows, which is quite expensive. For update-intensive tables, a properly chosen Fill factor value yields better update performance than an improper Fill factor value. <b>Pad Index</b> - If you specified a Fill factor of more than 0 percent, and you selected the option to create a unique index, you can specify to use the same percentage you specified in Fill Factor as the space to leave open on each interior node. By default, Microsoft SQL Server sets a 2 row index size.

Tab	Settings and tasks	
<b>Statistics</b>	Displays statistics in the following categories:	
	<b>Page Statistics</b>	<b>Data Pages, Pages Reserved, Used Pages, and Total Pages Modified.</b>
	<b>Row Statistics</b>	<b>Maximum Row Size, Minimum Row Size, Max Size of Non-Leaf Index Row, and Total Rows Modified Since Last.</b>
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## USERS EDITOR (SQL SERVER)

The Users Editor lets you manage properties, assign roles, and manage object bindings for a user.

### To edit a user

- Open an editor on the user. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	Lets you view the <b>Without Login, Login Name, Name, User Type, Certificate, Asymmetric Key, and Default Schema</b> properties.	
<b>Roles</b>	Lets you assign roles to a user by selecting the check boxes associated with the roles to be assigned.	
<b>Objects</b>	Lets you manage database objects (Defaults, Indexes, Procedures, Rules, Tables, Triggers, User Datatypes, Views) associated with the user. Objects are organized in a tree structure with folders containing the objects. Optionally, take one of the following actions:	
		Select an object and click <b>Edit</b> to open an object editor on the selected object. For information on each editor, use the links in <a href="#">Microsoft SQL Server Object Editors</a> .
		Select an object and click <b>Drop</b> to initiate dropping the selected object.
<b>Object Permissions and System Permissions</b>	For details on using these tabs, see <a href="#">Working with Privileges and Permissions</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).



## USER DATATYPES EDITOR (SQL SERVER)

The User Datatypes Editor lets you modify the base datatype, rule and default bindings, and referencing or referenced objects for a user datatype.

### To edit a user datatype

- 1 Open an editor on the user datatype. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	In addition to letting you change the <b>Schema</b> , this tab lets you work with settings in the following categories:	
	<b>Base Datatype</b>	Lets you modify the <b>Type</b> and <b>Allow Nulls</b> properties. Depending on the <b>Type</b> selected, <b>Precision</b> , <b>Scale</b> , or <b>Size</b> properties are available as appropriate.
	<b>Bindings</b>	<b>Default Binding</b> - Lets you select a default. For information on creating defaults, see <a href="#">Default Wizard (SQL Server)</a> . <b>Rule Binding</b> - Lets you select a rule. For information on creating rules, see <a href="#">Rule Wizard (SQL Server)</a> .
<b>Usage</b>	Lets you manage database objects referencing or referenced by a user datatype. Objects are grouped within object type folders. Optionally, you can:	
		Select an object and click <b>Edit</b> to open an object editor on the selected object.
		Select an object and click <b>Drop</b> to initiate dropping the selected object.
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## USER MESSAGES EDITOR (SQL SERVER)

**NOTE:** User Messages are only on the master database.

The User Messages Editor lets you view or modify basic properties of a user message. It also lets you add, edit, or delete individual language versions of the message.

**To edit a user message**

- 1 Open an editor on the user message. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	Lets you work with the following properties:	
	<b>Message Number</b>	Displays the message number chosen when the user message was created.
	<b>Severity</b>	Lets you modify the SQL Server error severity level (001-025).
	<b>Write to NT Event Log</b>	Lets you specify that the message is always written to the Windows NT Event Log.
<b>Information</b>	Displays the currently-defined language versions for this message number. Optionally you can:	
		Create a new language version for this message number. Click the <b>Add new text for the user message</b> button, and in the dialog box that opens, select a <b>Language</b> and provide the <b>Message Text</b> .
		Select a language version of the message and click the <b>Modify user message text</b> button. This opens a dialog that lets you modify the <b>Language</b> or <b>Message Text</b> .
		Select a language version of the message and click the <b>Remove user message text</b> button.
	NOTE: You cannot create two versions of a message for the same language. NOTE: For a given message number there should always be a least one, us_english version of the message.	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**VIEWS EDITOR (SQL SERVER)**

The Views Editor lets you view columns and change the schema of a view.

**To edit a view**

- 1 Open an editor on the view. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Displays a list of the columns that make up the view. Details include the column <b>Name</b> , its <b>Datatype</b> , and whether the definition allows <b>Nulls</b> .  This tab also lets you change the <b>Schema</b> owning the view.
<b>Definition</b>	Lets you view the SQL for the View.

Tab	Settings and tasks
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## ORACLE OBJECT EDITORS

Rapid SQL includes an Object Editor for all supported Oracle objects. To see an Editor for a specific object, click the corresponding link below.

**NOTE:** If an objects has dependent objects, such as tables, triggers, procedures and views, you can view and access their dependent objects in the editor.

- [Check Constraints Editor \(Oracle\)](#)
- [Clusters Editor \(Oracle\)](#)
- [Database Links Editor \(Oracle\)](#)
- [Foreign Keys Editor \(Oracle\)](#)
- [Functions Editor \(Oracle\)](#)
- [Indexes Editor \(Oracle\)](#)
- [Job Queue Editor \(Oracle\)](#)
- [Libraries Editor \(Oracle\)](#)
- [Materialized Views Editor \(Oracle\)](#)
- [Materialized View Logs Editor \(Oracle\)](#)
- [Outlines Editor \(Oracle\)](#)
- [Package Bodies Editor \(Oracle\)](#)
- [Packages Editor \(Oracle\)](#)
- [Primary Keys Editor \(Oracle\)](#)
- [Procedures Editor \(Oracle\)](#)
- [Profiles Editor \(Oracle\)](#)
- [Redo Log Groups Editor \(Oracle\)](#)
- [Roles Editor \(Oracle\)](#)
- [Rollback Segments Editor \(Oracle\)](#)

- [Sequences Editor \(Oracle\)](#)
- [Synonyms Editor \(Oracle\)](#)
- [Tables Editor \(Oracle\)](#)
- [Tablespaces Editor \(Oracle\)](#)
- [Triggers Editor \(Oracle\)](#)
- [Type Bodies Editor \(Oracle\)](#)
- [Types Editor \(Oracle\)](#)
- [Unique Keys Editor \(Oracle\)](#)
- [Users Editor \(Oracle\)](#)
- [Views Editor \(Oracle\)](#)

## CHECK CONSTRAINTS EDITOR (ORACLE)

The Check Constraints Editor lets you modify and enable/disable check constraints.

### To edit a check constraint

- 1 Open an editor on the check constraint. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Definition</b>	Use the <b>Enabled</b> control to enable/disable the check constraint.  You can also edit the condition in the <b>Check Condition</b> box. The <b>Table Columns</b> button acts a time saver in editing the condition. It opens a dialog that lets you select and paste column names into the check condition expression.
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## CLUSTERS EDITOR (ORACLE)

The Clusters Editor lets you manage cluster column and table details, view and modify storage and space parameters, and manage performance settings for a cluster.

**To edit a cluster**

- 1 Open an editor on the cluster. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	<p>Use the controls in the <b>General Properties</b> group to specify a hash or index <b>Cluster Type</b> and to specify a <b>Key Size</b> for the cluster.</p> <p>If you specify a hash cluster type, use the controls in the <b>Hash Specifications</b> group to specify the number of hash keys and to specify use of either the default hash function or a specified expression used as the hash function.</p>	
<b>Storage</b>	Lets you manage storage for a cluster:	
	<b>Data Block Storage</b>	<p>Each transaction that updates a data block requires a transaction entry.</p> <p><b>Percent Free</b> - minimum percentage of free space in a block</p> <p><b>Percent Used</b> - minimum percentage of used space in a block.</p> <p><b>Initial transactions</b> - The initial parameter ensures that a minimum number of concurrent transactions can update a data block, avoiding the overhead of allocating a transaction entry dynamically.</p> <p><b>Maximum transactions</b> - The maximum parameter limits concurrency on a data block.</p>
	<b>Extents</b>	<p>The unit of space allocated to an object whenever the object needs more space.</p> <p><b>Initial Extent</b> - The initial space extent (in bytes) allocated to the object.</p> <p><b>Next Extent</b> - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required.</p> <p><b>Minimum Extents</b> - The appropriate minimum extents value for the object.</p> <p><b>Maximum Extents</b> - The appropriate maximum extents value for the object.</p> <p><b>Percent Increase</b> - Magnifies how an object grows and, can materially affect available free space in a tablespace. Select a value in the corresponding box.</p>
	Optionally, you can modify the <b>Percent Free</b> , <b>Percent Used</b> , or <b>Max Transactions</b> values.	
<b>Performance</b>	Lets you modify the following performance settings:	
	<b>Parallel Query Option</b> group	Lets you specify <b>Degrees</b> and <b>Instances</b> settings for processing queries using many query server processes running against multiple CPUs, which provides substantial performance gains such as reduction of the query completion time.
	<b>Cache</b> group	The <b>Cache</b> setting keeps the blocks in memory by placing it at the most recently used end. This option is useful for small lookup tables.
<b>Space</b>	Lets you view the following usage and the space distribution details:	
	<b>Space Utilization</b> group	Displays the percent of space reserved for future updates.
	<b>FreeLists</b> group	Displays the allocation of data blocks when concurrent processes are issued against the cluster. Identifying multiple freelists can reduce contention for freelists when concurrent inserts take place and potentially improve the performance of the cluster.
	<b>Extents</b> group	The unit of space allocated to an object whenever the object needs more space.

Tab	Settings and tasks
<b>Tables</b>	Lets you view details for each cluster column. Details for each column include the key column in the cluster, the clustered table name, and the key column in the table.
<b>DDL</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## DATABASE LINKS EDITOR (ORACLE)

The Database Links Editor lets you view connection string information for a database link.

### To edit a database link

- 1 Open an editor on the database link. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Definition</b>	Displays connection string information for a database link. Details include the <b>Public</b> setting, the <b>User</b> name and <b>Password</b> associated with the database link, and a <b>Connect String</b> .
<b>DDL</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## FOREIGN KEYS EDITOR (ORACLE)

The Foreign Keys Editor lets you enable a foreign key, manage its column mappings, and specify a delete rule.

### To edit a foreign key

- 1 Open an editor on the foreign key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Column Mapping</b>	The existing column mapping for the foreign key is represented by selected columns in the <b>Main Table</b> and <b>Referenced Table</b> lists. Additional candidates in the <b>Referenced Table</b> list are indicated by enabled column check boxes.  If necessary, use the <b>Table</b> dropdown in the <b>Referenced Table</b> group to choose a new table for this foreign key. Select or deselect columns in the <b>Main Table</b> list and <b>Referenced Table</b> list to form the referential constraint between the two tables.
<b>Properties</b>	Lets you set the foreign key as <b>Enabled</b> and specify a <b>Delete Rule</b> of CASCADE or NO ACTION.
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## FUNCTIONS EDITOR (ORACLE)

The Functions Editor lets you modify basic properties and modify the CREATE/REPLACE DDL for a function.

### To edit a function

- 1 Open an editor on the function. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	Lets you view properties in the following categories:	
	<b>Properties</b>	Lets you view <b>Status</b> and <b>Last Modified</b> details for the function
	<b>Size Information</b>	Lets you view <b>Source Size</b> , <b>Parsed Size</b> , <b>Code Size</b> , and <b>Error Size</b> for the function
<b>Definition</b>	Lets you modify the CREATE OR REPLACE FUNCTION DDL for the function.	
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .	
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## INDEXES EDITOR (ORACLE)

The Indexes Editor lets you manage columns, basic properties, storage, space, and partitioning for an index.

### To edit an index

- 1 Open an editor on the index. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab		Settings and tasks
<b>Columns</b>		Displays a listing of the columns making up the index. Optionally, you can:
		Use the <b>Sort</b> column check box to specify a sort option.
		Click the New button to add a new column to the index.
		Select an existing column and click the Delete button to delete that column from the index.
<b>Properties</b>		Lets you work with settings in the following categories:
	<b>Attributes</b>	Lets you view <b>IsValid</b> and <b>Function-Based</b> properties. Lets you set <b>Index Type</b> (UNIQUE, NONUNIQUE, BITMAP), <b>No Sort</b> , <b>Logging</b> , and <b>Reverse</b> properties.
	<b>Parallel Query Option</b>	Lets you view the <b>No Parallel Execution</b> property. Lets you set the <b>Parallel Degree</b> and <b>Parallel Instances</b> values.
<b>Storage</b>		Lets you work with settings in the following categories:
	<b>Data Block Storage</b>	Lets you choose the <b>Tablespace</b> , and specify <b>Percent Free</b> , <b>Initial Transactions</b> , and <b>Max Transactions</b> values.
	<b>Extents</b>	Lets you specify <b>Initial Extent</b> , <b>Next Extent</b> , <b>Minimum Extents</b> , <b>Maximum Extents</b> , and <b>Percent Increase</b> values.
	<b>Freelists</b>	Lets you specify <b>Freelists</b> and <b>Freelist Groups</b> values.
	<b>Buffer Pool</b>	Lets you specify a buffer pool.
<b>Space</b>		Lets you view settings in the following categories:
	<b>Space Utilization</b>	Lets you view the <b>Size</b> and number of <b>Blocks</b> settings.
	<b>Statistics</b>	Lets you view <b>Index Level</b> , <b>Distinct Keys</b> , <b>Cluster Factor</b> , <b>Leaf Blocks</b> , <b>Leaf Blks/Key</b> and <b>Data Blks/Key</b> properties.



Tab		Settings and tasks
<b>Partition</b>	<p>If the index is not currently partitioned, you can click the <b>Convert To Partitioned</b> button to partition the index. For more information, see <a href="#">Partitioning Oracle indexes, primary keys, and unique keys</a>.</p> <p>If the index is currently partitioned, this tab displays the following partition details:</p>	
	<b>Properties</b>	Lets you view the <b>Locality</b> (Global/Local), <b>Alignment</b> (Prefixed/Non-Prefixed), <b>Partition Type</b> (RANGE or HASH), and <b>Subpartition type</b> properties.
		Click the Edit Partition button to edit partition details. or more information, see <a href="#">Partitioning Oracle indexes, primary keys, and unique keys</a> .
		Click the Drop Partition button to revert to an unpartitioned index.
		<p>Use the following <b>Partition Commands</b> menu items:</p> <p><a href="#">Allocate Extent</a></p> <p><a href="#">Analyze</a></p> <p><a href="#">Coalesce</a></p> <p><a href="#">Deallocate Unused Space</a></p> <p><b>Mark Unusable</b> - opens a dialog that lets you select one or more partitions/subpartitions to be marked as unusable.</p> <p><b>Rebuild</b> - opens a dialog that lets you select one or more unusable subpartitions to be rebuilt.</p> <p><b>Split</b> - opens a dialog that lets you divide a single partition into two partitions. You can split partitions if a single partition is causing maintenance problems because it is too large.</p> <p><b>NOTES:</b></p> <p>If you are preparing to drop or rebuild an index, mark local indexes as unusable. If you want to make unusable indexes valid or to recover space and improve performance, rebuild the unusable indexes.</p> <p>You cannot split a local index partition defined on a hash or composite table.</p> <p>Make sure that you specify an upper bound for the column that is lower than the upper bound for that column in the original partition.</p>
	<b>Columns</b>	Displays partitioning columns.
	<b>Partition Definitions</b>	Displays details for each partition.
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## JOB QUEUE EDITOR (ORACLE)

The Job Queue lets you modify the job's definition (code) and manage the schedule and status for a job.

**To edit a job queue**

- 1 Open an editor on the job queue. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Definition</b>	Lets you modify the statement, function call, or procedure call in the <b>Job Definition</b> area.
<b>Properties</b>	Use the settings in the <b>Schedule</b> group to specify that the job is to run once only or is to run according to a schedule. If it is to run on a schedule, specify the interval in days or using a custom date expression.  Use the settings in the <b>Status</b> group to <b>Enable</b> or <b>Disable</b> the job.
<b>DDL</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**LIBRARIES EDITOR (ORACLE)**

The Libraries Editor lets you view and modify library definitions and manage dependencies and privileges for the library.

**To edit a library**

- 1 Open an editor on the library. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Definition</b>	The <b>Definition</b> tab of the Libraries Editor lets you modify the file name and path for a library and view the current status and whether the library is dynamic.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Privileges</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**MATERIALIZED VIEWS EDITOR (ORACLE)**

The Materialized Views Editor lets you view and modify materialized view information and partitions.

**To edit a materialized view**

- 1 Open an editor on the materialized view. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Information</b>	Lets you work with properties in the following categories:	
	<b>Master</b>	Lets you work with the <b>Table Name</b> and <b>Master View</b> properties.
	<b>Last Refresh</b>	Lets you work with the <b>Last Date</b> and <b>Errors</b> properties.
	<b>Refresh Configuration</b>	Lets you work with the <b>Type</b> , <b>Refresh Method</b> , and <b>Mode</b> properties.
	<b>Rollback Usage</b>	Lets you work with the Local and Master properties.
	<b>Options</b>	Lets you work with the <b>Updatable</b> and <b>Enable Query Rewrite</b> properties.
<b>Storage</b>	Lets you work with properties in the following categories:	
	<b>Placement</b>	Lets you work with <b>Tablespace</b> and <b>Cluster</b> values.
	<b>Data Block Storage</b>	Lets you work with <b>Percent Free</b> , <b>Percent Used</b> , <b>Initial transactions</b> , and <b>Maximum transactions</b> values.
	<b>Extents</b>	Lets you work with <b>Initial Extent</b> , <b>Next Extent</b> , <b>Minimum Extents</b> , <b>Maximum Extents</b> , and <b>Percent Increase</b> values.
<b>Performance</b>	Lets you work with settings in the following categories:	
	<b>Parallel Query Option</b>	<p>The Parallel server query option lets you process queries using many query server processes running against multiple CPUs. This option provides substantial performance gains such as reduction of the query completion time.</p> <p><b>Degrees</b> - Lets you type a value indicating the number of query server processes that should be used in the operation.</p> <p><b>Instances</b> - Lets you type a value indicating how you want the parallel query partitioned between the Parallel Servers.</p>
	<b>Logging</b>	Select <b>Logging</b> to create a log for all Materialized View updates.
	<b>Cache</b>	Select <b>Cache</b> if you want Oracle to put data you access frequently at the most recently used end of the list in the buffer cache when a full table scan is performed. This option is useful for small lookup tables.
<b>Query</b>	Displays the associated query.	
<b>Partitions</b>	<b>Partitioning Method</b>	<p>Displays the partitioning method, including Range-Hash Composite or Range-List Composite.</p> <p>Hash partitions partition the table according to a hash function. Composite partitions use both range and hash types, first partitioning the data by a range of values, and then further dividing the partitions into subpartitions by way of a hash function.</p> <p>List partitioning lets you control how rows map to partitions.</p> <p>You can specify a list of discrete values for the partitioning column in the description for each partition.</p>
	<b>Row Movement</b>	If its key is updated, migrates the row to a new partition.

Tab	Settings and tasks	
	<b>Partitioning Columns</b>	Displays partitioning columns.
	<b>Subpartitioning Columns</b>	Displays subpartitioning columns.
	<b>Partitions</b>	Click Add or Edit to open the Partition dialog box. Click Drop to drop a partition.
	<b>Subpartition Template</b>	If the partitioning type is Range-Hash Composite, displays a list of subpartitions in the subpartition template. Click Add, Insert, or Edit to open the Subpartition dialog box. Click Drop to drop a subpartition.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .	
<b>Privileges</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## MATERIALIZED VIEW LOGS EDITOR (ORACLE)

The Materialized View Logs Editor lets you view replication log table details, and manage space, storage, and performance details for a materialized view log.

### To edit a materialized view log

- Open an editor on the materialized view log. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Information</b>	Displays information on the log table used for replication.	
<b>Storage</b>	Lets you work with settings in the following categories:	
	<b>Data Block Storage</b>	Lets you view the associated <b>Tablespace</b> and <b>Initial Transactions</b> value, and modify the <b>Percent Free</b> , <b>Percent Used</b> , and <b>Max Transactions</b> values.
	<b>Extents</b>	Lets you view <b>Initial Extent</b> , <b>Next Extent</b> , <b>Minimum Extents</b> , <b>Maximum Extents</b> , and <b>Percent Increase</b> values.
	<b>Column Filtering</b>	Lets you select the filter columns to be recorded in the materialized view log. You can specify only one primary key, one ROWID and one filter column list per materialized view log.  The ROWID is a globally unique identifier for a row in a database. It is created at the time the row is inserted into a table, and destroyed when it is removed from a table.

Tab	Settings and tasks	
<b>Performance</b>	Lets you work with settings in the following categories:	
	<b>Parallel Query Option</b>	The Parallel server query option lets you process queries using many query server processes running against multiple CPUs. This option provides substantial performance gains such as reduction of the query completion time.  <b>Degrees</b> - Lets you type a value indicating the number of query server processes that should be used in the operation.  <b>Instances</b> - Lets you type a value indicating how you want the parallel query partitioned between the Parallel Servers.
	<b>Logging</b>	Select <b>Logging</b> to create a log for all Materialized View updates.
	<b>Cache</b>	Select <b>Cache</b> if you want Oracle to put data you access frequently at the most recently used end of the list in the buffer cache when a full table scan is performed. This option is useful for small lookup tables.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## OUTLINES EDITOR (ORACLE)

The Outlines Editor lets you view information on an outline, and modify its category and associated SQL statement.

**NOTE:** The only SQL statements possible with stored outlines are SELECT, DELETE, UPDATE, INSERT...SELECT, and CREATE TABLE...AS SELECT.

### To edit an outline

- Open an editor on the outline. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	In addition to displaying basic creation and identification properties, this tab lets you select a new <b>Category</b> for the outline.
<b>Definition</b>	Lets you view and modify the SQL Statement associated with the outline.
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PACKAGE BODIES EDITOR (ORACLE)

While listed as separate objects in the Datasource Explorer, package bodies are created and edited on the **Body** tab of the Packages Editor. For details, see [Packages Editor \(Oracle\)](#).

## PACKAGES EDITOR (ORACLE)

The Packages Editor lets you view and modify header and body specifications of a package.

### To edit package

- 1 Open an editor on the package. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Header</b>	Lets you modify the package header specifications.	
<b>Body</b>	Lets you modify the package body specifications.	
<b>Information</b>	Lets you work with status and size properties in the following categories:	
	<b>Header and Body</b>	Lets you view <b>Status</b> , <b>Created</b> , and <b>Last Modified Source Size</b> , <b>Parsed Size</b> , <b>Code Size</b> , and <b>Error Size</b> details for the header and body of the package.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .	
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PRIMARY KEYS EDITOR (ORACLE)

The Primary Keys Editor lets you manage columns, basic properties, storage and space, and partitions for a primary key.

### To edit primary key

- 1 Open an editor on the primary key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Displays a listing of the columns making up the primary key. Optionally, you can:	
		Click the New button to add a new column to the index.
		Select an existing column and click the Delete button to delete that column from the index.

Tab		Settings and tasks
<b>Properties</b>	Lets you work with settings in the following categories:	
	<b>Enforcing Index</b>	Lets you view <b>User Defined</b> , <b>Index Owner</b> , and <b>Index Name</b> properties.
	<b>Attributes</b>	Lets you work with <b>No Sort</b> (only available if <b>Enabled</b> is set), <b>Logging</b> (YES, NO, or NONE), <b>Reverse</b> (disabled if <b>No Sort</b> is enabled), <b>Validate</b> , <b>Deferrable</b> , <b>Deferred</b> (IMMEDIATE or DEFERRED and only enabled if <b>Deferrable</b> is enabled), <b>Enabled</b> , <b>Cascade</b> (disabled if <b>Enabled</b> is set), <b>Rely</b> and <b>Update Date</b> properties.
<b>Storage</b>	Lets you work with settings in the following categories:	
	<b>Data Block Storage</b>	Lets you specify <b>Tablespace</b> , <b>Percent Free</b> , <b>Initial Transactions</b> , and <b>Max Transactions</b> values. NOTE: You should never place primary keys on the SYSTEM tablespace.
	<b>Extents</b>	Displays <b>Initial Extent</b> , <b>Next Extent</b> , <b>Percent Increase</b> , <b>Minimum Extents</b> and <b>Maximum Extents</b> values.
	<b>Freelists</b>	Lets you specify <b>Freelists</b> and <b>Freelist Groups</b> values.
	<b>Buffer Pool</b>	Lets you specify a buffer pool.
<b>Space</b>	Lets you work with settings in the following categories:	
	<b>Space Utilization</b>	Lets you view <b>Size</b> and <b>Blocks</b> properties.
	<b>Statistics</b>	Lets you view <b>Index Level</b> , <b>Distinct Keys</b> , <b>Cluster Factor</b> , <b>Leaf Blocks</b> , <b>Leaf Blks/Key</b> , and <b>Data Blks/Key</b> properties.
<b>Partition</b>	If the primary key is currently partitioned, this tab displays the following partition details:	
	<b>Properties</b>	Lets you view the <b>Locality</b> (Global/Local), <b>Alignment</b> (Prefixed/Non-Prefixed), <b>Partition Type</b> (RANGE or HASH), and <b>Subpartition type</b> properties.
		Click the Edit Partition button to edit partition details. or more information, see <a href="#">Partitioning Oracle indexes, primary keys, and unique keys</a> .
		Click the Drop Partition button to revert to an unpartitioned primary key.
		Use the following <b>Partition Commands</b> menu items: <a href="#">Allocate Extent</a> <a href="#">Analyze</a> <a href="#">Coalesce</a> <a href="#">Deallocate Unused Space</a> <b>Mark Unusable</b> - opens a dialog that lets you select one or more partitions to be marked as unusable. <b>Rebuild</b> - opens a dialog that lets you select a partition to be rebuilt. <b>Split</b> - opens a dialog that lets you divide a single partition into two partitions. You can split partitions if a single partition is causing maintenance problems because it is too large.
	<b>Columns</b>	Displays partitioning columns.
	<b>Partition Definitions</b>	Displays details for each partition.
	If the primary key is not currently partitioned, you can click the <b>Convert To Partitioned</b> button to partition the primary key. For more information, see <a href="#">Partitioning Oracle indexes, primary keys, and unique keys</a> .	

Tab	Settings and tasks
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PROCEDURES EDITOR (ORACLE)

The Procedures Editor lets you view and modify the SQL code and properties of a procedure.

### To edit a procedure

- Open an editor on the procedure. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	Lets you work with properties in the following categories:	
	<b>Properties</b>	Lets you view <b>Status</b> , and <b>Last Modified</b> properties.
	<b>Size Information</b>	Lets you view <b>Source Size</b> , <b>Parsed Size</b> , <b>Code Size</b> , and <b>Error Size</b>
<b>Definition</b>	Lets you modify the SQL code for a procedure.	
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .	
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PROFILES EDITOR (ORACLE)

The Profiles Editor lets you manage limits and manage user assignments for the profile.



**To edit a profile**

- 1 Open an editor on the profile. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Resources</b>	Lets you work with settings in the following categories:	
	<b>General Limits</b>	Lets you specify <b>Composite Limit</b> and <b>Private SGA</b> settings.
	<b>Session Limits</b>	Lets you specify the limit on the amount of private space a session can allocate in the shared pool of the SGA. Specific settings are <b>Sessions Per User</b> , <b>CPU Per Session</b> , <b>Logical Reads</b> .
	<b>Time Limits</b>	Lets you specify the limit on total connection time per session. Specific settings are <b>Connect Time</b> and <b>Idle Time</b> .
	<b>Call Limits</b>	Lets you specify the CPU time limit for a call (a parse, execute, or fetch), expressed in hundredths of seconds. Specific settings are <b>CPU Per Call</b> and <b>Logical Reads</b> .
	<b>Login Limits</b>	Lets you specify the number of <b>Failed Login Attempts</b> on the user account before the account is locked and the <b>Account Lock Time</b> .
	<b>Password Limits</b>	Lets you specify the <b>Lifetime</b> , <b>Reuse Time</b> , <b>Reuse Max</b> , <b>Grace Period</b> and a <b>Verify Function</b> for passwords.
<b>Users</b>	Use the <b>Assign</b> button to open a dialog that lets you assign a user to this profile or select a user from the list, and click the <b>Unassign</b> button to open a dialog prompting you to confirm that the user is to be unassigned.	
<b>DDL</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**REDO LOG GROUPS EDITOR (ORACLE)**

The Redo Log Groups Editor lets you manage the members in a redo log group.

**To edit a redo log group**

- 1 Open an editor on the redo log group. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Redo Log Members</b>	Lets you add new members to the redo log group, edit existing members, and delete members from the redo log group.
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## ROLES EDITOR (ORACLE)

The Roles Editor lets you manage authentication and grant/revoke profiles for users and logins.

### To edit a role

- 1 Open an editor on the role. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Authentication</b>	Lets you manage role identity. When creating a role, you must establish certain rules governing its use. You can specify whether or not a role must be identified when used. If you require role identification, you can authenticate the user externally through the operating system, or with a specific password.  If you specified that the role requires identification, provide a <b>Password</b> and specify whether the role is to be authenticated <b>Globally</b> or <b>Externally</b> .
<b>User/Roles</b>	Displays permissions for this role to logins or users. Click <b>Grant</b> to open a dialog that lets you grant this role to a login or another role. Select a role or login and click <b>Revoke</b> to revoke the role or login.
<b>Object Privileges and System Privileges</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## ROLLBACK SEGMENTS EDITOR (ORACLE)

The Rollback Segments Editor lets you view rollback segment status, manage rollback segment storage, and view activity levels.

### To edit a rollback segment

- 1 Open an editor on the rollback segment. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Status</b>	Lets you enable and disable a rollback segment and displays status details for the rollback segment. The tab displays whether the rollback segment is online or offline and provides the associated <b>Tablespace</b> , <b>Size</b> , and <b>No. of Extents</b> .

Tab	Settings and tasks	
<b>Storage</b>	Lets you work with settings in the following categories:	
	<b>Extents</b>	<p>The unit of space allocated to an object whenever the object needs more space.</p> <p><b>Initial Extent</b> - The initial space extent (in bytes) allocated to the object.</p> <p><b>Next Extent</b> - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required.</p> <p><b>Optimal Size</b> - optimal extent size</p> <p><b>Minimum Extents</b> - The appropriate minimum extents value for the object.</p> <p><b>Maximum Extents</b> - The appropriate maximum extents value for the object.</p>
	<b>Extent Detail</b>	Displays extent details.
<b>Activity</b>	Lets you work with settings in the following categories:	
	<b>Activity Levels</b>	Displays <b>Active Transactions</b> , <b>Writes</b> , <b>Gets</b> and <b>Waits</b> values.
	<b>Dynamic Sizing</b>	Displays <b>High Watermark</b> , <b>Extends</b> , <b>Shrinks</b> , and <b>Wraps</b> values.
<b>DDL</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## SEQUENCES EDITOR (ORACLE)

The Sequences Editor lets you manage parameters for a sequence, manage database objects dependent on the sequence, and manage privileges for the sequence.

### To edit a sequence

- Open an editor on the sequence. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Definition</b>	Lets you work with settings in the following categories:	
	<b>Parameters</b>	Lets you specify increment, minimum value and maximum value settings.
	<b>Current/Next Sequence Numbers</b>	Lets you work with sequence cycle numbers.
	<b>Options</b>	Lets you specify <b>Cache Size</b> , <b>Cycle When Reach Max/Min</b> , and <b>Generate Numbers in Order</b> (useful when you are using the sequence number as a timestamp) values.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .	
<b>Privileges</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .	

Tab	Settings and tasks
<b>DDL</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## SYNONYMS EDITOR (ORACLE)

The Synonyms Editor lets you view base object information and manage database objects dependent on a synonym.

### To edit a synonym

- 1 Open an editor on the synonym. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Displays the type, owner, name, and other details of the object referenced by the synonym.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## TABLES EDITOR (ORACLE)

The Tables Editor lets you manage columns, constraints, storage and space, and partitions for a table.

**To edit a table**

- 1 Open an editor on the table. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Displays the currently defined columns in the table. For any selected column, the Property/Value list provides additional detail on that column. Available properties depend on the datatype you choose as well as on the property values you select:  <b>Datatype</b> properties- Lets you select a <b>type</b> and depending on your selection, additional properties such as <b>Scale</b> , <b>Width</b> , and <b>Unused</b> may be available). <b>Allow Nulls</b> - Select this check box to allow nulls in this column. <b>Encryption</b> properties include <b>Password</b> , <b>Salted</b> , and <b>EncryptionAlgorithm</b> . <b>Default Value</b> lets you choose among CURRENT_TIMESTAMP, USER, SYSDATE, and UID. <b>Comment</b> lets you add a comment to the column. <b>LOB Storage</b> settings are available for bfile, blob, clob, and nclob types. They include <b>Segment Name</b> , <b>Configuration</b> properties ( <b>Tablespace</b> , <b>Chunk</b> , <b>Percent Version</b> , <b>Enable Storage In Row</b> , <b>Cache</b> , and <b>Logging</b> ), and <b>Storage</b> properties ( <b>Initial Extent</b> , <b>Next Extent</b> , <b>Percent Increase</b> , <b>Minimum Extents</b> , <b>Maximum Extents</b> , <b>Free Lists</b> , and Free List Groups). Optionally, you can:	
		Select a column, and in the <b>Property/Value</b> list modify property values for that column.
		Click <b>Add Column</b> , provide a name for the new column, and set property values for the column.
		Select a column and click <b>Delete</b> to remove the column from the table.
<b>Properties</b>	Lets you work with settings in the following categories:	
	Table	Lets you work with the following properties, corresponding to CREATE TABLE clauses: <b>Cache</b> , <b>Row Movement</b> , <b>Parallel Degree</b> , and <b>Parallel Instances</b> .
	Physical	Lets you view the <b>Row Organization</b> property. Lets you set the <b>Logging</b> and <b>Table Compression</b> properties.
<b>Indexes</b>	Lets you manage indexes for a table. On opening, the list of current indexes for the table is displayed. Optionally, take one of the following actions:	
		Click <b>Add</b> to open a dialog that lets you add a new index to the table.
		Select an index and click <b>Edit</b> to open a dialog that lets you edit index properties.
		Select an index and click <b>Drop</b> to open a dialog that lets you remove the index from the table.
<b>Constraints</b>	Lets you manage constraints for the table. Constraints are grouped by type, under folders. Optionally take one of the following actions:	
		Select a constraint type folder and click <b>Add</b> to open a dialog that lets you add a constraint of that type.
		Select a constraint and click <b>Edit</b> to open a dialog that lets you modify the constraint details.
		Select a constraint and click <b>Drop</b> to remove the constraint.

Tab	Settings and tasks	
<b>Storage</b>	Lets you work with settings in the following categories:	
	<b>Data Block Storage</b>	Lets you choose the <b>Tablespace Name</b> , and specify <b>Percent Free</b> , <b>Initial Transactions</b> , and <b>Max Transactions</b> values.
	<b>Extents</b>	Lets you view <b>Initial Extent</b> , <b>Next Extent</b> , <b>Minimum Extents</b> , <b>Maximum Extents</b> , and <b>Percent Increase</b> values.
	<b>Freelists</b>	Lets you specify <b>Freelists</b> and <b>Freelist Groups</b> values.
	<b>Bufferpool</b>	Lets you specify a <b>Buffer Pool</b> .
<b>IOT Properties</b>	Lets you work with settings in the following categories:	
	Ungrouped	Lets you specify an overflow segment.
	<b>Percent Threshold</b>	Lets you specify the percentage of space reserved for an index-organized table.
	<b>Key Compression</b>	Lets you enable or disable compression and provide a compression value.
<b>Space</b>	Lets you work with settings in the following categories:	
	<b>Space Utilization</b>	Lets you view <b>Size</b> and number of <b>Blocks</b> .
	<b>Row Information</b>	Lets you view the <b>Number of Rows</b> , <b>Average Row Length</b> , and <b>Chain Rows</b> properties.
	<b>Extents</b>	Let you view the <b>Number of Extents</b> and <b>Maximum Extents</b> values.
<b>LOB columns</b>	Lets you work with settings in the following categories:	
	<b>LOB Column Segment and LOB Column Index</b>	Lets you view the <b>Segment Name</b> , <b>No. of Extents</b> , <b>Max Extents</b> , <b>Size</b> , and <b>Blocks</b> properties.
	<b>Segment Extents and Index Extents</b>	Lets you view the <b>Extent ID</b> , <b>File ID</b> , <b>Block ID</b> and number of <b>Blocks</b> properties for segment and index extents.
<b>Partition</b>	Lets you work with table partitions	
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .	
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .	
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## TABLESPACES EDITOR (ORACLE)

The Tablespaces Editor lets you manage datafiles, space, storage, quotas, and objects for a tablespace.

## To edit a tablespace

- 1 Open an editor on the tablespace. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Datafile</b>	The tab lists details for each datafile on the tablespace. Optionally you can <b>Add</b> a datafile or <b>Edit</b> a selected datafile. For more information, see <a href="#">Adding or Editing Tablespace Datafiles</a> .	
<b>Storage</b>	Lets you work with settings in the following categories:	
	<b>Extents</b>	Lets you view <b>Initial Extent</b> , <b>Next Extent</b> , <b>Minimum Extent Size</b> , <b>Minimum Extents</b> , <b>Maximum Extents</b> , and <b>Percent Increase</b> values.
	<b>Segment Management</b>	Lets you view whether segment management is <b>Automatic</b> or <b>Manual</b> , and if manual, the <b>TableSpace Block Size</b> .
	<b>Extent Details</b>	Lists details for each extent.
	<b>Extent Management</b>	Lets you view whether the <b>Type</b> of extent management is LOCAL or DICTIONARY, whether the <b>Allocation Type</b> is AUTO, UNIFORM, or USER, and the <b>Extent Size</b> .
	For background information, see <a href="#">About Oracle Tablespace Storage</a> .	
<b>Space</b>	Displays a <b>Free Space</b> pie chart and free space, used space and total space statistics. This tab also displays the current <b>Fragmentation Index</b> and offers the option to coalesce the tablespace	
<b>Map</b>	The <b>Map</b> tab of the tablespace editor lets you view a tablespace map. Options on the tab let you perform tasks such as scanning for honeycomb or bubble space fragmentation.	
<b>Objects</b>	Displays the objects currently stored on the tablespace, group under object type folders. Optionally you can:	
		Specify a logging option using the <b>Log Changes When Scheme/Data is Modified?</b> radio set.
		Select an object under one of the object folders and click <b>Edit</b> to open an object editor on that object.
<b>Quotas</b>	Oracle limits the amount of space that can be allocated for storage of a user's objects within the specified tablespace to the amount of the quota. Users with privileges to create certain types of objects can create those objects in the specified tablespace.  The <b>Quotas</b> tab of the Tablespace editor lets you manage user space quotas for tablespaces on the current datasource. Optionally, you can:	
		Click <b>Add</b> or select a user and click <b>Edit</b> to assign a user unlimited or a specific space usage quota on the tablespace. For details, see <a href="#">Adding or Editing User Tablespace Quotas</a> .
		Select an existing user and click <b>Drop</b> to delete the quota for that user
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## ADDING OR EDITING USER TABLESPACE QUOTAS

When you assign a quota:

- Users with privileges to create certain types of objects can create those objects in the specified tablespace.
- Oracle limits the amount of space that can be allocated for storage of a user's objects within the specified tablespace to the amount of the quota.

**NOTE:** This functionality is available for Oracle only.

User tablespace quotas are added and modified from the Tablespaces editor. For more information, see [Tablespaces Editor \(Oracle\)](#).

The table below describes the options and functionality on the **Add User Quota...** or **Edit User Quota...** dialog boxes:

Option	Description
<b>User selection list</b> (Add only)	Lets you select one or more users to assign a quota.
<b>Quota</b>	Lets you set a quota for the selected user or users. You can select an unlimited, or a specified size. Unlimited - Lets you place an unlimited quota on the tablespace. Other - Lets you place a specified quota in KB or MB on the tablespace.

## ABOUT ORACLE TABLESPACE STORAGE

The **Storage** tab of the Tablespace Editor lets you view storage details for tablespaces on the current datasource.

**TIP:** Always create tablespaces for user data and never place user tables and indexes in the SYSTEM tablespace. Placing user objects in the SYSTEM tablespace can degrade performance and introduce space-related headaches to the database.

Oracle8i or later supports locally managed tablespaces, which can all but eliminate the problem of tablespace fragmentation. It totally does away with the storage parameters of MINEXTENTS, MAXEXTENTS, PCTINCREASE, and NEXT. With locally managed tablespaces you either specify the initial extent size and let Oracle automatically size all other extents, or specify a uniform extent size for everything.

For users using a version earlier than Oracle 8i and locally managed tablespaces, there are manual methods can employ to assist in the fight against tablespace fragmentation. They include:

- Setting PCTINCREASE to zero for all tablespaces and objects to promote same-sized extents.
- Specifying equal-sized allotments for your INITIAL and NEXT object storage parameters.



- Grouping objects with like growth and storage needs together in their own tablespaces.

**TIP:** One of the best ways to avoid fragmentation in a tablespace is to pre-allocate the space that your objects will use. If possible, plan for one to two years' growth for each object and allocate your space accordingly. Having initial empty objects will not affect table scan times as Oracle only scans up to the high-water mark (the last used block) in a table.

Of all your tablespaces, you want to avoid fragmentation problems in your SYSTEM tablespace the most as this is the major hotbed tablespace for Oracle activities. The easiest way to avoid this is to not allow any user (even the default DBA ID's SYS and SYSTEM) to have access to it. There are three ways to do this:

- Ensure no user has a DEFAULT or TEMPORARY tablespace assignment of SYSTEM.
- Ensure no user has a quota set for SYSTEM.

Ensure no user has been granted the UNLIMITED TABLESPACE privilege.

#### ADDING OR EDITING TABLESPACE DATAFILES

Rapid SQL lets you edit datafile details for a tablespace or add a new datafile.

**NOTE:** This functionality is available for Oracle.

#### To add or edit tablespace datafiles

- 1 Open an editor on a tablespace. For details, see [Opening an Object Editor](#).
- 2 Click the **Partition** tab and take one of the following actions:
  - Click **Add**.
  - Select a datafile from the list and click **Edit**.

Rapid SQL opens an **Add Datafile** or **Modify Datafile** dialog.

- 3 Use the following table as a guide to understanding and modifying the settings on the tabs of this dialog:

Control	Description
<b>Datafile Name</b>	Type the name of the datafile in the box.
<b>Datafile Size</b>	Type the datafile size and then click the list to indicate MB or KB.
<b>Datafile Status</b>	Select ONLINE or OFFLINE availability for the datafile
<b>Reuse Existing File?</b> (available when adding a datafile only)	The default is <b>No</b> .
<b>Autoextend</b>	Select this check box if you want to autoextend the datafile when more extents are required. This enables the options below.
<b>Disk Space to Allocate to the Datafile When More Extents are Required</b>	Type the disk space to allocate to the datafile when more extents are required by typing it and then click the list to indicate MB or KB.

Control	Description
<b>Maximum Disk Space Allowed for Allocation to the Datafile</b>	Specify the maximum disk space allowed for allocation to the datafile. Select <b>Unlimited</b> or <b>Other</b> . If you select <b>Other</b> , in the <b>Other</b> box, type the amount and then click the list to indicate MB or KB.

5 Click **OK**.

Rapid SQL makes the datafile changes and returns you to the Tablespace editor.

## TRIGGERS EDITOR (ORACLE)

The Triggers Editor lets you modify actions, events, and other details for a trigger.

### To edit a trigger

- 1 Open an editor on the trigger. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	In addition to displaying basic identification and creation properties, this tab lets you work with settings in the following categories:	
	<b>Attributes</b>	Lets you specify whether the trigger is <b>Enabled</b> , select the <b>Trigger Timing</b> (BEFORE, AFTER), and the <b>Trigger Type</b> (ROW, STATEMENT).
	<b>Correlation Names</b>	Lets you provide an <b>Old Table Alias</b> and a <b>New Table Alias</b> as well as construct a <b>When Clause</b> .
	<b>Status</b>	Lets you view <b>Object Status</b> , <b>Create Date</b> , <b>Last Modified</b> , and <b>Base Object Type</b> properties.
	<b>Size Information</b>	Lets you view <b>Source Size</b> , <b>Parsed Size</b> , <b>Code Size</b> , and <b>Error Size</b> properties.
<b>Events</b>	Lets you select the INSERT, UPDATE, or DELETE event for the trigger.	
<b>Column Selection</b>	Lets you select the associated columns.	
<b>Action</b>	Lets you modify the trigger action PL/SQL block for any trigger on the datasource. You modify the body of the trigger in the Trigger Action (PL/SQL Block) area.	
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## TYPE BODIES EDITOR (ORACLE)

The Type Bodies Editor contains the code for the methods that implement an object type. To create or replace a type body in one of your schema, you must have the CREATE TYPE or CREATE ANY TYPE system privilege. To replace a type in another user's schema, you must have the DROP ANY TYPE system privilege.

While listed as separate objects in the Datasource Explorer, type bodies are created and edited on the **Body** tab of the Types Editor. For details, see [Types Editor \(Oracle\)](#).

## TYPES EDITOR (ORACLE)

The Types Editor lets you manage header text and body text for a type.

### To edit a type

- 1 Open an editor on the type. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Header</b>	Lets you manage the type header text.
<b>Body</b>	Lets you create and modify type body text for the type. For information on creating type bodies, see <a href="#">Object Type Wizard (Oracle)</a> .
<b>Information</b>	Displays the header and body information for a type.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Privileges</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## UNIQUE KEYS EDITOR (ORACLE)

The Unique Keys Editor lets you manage columns, basic properties, storage and space, and partitions for a unique key.

**To edit a unique key:**

- 1 Open an editor on the unique key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Displays a listing of the columns making up the index. Details include the <b>Column</b> name, the <b>Datatype</b> , and whether the definition allows <b>Nulls</b> . Optionally, you can:	
		Click the New button to add a new column to the index.
		Select an existing column and click the Delete button to delete that column from the index.
<b>Properties</b>	In addition to displaying basic identification information, this tab lets you work with settings in the following categories:	
	<b>Enforcing Index</b>	Lets you view <b>User-Defined</b> , <b>Index Owner</b> , and <b>Index Name</b> settings.
	<b>Attributes</b>	Lets you set <b>No Sort</b> , <b>Logging</b> , <b>Reverse</b> , <b>Validate</b> , <b>Deferrable</b> , <b>Deferred</b> , <b>Enabled</b> , <b>Cascade</b> , <b>Rely</b> , and <b>Update Date</b> properties.
<b>Storage</b>	Lets you work with settings in the following categories:	
	<b>Data Block Storage</b>	Lets you choose the <b>Tablespace Name</b> , and specify <b>Percent Free</b> , <b>Initial Transactions</b> , and <b>Max Transactions</b> values.
	<b>Extents</b>	Lets you specify <b>Initial Extent</b> , <b>Next Extent</b> , <b>Minimum Extents</b> , <b>Maximum Extents</b> , and <b>Percent Increase</b> values.
	<b>Freelists</b>	Free lists let you manage the allocation of data blocks when concurrent processes are issued against the cluster. Identifying multiple free lists can reduce contention for free lists when concurrent inserts take place and potentially improve the performance of the cluster.  This tab lets you specify <b>Freelists</b> and <b>Freelist Groups</b> values.
	<b>Buffer Pool</b>	Lets you select a <b>Buffer Pool</b> .
<b>Space</b>	Lets you work with settings in the following categories:	
	<b>Space Utilization</b>	Displays the <b>Size</b> and number of <b>Blocks</b> properties.
	<b>Statistics</b>	Lets you view <b>Index Level</b> , <b>Distinct Keys</b> , <b>Cluster Factor</b> , <b>Leaf Blocks</b> , <b>Leaf Blks/Key</b> and <b>Data Blks/Key</b> settings.
	<b>Extents</b>	The unit of space allocated to an object whenever the object needs more space. This tab lets you view <b>Number of Extents</b> and <b>Maximum Extents</b> properties.

Tab	Settings and tasks	
<b>Partition</b>	If the unique key is currently partitioned, this tab displays the following partition details:	
	<b>Properties</b>	Lets you view the <b>Locality</b> (Global/Local), <b>Alignment</b> (Prefixed/Non-Prefixed), <b>Partition Type</b> (RANGE or HASH), and <b>Subpartition type</b> properties.
		Click the Edit Partition button to edit partition details. For more information, see <a href="#">Partitioning Oracle indexes, primary keys, and unique keys</a> .
		Click the Drop Partition button to revert to an unpartitioned unique key.
		Use the following <b>Partition Commands</b> menu items: <a href="#">Allocate Extent</a> <a href="#">Analyze</a> <a href="#">Coalesce</a> <a href="#">Deallocate Unused Space</a> <b>Mark Unusable</b> <b>Rebuild</b> <b>Split</b>
	<b>Columns</b>	Displays partitioning columns.
	<b>Partition Definitions</b>	Displays details for each partition.
	If the unique key is not currently partitioned, you can click the <b>Convert To Partitioned</b> button to partition the unique key. For more information, see <a href="#">Partitioning Oracle indexes, primary keys, and unique keys</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .	

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## USERS EDITOR (ORACLE)

The Users Editor lets you manage basic properties, roles, tablespace quotas, and associated objects for a user.

**To edit a user**

- 1 Open an editor on the user. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	Lets you work with settings in the following categories:	
	<b>Defaults</b>	Lets you select a <b>Default Tablespace</b> and Temporary <b>Tablespace</b> as well as select the <b>Profile</b> used by this user. For more on creating profiles, see <a href="#">Profile Wizard (Oracle)</a> .
	<b>Account</b>	<p><b>Identified By</b> - Lets you select among REQUIRED YES, REQUIRED EXTERNAL, and REQUIRED_GLOBAL.</p> <p><b>Password</b> - (available with <b>Identified By</b> value of REQUIRED_YES only) lets you specify a password.</p> <p><b>External Name</b> - (available with <b>Identified By</b> value of REQUIRED_GLOBAL only) lets you specify the name of the user in the enterprise directory service.</p> <p><b>Account Locked</b> - lets you lock this user account.</p> <p><b>Password Expired</b> - (available with <b>Identified By</b> value of REQUIRED_YES only) lets you mark the password as expired, forcing the user to change their password before being allowed to connect.</p>
<b>Role</b>	Lets you select the roles that are to be assigned to this user.	
<b>Quota</b>	Lets you set tablespace quotas for the user. Set an unlimited quota by selecting a tablespace and selecting the <b>Unlimited</b> radio button. Set a specific quota by selecting a tablespace, selecting the <b>Other</b> radio button, and providing a <b>Quota</b> value in kilobytes or megabytes.	
<b>Objects</b>	Lets you manage database objects associated with a user. Objects are grouped within object type folders. Optionally, you can:	
		Select an object and click <b>Edit</b> to open an object editor on the selected object.
		Select an object and click <b>Drop</b> to initiate dropping the selected object.
<b>Object Privileges and System Privileges</b>	For details on using these tabs, see <a href="#">Working with Privileges and Permissions</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**VIEWS EDITOR (ORACLE)**

The Views Editor lets you work with view columns and manage the dependencies and permissions for the view.

**To edit a view**

- 1 Open an editor on the view. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Displays a listing of the columns making up the view.
<b>Definition</b>	Lets you view and modify the DDL that will implement any changes made with the editor.
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PARTITIONING ORACLE INDEXES, PRIMARY KEYS, AND UNIQUE KEYS

The **Add Partition** wizard lets you partition an Oracle index, primary key, or unique key.

**To partition an index, primary key, or unique key**

- 1 Open an editor on the index, primary key, or unique key. For details, see [Opening an Object Editor](#).
- 2 On the **Partition** tab, click **Convert to Partitioned**.

**NOTE:** If the object is already partitioned, the **Convert to Partitioned** button is not present.

- 3 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	In addition to displaying basic identification and ownership properties, this pane lets you work with the following settings: <b>Locality</b> - GLOBAL or LOCAL. <b>Alignment</b> - displays whether this partitioned object is prefixed or non-prefixed. <b>Partition Type</b> - RANGE or HASH. <b>Subpartition Type</b> - displays the subpartition type.
<b>Columns</b>	For each column to be added, click New to add a column, select a <b>Column</b> name, and specify a <b>Sort</b> order.
<b>Range Definitions</b> (only available for a <b>Partition Type</b> of RANGE)	Lets you create an ordered list of partitions. For each partition, click New to open the <b>Add Partition Definition</b> wizard. On the <b>Partition Definition</b> pane, provide a <b>Name</b> , select a <b>Tablespace</b> , and enable or disable <b>Logging</b> . On the <b>Storage</b> pane, provide <b>Data Block Storage</b> , <b>Extents</b> , and <b>Freelists</b> property values. Click <b>Add</b> . The <b>Range Definitions</b> pane also lets you edit and remove partitions.

Tab	Settings and tasks
<b>Hash Definitions</b> (only available for a <b>Partition Type</b> of HASH)	Lets you specify a Partition method: <b>None</b> <b>Number of Partitions</b> - Specify a <b>Number of Partitions</b> , and then for each partition, click the New button and select a tablespace. <b>By Partition Name</b> - For each partition, click New, specify a <b>Name</b> and select a <b>Tablespace</b> .

4 When ready, click **Finish**.

## SYBASE ASE OBJECT EDITORS

Rapid SQL includes an Object Editor for all supported Sybase ASE objects. To see an Editor for a specific object, click the corresponding link below:

**NOTE:** If an objects has dependent objects, such as tables, triggers, procedures and views, you can view and access their dependent objects in the editor.

- [Aliases Editor \(Sybase\)](#)
- [Check Constraints Editor \(Sybase\)](#)
- [Databases Editor \(Sybase\)](#)
- [Defaults Editor \(Sybase\)](#)
- [Extended Procedures Editor \(Sybase\)](#)
- [Foreign Keys Editor \(Sybase\)](#)
- [Groups Editor \(Sybase\)](#)
- [Indexes Editor \(Sybase\)](#)
- [Logins Editor \(Sybase\)](#)
- [Primary Keys Editor \(Sybase\)](#)
- [Procedures Editor \(Sybase\)](#)
- [Roles Editor \(Sybase\)](#)
- [Rules Editor \(Sybase\)](#)
- [Segments Editor \(Sybase\)](#)
- [Tables Editor \(Sybase\)](#)
- [Triggers Editor \(Sybase\)](#)
- [Unique Keys Editor \(Sybase\)](#)
- [User Datatypes Editor \(Sybase\)](#)



- [Users Editor \(Sybase\)](#)
- [Users Editor \(Sybase\)](#)

## ALIASES EDITOR (SYBASE)

The Aliases Editor lets you view and modify an alias definition.

### To edit an alias

- 1 Open an editor on the alias. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you change the user referenced by the alias.
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## CHECK CONSTRAINTS EDITOR (SYBASE)

The Check Constraints Editor lets you manage a check constraint definition and modify the condition expression.

### To edit a check constraint

- 1 Open an editor on the check constraint. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Definition</b>	Lets you edit the condition in the <b>Check Condition</b> box. The <b>Table Columns</b> button acts a time saver in editing the condition. It opens a dialog that lets you select and paste column names into the check condition expression.
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## DATABASES EDITOR (SYBASE)

The Databases Editor lets you manage properties, placement, and binds for a database as well as view associated space statistics.

### To edit a database

- 1 Open an editor on the database. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	Lets you work with the following categories:	
	<b>Creation Properties</b>	Has the <b>Database Name</b> , <b>Owner</b> , <b>Transfer aliases and permissions</b> , and <b>Type</b> properties.
	<b>Options</b>	Lets you work with the <b>abort tran on log full</b> , <b>allow nulls by default</b> , <b>auto identity</b> , <b>dbo use only</b> , <b>ddl in tran</b> , <b>identity in nonunique index</b> , <b>no chkpt on recovery</b> , <b>no free space acctg</b> , <b>read only</b> , <b>select into/bulkcopy/pilsort</b> , <b>single user</b> , and <b>trunc log on chkpt</b> properties.
<b>Placement</b>	Displays device fragments. Optionally you can:	
		Select a fragment to view the <b>Device Name</b> , <b>Size</b> , and <b>Device Type</b> (DATA AND LOG or DATA ONLY) for the fragment.
		Click the New button to open a dialog that lets you provide the data device and log device details for a new database fragment.  NOTE: We strongly recommend that you place the transaction log on a separate device from all other database fragments.
<b>Space</b>	Lets you view pie charts showing the data space usage and the transaction log (if available) space usage for the database.	
<b>Bindings</b>	This tab is only available for databases created with a <b>Type</b> of TEMP. It displays all applications or logins bound to the database. Optionally, you can:	
		Bind another application or login to the database by selecting the APPLICATION or LOGIN folder as appropriate, clicking <b>Add</b> , and in the dialog that opens select a login or provide an application name.
		Remove bindings by selecting the specific application or login and clicking <b>Drop</b> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## DEFAULTS EDITOR (SYBASE)

The Defaults Editor lets you modify the value and manage bindings for a default.

### To edit a default

- 1 Open an editor on the default. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you change the <b>Value</b> of the default.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## EXTENDED PROCEDURES EDITOR (SYBASE)

The Procedures Editor lets you view and modify an extended procedure definition.

### To edit an extended procedure

- 1 Open an editor on the extended procedure. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you specify the <b>Library Name</b> for the extended procedure.
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## FOREIGN KEYS EDITOR (SYBASE)

The Foreign Keys Editor lets you modify the column mapping and specify properties of a foreign key.

### To edit a foreign key

- 1 Open an editor on the foreign key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Column Mappings</b>	The existing column mapping for the foreign key is represented by selected columns in the <b>Main Table</b> and <b>Referenced Table</b> lists. Additional candidates in the <b>Referenced Table</b> list are indicated by enabled column check boxes.  If necessary, use the <b>Table</b> dropdown in the <b>Referenced Table</b> group to choose a new table for this foreign key. Select or deselect columns in the <b>Main Table</b> list and <b>Referenced Table</b> list to form the referential constraint between the two tables.
<b>Properties</b>	Lets you specify a <b>Match Full</b> value of TRUE or FALSE.
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## GROUPS EDITOR (SYBASE)

The Groups Editor lets you manage users in a group and the group's associated privileges.

**TIP:** The [Refresh](#) button lets you refresh or clear the editor's contents, and log SQL.

### To edit a group

- 1 Open an editor on the group. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Users</b>	Lets you add users to, or remove users from the group. To move a user from one of the <b>Users Not In Group</b> or <b>Users In Group</b> lists to the other list, select the user in the list and click the <b>Join Group</b> or <b>Leave Group</b> button.
<b>Object Privileges</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>System Privileges</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>DDL</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## INDEXES EDITOR (SYBASE)

The Indexes Editor lets you manage columns, properties, and partitions, and view statistics for an index.

### To edit an index

- 1 Open an editor on the index. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Lets you manage columns that make up the index. On opening, this tab shows the existing columns. For each column, the datatype (and if applicable the precision in brackets) and whether the table definition permits nulls in the target table column. Optionally you can:	
		Change the <b>Sort</b> order of a column.
		Click the New button and select a column name from the dropdown, to add a column to the index.
		Select a column and click the Drop button to delete the column from the index.
		Select a column and use the arrow buttons to reorder the columns in the index.
<b>Properties</b>	Lets you work with settings in the following categories:	
	<b>Attributes</b>	Has the <b>Index Type</b> (UNIQUE or NONUNIQUE), <b>Clustered</b> , <b>Ignore Duplicate Key</b> (available with <b>Index Type</b> of UNIQUE), <b>Ignore Duplicate Rows</b> (available with <b>Clustered</b> enabled), <b>Allow Duplicate Rows</b> (available with <b>Clustered</b> enabled), and <b>Maximum Rows Per Page</b> properties.
	<b>Storage</b>	Lets you set <b>Reserve Page Gap</b> , <b>Segment Name</b> , <b>Fill Factor</b> , <b>Prefetch Strategy</b> , <b>MRU Replacement Strategy</b> settings
<b>Partition</b>	Lets you work with settings in the following categories:	
	<b>Properties</b>	Lets you view the <b>Locality</b> (Global/Local), <b>Alignment</b> (Prefixed/Non-Prefixed), <b>Partition Type</b> (including Range-Hash Composite or Range-List Composite), and <b>Subpartition type</b> properties.
	<b>Columns</b>	Displays partitioning columns.
	<b>Partition Definitions</b>	Displays details for each partition.
<b>Statistics</b>	Lets you view statistics in the following categories:	
	<b>Page Statistics</b>	<b>Data Pages</b> , <b>Reserved Pages</b> , <b>Used Pages</b> , <b>OAM Page Ratio</b> , and <b>Index Page Ratio</b> .
	<b>Row Statistics</b>	<b>Maximum Row Size</b> , <b>Minimum Row Size</b> , <b>Max Size of Non-Leaf Index Row</b> , and <b>Maximum Rows Per Page</b> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## LOGINS EDITOR (SYBASE)

The Logins Editor lets you manage basic properties for a login, associated users and roles, and accounting.

**To edit a login**

- 1 Open an editor on the login. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Definitions</b>	Lets you work with the following settings:	
	<b>Full Name</b>	OPTIONAL: A user-defined descriptive name for the target login. <b>TIP:</b> To identify the login, in the Full Name box, type the actual name of the user who uses the target login.
	<b>Default Database</b>	The database that Sybase ASE places the target user into when they log in to the target server.
	<b>Default Language</b>	The language that Sybase ASE uses to displays the target login's system messages.
	<b>Password Status</b>	The status, current, expired, of the password.
	<b>Last Password Change</b>	The last date and time that the target login's password changed.
	<b>Password Expiration</b>	Duration, in days, before passwords expire.
	<b>Minimum Password Length</b>	Minimum number of characters in a password.
	<b>Maximum Login Attempts</b>	Maximum number of login attempts before the login (account) is suspended.
	<b>Currently Logged In</b>	The login status of the target login.
	<b>Account Status</b>	Whether the login is locked or unlocked.
<b>Users</b>	<p>Before a login name can access a user database, it must be added as a user or alias to that database by the system administrator or database owner.</p> <p>NOTE: User IDs and passwords are not shared across Windows servers unless the Security Account Delegation option is turned on in the AD Security groups.</p> <p>Use the <b>Add User</b> button to add a selected database from the <b>Databases where the login does NOT have a user account</b> to the <b>Databases where the login HAS a user account</b> list. Conversely, remove a login from a database by using the <b>Drop User</b> button to move a selected database from the <b>Databases where the login HAS a user account</b> list.</p>	
<b>Roles</b>	Displays roles currently assigned to this login. Optionally, you can grant or revoke roles as needed. Also, you can use the <b>Def. Roles</b> button to open a dialog that lets you add or delete roles from the set of default roles assigned to this login.	

Tab	Settings and tasks
<b>Accounting</b>	The <b>Accounting</b> tab of the Logins Editor lets you manage chargeback accounting statistics for every login on the current server. Chargeback accounting statistics are CPU and I/O usage statistics that Sybase ASE accumulates for every login. To start a new accounting period, the system administrator must clear all previous login statistics. Optionally, you can click the <b>Clear Statistics</b> button to start a new accounting interval.
<b>DDL</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PRIMARY KEYS EDITOR (SYBASE)

The Primary Keys Editor lets you manage the columns that make up the primary key and its basic properties, as well as view statistics for the primary key.

### To edit a primary key

- Open an editor on the primary key. For details, see [Opening an Object Editor](#).
- Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Lets you manage columns that make up the primary key. On opening, this tab shows the existing columns. For each column, the datatype (and if applicable the precision in brackets) and whether the table definition permits nulls in the target table column. Optionally you can:	
		Change the <b>Sort</b> order of a column.
		Click the New button and select a column name from the dropdown, to add a column to the primary key.
		Select a column and click the Drop button to delete the column from the primary key.
		Select a column and use the arrow buttons to reorder the columns in the primary key.
<b>Properties</b>	Lets you work with settings in the following categories:	
	<b>Attributes</b>	Lets you set the <b>Clustered</b> and <b>Maximum Rows Per Page</b> properties.
	<b>Storage</b>	Lets you set <b>Reserve Page Gap</b> , <b>Segment</b> (DEFAULT, LOGSEGMENT or SYSTEM), and <b>Fill Factor</b> properties.
<b>Statistics</b>	Lets you view statistics in the following categories:	
	<b>Page Statistics</b>	<b>Data Pages, Reserved Pages, Used Pages, OAM Page Ratio, and Index Page Ratio.</b>
	<b>Row Statistics</b>	<b>Maximum Row Size, Minimum Row Size, Max Size of Non-Leaf Index Row, and Maximum Rows Per Page.</b>
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PROCEDURES EDITOR (SYBASE)

The Procedures Editor lets you view and modify procedure definitions and modify associated privileges and dependencies.

### To edit a procedure

- 1 Open an editor on the procedure. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Definition</b>	Lets you modify the SQL code for a procedure.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## ROLES EDITOR (SYBASE)

The Roles Editor lets you manage authentication, logins, and privileges for a role.

**NOTE:** User-defined roles are available in Sybase ASE 11.5 or later.



**To edit a role**

- 1 Open an editor on the role. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Authentication</b>	Lets you manage role identification for a role. To identify the role, supply a password. Property categories are as follows:	
	<b>Identification</b>	Lets you specify a <b>Password</b> .
	<b>Security</b>	Lets you specify <b>Role Locked</b> .
	<b>Options</b>	Lets you specify <b>Password Expiration</b> , <b>Minimum Password Length</b> , and <b>Maximum Failed Logins</b> .
<b>Logins/Roles</b>	Displays the current permissions for this role to logins or other roles. Optionally, you can:	
		Click <b>Grant</b> to open a dialog that lets you grant this role to a login or another role.
		Select a role or login and click <b>Revoke</b> .
<b>Object Privileges and System Privileges</b>	For details on using these tabs, see <a href="#">Working with Privileges and Permissions</a> .	
<b>DDL</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**RULES EDITOR (SYBASE)**

The Rules Editor lets you modify basic properties of a rule.

**To edit a rule**

- 1 Open an editor on the rule. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you specify a <b>Restriction</b> and select a <b>Type</b> (STANDARD_RULE, AND_ACCESS_RULE, OR_ACCESS_RULE)
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## SEGMENTS EDITOR (SYBASE)

The Segments Editor lets you manage segment location, associated objects, segment space, and segment thresholds.

### To edit a segment

- 1 Open an editor on the segment. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Location</b>	Lets you <b>Drop</b> or <b>Extend</b> selected segments.	
<b>Objects</b>	Lets you manage database objects associated with the segment. Objects are organized in a tree structure. You can open an object editor on a table, index, or constraint by double-clicking that object	
<b>Space</b>	Lets you view <b>Segment Usage</b> , <b>Distribution of Segment Space</b> , <b>AutoGrowth Properties</b> , and <b>Object Space Usage</b> for the segment.	
	Shows the <b>Hysteresis</b> value for the segment and for each threshold for the segment, the number of <b>Free Pages</b> and an associated <b>Procedure Name</b> . Manage thresholds as follows:	
		Click the <b>Add</b> button to specify the number of free pages and procedure name for a new segment threshold
		Select an existing threshold and click the <b>Edit</b> button to modify the threshold's procedure name value
<b>Threshold</b>	Shows the <b>Hysteresis</b> value for the segment and for each threshold for the segment, the number of <b>Free Pages</b> and an associated <b>Procedure Name</b> . Manage thresholds as follows:	
		Click the <b>Add</b> button to specify the number of free pages and procedure name for a new segment threshold
		Select an existing threshold and click the <b>Edit</b> button to modify the threshold's procedure name value
		Select an existing threshold and click the <b>Drop</b> button to delete that threshold
<b>DDL</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## TABLES EDITOR (SYBASE)

The Tables Editor lets you manage columns, basic properties, constraints, storage and space, and partitions of a table.

**To edit a table**

- 1 Open an editor on the table. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Displays the currently defined columns in the table. For any selected column, the Property/Value list provides additional detail on that column. Available properties depend on the datatype you choose as well as on the property values you select:  <b>Computed</b> and <b>Computed Expression</b> - Let you define a column as a computed column and provide the computed column expression.  <b>Type</b> - Lets you select a datatype (depending on the type, additional properties such as <b>Size</b> , <b>Width</b> , and <b>Scale</b> may be available).  <b>Identity Column</b> - Select this check box to define the column as an identity column.  <b>Allow Nulls</b> - Select this check box to allow nulls in this column.  <b>Default Collation</b> - available for text/character datatypes, lets you specify a default collation.  <b>Default Value</b> - Lets you type a constant value or select a function returning a constant value to serve as the default for the column  <b>Default Binding</b> and <b>Rule Binding</b> - Let you bind a rule or default to a column.  Optionally, you can:	
		Select a column, and in the <b>Property/Value</b> list modify property values for that column.
		Click <b>Add Column</b> , provide a name for the new column, and set property values for the column.
		Select a column and click <b>Delete</b> to remove the column from the table.
<b>Properties</b>	Lets you work with properties in the following categories:	
	<b>Physical Storage</b>	Lets you set the <b>Segment Name</b> (DEFAULT, SYSTEM), <b>Maximum Rows Per Page</b> , <b>Reserve Page Gap</b> , and <b>Identity Gap</b> properties.
	<b>Cache Strategy</b>	Lets you set <b>MRU Replacement Strategy</b> and <b>Prefetch Strategy</b> properties.
	<b>Row Locking Strategy</b>	Lets you view the <b>Expected Row Size</b> property and set the <b>Lock Scheme</b> (ALLPAGES, DATAPAGES, ROWPAGES) property.
<b>Indexes</b>	Lets you manage indexes for a table. On opening, the list of current indexes for the table is displayed. Optionally, take one of the following actions:	
		Click <b>Add</b> to open a dialog that lets you add a new index to the table.
		Select an index and click <b>Edit</b> to open a dialog that lets you edit index properties.
		Select an index and click <b>Drop</b> to open a dialog that lets you remove the index from the table.

Tab	Settings and tasks
<b>Constraints</b>	Lets you manage constraints for the table. Constraints are grouped by type, under folders. Optionally take one of the following actions:
	Select a constraint type folder and click <b>Add</b> to open a dialog that lets you add a constraint of that type.
	Select a constraint and click <b>Edit</b> to open a dialog that lets you modify the constraint details.
	Select a constraint and click <b>Drop</b> to remove the constraint.
<b>Space</b>	Lets you view the table usage and the distribution of table space for a table. Optionally, double-click a slice in the pie chart for detailed statistics.
<b>Partitions</b>	Displays partition details for the table including the <b>Partition Type</b> (ROUNDROBIN, HASH, LIST, RANGE), partitioning <b>Columns</b> (name and datatype) and <b>Partition Definitions</b> (Name, values, segments).  You can also add or modify partition definitions. For more information, see <a href="#">Partitioning Sybase Tables</a> .
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDl for an Object</a> .

- When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## PARTITIONING SYBASE TABLES

The **Create Partition** and **Edit Partition** wizards let you work with table partitions.

### To create or edit table partitions

- Open an editor on a table. For details, see [Opening an Object Editor](#).
- On the **Partition** tab, click the **Create Partition** or **Edit Partition** button.
- Use the following table as a guide to understanding and modifying the settings on the tabs of this wizard:

Tab	Settings and tasks
<b>Properties</b>	In addition to displaying basic identification and ownership properties, this pane lets you select a <b>Partition Type</b> (ROUNDROBIN, HASH, LIST, RANGE).
<b>Columns</b> (not available for a <b>Partition Type</b> of ROUNDROBIN)	For each column to be added, click New to add a column and select a <b>Column</b> name.
<b>Range Definitions</b> (only available for a <b>Partition Type</b> of RANGE)	Lets you create an ordered list of partitions. For each partition, click New to open the <b>Add Partition Definition</b> wizard. Provide a <b>Name</b> , select a <b>Segment</b> , (DEFAULT, SYSTEM, LOGSEGMENT) and click <b>Add</b> .  The <b>Range Definitions</b> pane also lets you edit and remove partitions.

Tab	Settings and tasks
<b>Partitions Definitions</b> (only available for a <b>Partition Type</b> of HASH)	Lets you specify a Partition method: <b>None</b> <b>Number of Partitions</b> - Specify a <b>Number of Partitions</b> , and then for each partition, click the New button and select a segment. <b>By Partition Name</b> - For each partition, click New, specify a <b>Name</b> and select a <b>Segment</b> .
<b>List Definitions</b> (only available for a <b>Partition Type</b> of LIST)	Lets you create a list of partition definitions. For each partition, click New to open the <b>Add Partition Definition</b> wizard. Provide a <b>Name</b> , select a <b>Segment</b> , (DEFAULT, SYSTEM, LOGSEGMENT) and click <b>Add</b> . The <b>List Definitions</b> pane also lets you edit and remove partitions.

- 4 When ready, click **Finish**.

## TRIGGERS EDITOR (SYBASE)

The Triggers Editor lets you modify the basic properties of a trigger.

### To edit a trigger

- 1 Open an editor on the trigger. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Lets you specify whether the trigger is <b>Enabled</b> and lets you view the <b>Fire On Insert</b> , <b>Fire On Update</b> , and <b>Fire On Delete</b> properties.
<b>Definition</b>	Lets you modify the CREATE TRIGGER body for a trigger. To modify a trigger, Rapid SQL must drop then create the trigger.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## UNIQUE KEYS EDITOR (SYBASE)

The Unique Keys Editor lets you manage the columns and basic properties for a unique key as well as view page and row statistics.

**To edit a unique key**

- 1 Open an editor on the unique key. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Columns</b>	Lets you manage columns that make up the unique key. On opening, this tab shows the existing columns. For each column, the datatype (and if applicable the precision in brackets) and whether the table definition permits nulls in the target table column. Optionally you can:	
		Change the <b>Sort</b> order of a column.
		Click the New button and select a column name from the dropdown, to add a column to the index.
		Select a column and click the Drop button to delete the column from the index.
		Select a column and use the arrow buttons to reorder the columns in the index.
<b>Properties</b>	Lets you work with properties in the following categories:	
	<b>Attributes</b>	Lets you set <b>Clustered</b> and <b>Maximum Rows Per Page</b> properties.
	<b>Storage</b>	Lets you set the <b>Reserve Page Gap</b> , and <b>Segment</b> (DEFAULT, LOGSEGMENT, SYSTEM) properties.
<b>Statistics</b>	Lets you view statistics in the following categories:	
	<b>Page Statistics</b>	<b>Data Pages, Reserved Pages, Used Pages, OAM Page Ratio, and Index Page Ratio.</b>
	<b>Row Statistics</b>	<b>Maximum Row Size, Minimum Row Size, Max Size of Non-Leaf Index Row, and Maximum Rows Per Page.</b>
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

**USER DATATYPES EDITOR (SYBASE)**

The User Datatypes Editor lets you view and modify user datatype definitions and view objects associated with a user datatype.

### To edit a user datatype

- 1 Open an editor on the user datatype. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	Lets you work with properties in the following categories:	
	<b>Base datatype</b>	Lets you set <b>Type</b> , <b>Size</b> , <b>Allow Nulls</b> , and <b>Identity</b> settings.
	<b>Bindings</b>	Lets you set <b>Default Binding</b> and <b>Rule Binding</b> properties.
<b>Usage</b>	Displays all objects referring to or referenced by a user datatype.	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## USERS EDITOR (SYBASE)

The Users Editor lets you view and modify user definitions and manage associated logins and objects.

### To edit a user

- 1 Open an editor on the user. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks	
<b>Properties</b>	Lets you view the user <b>Name</b> and <b>Login Name</b> for the user. Lets you specify a <b>Group Name</b> for the user.	
<b>Logins</b>	Lets you assign logins to a user by selecting associated check boxes.	
<b>Objects</b>	Lets you manage database objects associated with the user. Objects are organized in a tree structure with folders containing the objects. Optionally, take one of the following actions:	
		Select an object and click <b>Edit</b> to open an object editor on the selected object.
		Select an object and click <b>Drop</b> to initiate dropping the selected object.
<b>Object Permissions and System Permissions</b>	For details on using these tab, see <a href="#">Working with Privileges and Permissions</a> .	
<b>Comment</b>	For details on using this tab, see <a href="#">Adding a Comment to an object</a> .	
<b>DDL View</b>	For details on using this tab, see <a href="#">Viewing the SQL/DDDL for an Object</a> .	

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).

## VIEWS EDITOR (SYBASE)

The Views Editor lets you view the columns for the view and work directly with the SQL to create the view.

### To edit a view

- 1 Open an editor on the view. For details, see [Opening an Object Editor](#).
- 2 Use the following table as a guide to understanding and modifying the settings on the tabs of this editor:

Tab	Settings and tasks
<b>Properties</b>	Displays all columns for the view. Details for each column include the <b>Column Name</b> , the <b>Datatype</b> (and if applicable, with the precision in parentheses), and whether or not <b>Nulls</b> are allowed for that column.
<b>Definition</b>	Displays the SQL (CREATE VIEW) for the view.
<b>Dependencies</b>	For details on using this tab, see <a href="#">Working with Object Dependencies</a> .
<b>Permissions</b>	For details on using this tab, see <a href="#">Working with Privileges and Permissions</a> .

- 3 When finished, you can submit your changes. For details, see [Previewing and Submitting Object Editor Changes](#).



## CREATING OBJECTS

Rapid SQL offers easy-to-use wizards and dialog boxes for creating objects. For an introduction to object creation wizards, see [Overview and Common Usage of Object Wizards](#).

Platform-by-platform discussions of wizards are discussed as follows:

- [IBM DB2 for Linux, Unix, and Windows Object Wizards](#)
- [IBM DB2 for z/OS Object Wizards](#)
- [InterBase/Firebird Object Wizards](#)
- [Microsoft SQL Server Object Wizards](#)
- [Oracle Object Wizards](#)
- [Sybase ASE Object Wizards](#)

## OVERVIEW AND COMMON USAGE OF OBJECT WIZARDS

An object wizard lets you create new database and server objects on a datasource. For example, a Tables wizard lets you create columns, set up permissions to work with that table, provide details on physical storage for the table, and so on. Each pane on an object wizard lets you perform a logical task or collection of logical tasks for that object type. For example:

**Create Index Wizard - Step 1 of 6**

**Properties**  
Specify the index properties.

Property	Value
<b>Creation</b>	
Table Schema	
Table Name	
Schema	
Name	
<b>Attributes</b>	
Index Type	NONUNIQUE
Clustered	False
Number of Partitions	0
Buffer Pool	
Defer	False
Close	True
Copy	False
Define	True
Piece Size (KB)	2097152

< Back   Next >   Finish   Cancel   Help

In order to work with object wizards, you must know how to:

- [Open an Object Wizard](#) for a server or database object.
- [Navigate and Set Properties in an Object Wizard](#)
- [Create the New Object](#).

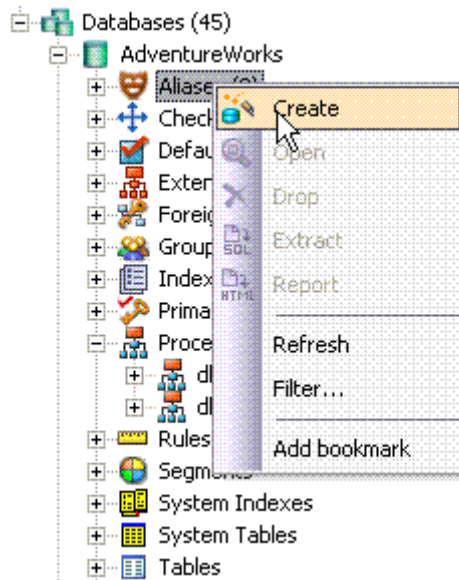
## OPEN AN OBJECT WIZARD

Object wizards are accessed from the Datasource Explorer pane of the Rapid SQL main window.

### To open an Object wizard for a server or database object

- 1 Connect to the datasource where you want to create a new resource. For more information, see [Datasource Management](#).
- 2 On the Datasource Explorer, expand the target datasource.
- 3 Continue to expand folders under the datasource until the type of object you want to create is visible.

- 4 On the Datasource Explorer, right-click on the specific type of object that you want to create and select **Create** from the context menu.



Rapid SQL opens the object wizard for that object type.

## NAVIGATE AND SET PROPERTIES IN AN OBJECT WIZARD

When you invoke an object wizard, it opens on the first pane of the wizard, typically a **Properties** pane. As you select options or provide details, you use the navigation buttons at the bottom of the window and the pane controls at the left of the window to navigate through the wizard.

- Use the **Next** button to move to the next pane of the wizard.

**NOTE:** In some cases, the **Next** button is not activated until required information is provided. Similarly, some panes of a wizard do not become accessible until activated by choice of an option or property on a previous pane. For example, a **Partitions** tab in a table wizard may only become available if the **Clustered** and **Number of Partitions** options on a prior tab are set accordingly.

- Use the **Back** button to move to the previous pane of the wizard.
- Use the pane controls at the left of the window to move to a specific pane of the wizard.

## ADD A COMMENT TO AN OBJECT

The object wizards for certain object types feature a **Comment** tab that lets you add an explanatory note to specific object definitions. Comments are stored in the REMARKS column of the object's system-catalog.

**To add a comment to an object:**

- 1 [Open an Object Wizard](#) on an object type that permits comments. See the topics for specific Object editors later in this chapter for information on whether that object type supports comments.
- 2 Click the **Comment** panel when enabled.

The screenshot shows a user interface for adding a comment. At the top, there is a tab labeled "Comment". Below the tab is a large text input area. Inside this area, the text "Comment:" is followed by a small box containing the number "1".

- 3 In the **Comment** area, type an explanatory note of up to 254 characters long.

**SET PERMISSIONS FOR AN OBJECT**

When you open an Object wizard to create an object with associated privileges, the **Permissions** (or **Privileges** or **Object Permissions** or **System Permissions**) panel for that editor displays the relevant privileges and lets you make changes accordingly:

**To set permissions for an object:**

- 1 [Open an Object Wizard](#) on an object type with associated permissions or privileges. See the topics for specific Object wizards later in this chapter for information on whether that object type supports permissions/privileges.
- 2 Click the **Permissions (Privileges, Object Permissions or System Permissions)** panel when enabled.
- 3 For each specific permission to be granted to an entity such as a user, login, or group, select the cell corresponding to the entity and specific permission, and click the **Grant** button. To revoke a privilege, select a cell showing a Granted permission and click **Revoke**.

**CREATE THE NEW OBJECT**

The final pane of a wizard is most commonly a **DDL View** or **Definition** pane. That pane lets you view the SQL or DDL generated to create the object with which you are working. In most cases you arrive at that pane by navigating through panes, choosing options as you proceed. Alternatively, if you have provided all required information for that object type, you can click **Finish** to view the generated SQL or DDL.

You can inspect the SQL or DDL in a **DDL View** or **Definition** pane and if necessary navigate to other panes to modify your option choices. In some **Definition** panes, you have to complete the DDL or SQL to create the object.

Lastly, when your option selections are complete, use the **Execute** button to create the new object.

## IBM DB2 FOR LINUX, UNIX, AND WINDOWS OBJECT WIZARDS

Rapid SQL lets you create DB2 LUW objects using the following wizards:

- [Alias Wizard \(DB2 LUW\)](#)
- [Database Wizard \(IBM DB2 LUW\)](#)
- [Foreign Key Wizard \(IBM DB2 LUW\)](#)
- [Function Wizard \(IBM DB2 LUW\)](#)
- [Index Wizard \(DB2 LUW\)](#)
- [Materialized Query Table Wizard \(DB2 LUW\)](#)
- [Nodegroup Wizard \(DB2 LUW\)](#)
- [Primary Key Wizard \(DB2 LUW\)](#)
- [Procedure Wizard \(DB2 LUW\)](#)
- [Schema Wizard \(DB2 LUW\)](#)
- [Sequence Wizard \(DB2 LUW\)](#)
- [Structured Type Wizard \(DB2 LUW\)](#)
- [Table Wizard \(DB2 LUW\)](#)
- [Tablespace Wizard \(DB2 LUW\)](#)
- [Trigger Wizard \(DB2 LUW\)](#)
- [Unique Key Wizard \(DB2 LUW\)](#)
- [User Datatype \(DB2 LUW\)](#)
- [User Wizard \(DB2 LUW\)](#)
- [View Wizard \(DB2 LUW\)](#)

In addition, see [Create Synonym](#).

### ALIAS WIZARD (DB2 LUW)

An alias offers you security and convenience so that you can refer to an object without revealing who owns it or what database it belongs to. You can create aliases for tables, views, and even other aliases. The Alias Wizard lets you create an alias without knowing the underlying commands. As you complete the Alias Wizard process, Rapid SQL constructs the necessary CREATE ALIAS statement based on the information that you supply.

**To create a new alias using a wizard:**

- 1 [Open an Object Wizard](#) for an alias.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details see [Aliases \(DB2 LUW\) - Properties](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**Aliases (DB2 LUW) - Properties**

When creating or editing an alias, this tab/panel lets you work with the following settings:

Setting	Description
<b>Schema</b>	Select the schema that is to own the alias.
<b>Name</b>	Provide a name for the alias
<b>Target Owner</b>	Select the owner of the object to which you are creating an alias.
<b>Target Type</b>	Select the type of object to which you want to create an alias.
<b>Target Name</b>	Select the specific object to which you are creating an alias.

**DATABASE WIZARD (IBM DB2 LUW)**

The Database Wizard lets you create a database (a structured collection of data that can be updated or queried) without knowing the underlying commands. Databases can be simple, that is one file with many records and the same fields, or much more complicated with multiple files with different fields.

**To open the IBM DB2 for Linux, Unix, and Windows Database Wizard**

On the *Registration* tool bar, click New UDB Database.

OR

On the *Datasource Explorer*, right-click an instance node or the Datasources node, and then click New UDB Database.

The tables below describe the fields you may encounter as you complete the wizard.

## Naming the New Database

Required Information	Description
What is the name of the database?	Each database requires a unique name. The database name should be between 1 - 8 characters and must begin with a letter A-Z. It's wise to avoid the symbols @, #, and \$ if the database will exist in a communications environment or be summoned from an area where those characters do not exist in the language (and hence, on a keyboard).
What is the datasource name for the new database?	Each datasource should have a unique name. This is the name that will appear on the Explorer.
What is the alias name of the database.	OPTIONAL: A database alias is the local synonym given to the database you are creating, and must be unique. If you don't assign an alias, the database name becomes the default alias.  Note that a DB2 client can have connections to two different databases with the same name if those databases live on different servers and each database has its own alias.
What is the comment used for the database?	OPTIONAL: Lets you type a comment up to 30 characters. Any comment you enter can be changed later.

## Drive Path and Parameters

Required Information	Description
On what drive/path will the database reside?	OPTIONAL: Leave blank if you want Rapid SQL to create the database using the DFTBPATH (default database path configuration) parameter.
What default tablespace parameters should be used?	OPTIONAL Extent size: This is the number of pages of table data that will be written before data is written to the next container. Number of segments: The default is 0. The number you designate specifies the number of extended memory segments available for the database to use. You should only set this number if you have a large amount of real memory.
What global default parameters should the database use?	Territory: The territory where the database was created. This is purely informational, for example, en_US and is related to region-specific support. Codeset: IBM-1252 or UTF-8 are the options for language support. Collating Sequence: Compatibility, Identity, System are your choices for comparing character data. Note: The collation sequence can't be changed after the database has been created.
Finish	Opens the <b>Preview SQL</b> dialog box. For more information, see <a href="#">Preview</a> .

## Add Catalog/Temporary/User Containers

There are three distinct wizard panels, one for each container type, that enable you to customize your database. Click Next or Finish depending on how you want to proceed.

Required Information	Description
What container(s) will be used to store the catalog/temporary tables?	<p>Use System Managed Space: If you select this option, the operating system's file system manager allocates and manages the space where the catalog tables are stored. This is the default.</p> <p>Use Database Managed Space: If you select this option, the database manager controls the storage space for catalog tables.</p> <p>Add or Edit - For details, see <a href="#">Add/Edit Container for Tablespace</a>.</p>
What optional default storage parameters should be used.	Optionally identify values for extent size, pre-fetch size, transfer rate, and overhead.
Finish	Opens the <b>Preview SQL</b> dialog box. For more information, see <a href="#">Preview</a> .

## ADD/EDIT CONTAINER FOR TABLESPACE

The table below describes the options and functionality on this dialog box:

Required Information	Description
<b>Container Parameters</b>	These parameters are enabled for modification only if you have opted to use database managed space. When enabled, you can either indicate file size parameters or specify device information.
<b>Directory</b>	Identify the directory where the catalog table container will be stored.
<b>File Name</b>	Identify a name for the file where the catalog table container will be stored.
<b>Device Information</b>	Enable either raw partition and type a value, or raw drive, and select a value.

## FOREIGN KEY WIZARD (IBM DB2 LUW)

Foreign keys are unique values that refer to specific columns of other tables. Thus, a foreign key links two tables together. Embarcadero Rapid SQL's Foreign Key Wizard makes it easy for you to create a relational link between two tables, thereby speeding queries and giving you faster access to data. The Foreign Key Wizard lets you create a foreign key without knowing the underlying commands.



**To create a new foreign key using a wizard:**

- 1 [Open an Object Wizard](#) for a foreign key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Foreign Keys \(DB2 LUW\) - Properties](#)
  - **Column Mapping** panel - for details, see [Foreign Keys \(DB2 LUW\) - Column Mapping](#)
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**Foreign Keys (DB2 LUW) - Properties**

When creating or editing a foreign key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Schema</b>	The schema owning the table where the foreign key is being created.
<b>Table Name</b>	This is the table where the foreign key link originates--the child table.
<b>Name</b>	Lets you select a constraint name. System Generated Name - DB2 automatically generates a name. User Specified Constraint Name - You type the name.
<b>Enabled</b>	TRUE enables the Foreign Key while FALSE disables the Foreign Key.
<b>Delete Rule</b>	Select an action: <b>NO ACTION</b> - ensures that referenced values cannot be updated or deleted if to do so would violate referential integrity. <b>CASCADE</b> permits a referenced row in a child table to be deleted/updated if it is deleted/updated in the parent table. A row in the child table is SET NULL when rows in the parent table are deleted/updated. <b>RESTRICT</b> prevents a command from executing if changes to referential data prompts an error.
<b>Update Rule</b>	Select an action. <b>NO ACTION</b> - ensures that referenced values cannot be updated or deleted if to do so would violate referential integrity. <b>RESTRICT</b> - prevents a command from executing if changes to referential data prompts an error

**Foreign Keys (DB2 LUW) - Column Mapping**

- 1 Under **Referenced Table**, choose the **Owner** and then the **Name** of the referenced, or parent, table.
- 2 Under the **Main Table**, select checkboxes corresponding to the columns that are to reference columns in the referenced table.

## FUNCTION WIZARD (IBM DB2 LUW)

To create a relationship between one set of values and another, Rapid SQL offers the Function Wizard. You can develop reusable subroutines so you can control, access, and manipulate the data that underlies an object. As you complete the Function Wizard process, Rapid SQL constructs the necessary CREATE FUNCTION statement based on the information that you supply. The Function Wizard lets you create a function without knowing the underlying commands.

**NOTE:** To create a user-defined function, you need CREATE ANY privileges or IMPLICIT\_SCHEMA authority on the database if the schema does not already exist.

### To create a new function using a wizard:

- 1 [Open an Object Wizard](#) for a function.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Functions \(DB2 LUW\) - Properties](#).
  - **Advanced** panel - for details, see [Functions \(DB2 LUW\) - Advanced](#).
  - **Source** panel (only available for a **Function Type** of SOURCED) - for details, see [Functions \(DB2 LUW\) - Source](#).
  - **Parameters** panel (only available for a **Function Type** of SQL, SOURCED, or TEMPLATE) - for details, see [Functions \(DB2 LUW\) - Parameters](#).
  - **Return Scalar** panel - for details, see [Functions \(DB2 LUW\) - Return Scalar](#).
  - **Return Columns** panel (only available when you choose a **Function Type** of EXTERNAL TABLE) - for details, see [Functions \(DB2 LUW\) - Return Columns](#).
  - **Body** panel - for details, see [Functions \(DB2 LUW\) - Body](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

### FUNCTIONS (DB2 LUW) - PROPERTIES

When creating or editing a foreign key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Schema, Name, and Specific Name</b>	Let you select the owner of the function, provide a name for the function, and provide the Specific name to be used by some SQL statements and DB2 commands for this function.

Setting	Description
<b>Function Type</b>	Select the type of function: External Scalar - written in a programming language and returns a scalar value. External Table - written in a programming language and returns a complete table. OLEDB - accesses OLE DB data in user-defined OLE DB external tables. Sourced - another function is invoked to implement the function you are creating. SQL Language - written in SQL and returns a table, scalar value, or single row. Template: - this is a partial function and can only be invoked from a federated datasource. For more information, see <a href="#">About function types</a> .
<b>Language</b>	If you chose a <b>Function Type</b> of EXTERNAL SCALAR or EXTERNAL TABLE, specify a language of C, JAVA, or OLE. For more information, see <a href="#">About function types</a> .
<b>Return Type</b>	For a <b>Function Type</b> of SQL, select ROW, TABLE, or SCALAR. For other <b>Function Type</b> choices, this indicates the default return type for that choice.
<b>External Name</b>	Provide the External Name of the function.
<b>SQL Access Level</b>	Indicates whether the function can execute SQL statements. CONTAINS SQL: Statements that don't read or modify SQL can be executed. NO SQL: No SQL statements can be executed. READS SQL: Statements that cannot modify SQL can be executed.

### Functions (DB2 LUW) - Advanced

This tab is only available after clicking the **Advanced** button on the Function wizard's **Properties** panel. It lets you work with the **Threadsafe**, **Fenced**, **Scratchpad**, **Scratchpad Length**, **Allow Parallel**, **Final Call**, **Parameter Style**, **Inherit Special Registers**, **DBINFO**, **Deterministic**, **External Action**, **Called on Null Input**, and **Parameter CCSID** properties.

### Functions (DB2 LUW) - Source

**NOTE:** This panel is only available for a **Function Type** of SOURCED.

Select the **Schema**, **Name**, and **Specific Name** of the source function. Function **Parameters** and **Return Type** for the selected function are displayed. For more information, see [About function types](#).

### Functions (DB2 LUW) - Parameters

**NOTE:** This panel is only available for a **Function Type** of SQL, SOURCED, or TEMPLATE.

For each parameter for this function, use the New button to add a new parameter, provide a name for the parameter, and in the **Attributes** area, select a **Type**, and if appropriate, the **Precision**, **Scale**, and **Size** options.

### Functions (DB2 LUW) - Return Scalar

Under **Return Datatype**, select a Type and depending on your choice, provide or select **Precision**, **Scale**, **Size**, and **As Locator** options.

To make use of a CAST FROM clause, under **Cast Datatype** set **Enabled** to True, select a **Type**, and if appropriate, the **Scale**, **Size**, and **As Locator** options

### Functions (DB2 LUW) - Return Columns

**NOTE:** This panel is only available when you choose a **Function Type** of EXTERNAL TABLE.

For each column returned by this function, use the New button to add a new parameter, provide a name for the parameter, and in the **Attributes** area, select a **Type**, and if appropriate, the **Precision**, **Scale**, **Size**, and **As Locator** options.

### Functions (DB2 LUW) - Body

Enter the return statement for the function.

## ABOUT FUNCTION TYPES

### External Scalar/Table/OLE DB Function

External scalar user-defined functions are implemented in an external programming language. The functions are executed on the server and can read SQL data but cannot make changes to the data. These functions are often used to extend the set of built-in functions for DB2, perform logic inside a SQL query that SQL can't perform on its own, and, encapsulate a scalar query that is often used as a subquery in SQL statements, for example, if given an ingredient, search a table for a recipe that uses that ingredient.

When specifying the **Specific Name** for one of these function types:

- If you are using C language, specify the full library path and the function name, otherwise IBM DB2 Database Manager assumes the function is under the IBM DB2 library.
- If you are using Java script, specify the Class ID and the function name, otherwise IBM DB2 Database Manager assumes the function is under the IBM DB2 library.
- If you are using OLE language, specify the full library path and the function name, otherwise IBM DB2 Database Manager assumes the function is under the IBM DB2 library.

### Sourced Functions

When you create a sourced function, the new function you are creating will be implemented by a preexisting (source) function that's known to the database manager. The source function can be a built-in function or a previously created user-defined scalar function.

When you select the appropriate schema, you are really specifying an implicit schema privilege. In other words, you're selecting a schema/function that belongs to a user with DBAdmin privileges. If you want to use a built-in function, you must specify the function's specific name.

## SQL Language Function

The function you are creating is written in SQL. A table, scalar value, or single row is returned.

## Template Function

A template function is better thought of as a template for a function. It's a partial function that has no executable code. You create a function template for mapping it to a datasource function. When the mapping is created, you can specify the function template be used in queries submitted to the federated server. When the query is processed, the federated server invokes datasource function where the template is mapped and returns the appropriate values.

## INDEX WIZARD (DB2 LUW)

Comparable to an index in a book, an index gives you speedy access to particular records in a table. The Index Wizard lets you create an index without knowing the underlying commands.

### To create a new index using a wizard:

- 1 [Open an Object Wizard](#) for an index.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Indexes \(DB2 LUW\) - Properties](#).
  - **Columns** and **Include Columns** panels - for details, see [Indexes \(DB2 LUW\) - Columns and Include Columns](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## INDEXES (DB2 LUW) - PROPERTIES

When creating or editing an index, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Schema</b> and <b>Table Name</b>	Choose the owner and name of the table in which the index is being created.
<b>Schema</b> and <b>Name</b>	Choose the owner and name of the index being created.
<b>Index Type</b>	Index enforces uniqueness on the values of the table's index key.
<b>Clustered</b>	Specifies that the index is the clustering index of the table. The cluster factor of a clustering index is maintained or improved dynamically as data is inserted into the associated table, by attempting to insert new rows physically close to the rows for which the key values of this index are in the same range.

Setting	Description
<b>Percent Free</b>	Lets you type or select the percentage of each index page to leave as free space when building the index, from 0 to 99.

### INDEXES (DB2 LUW) - COLUMNS AND INCLUDE COLUMNS

Index columns can be segregated into unique key columns (**Columns** pane) and Include columns that are to be part of the index but do not form part of the unique key.

The steps in completing the panes for the two column types are identical.

- From the **Column** dropdown, select a column for the index and specify a **Sort** option. To add more columns, click the New button and then follow the steps in the last instruction. Use the Delete button to drop columns.

### MATERIALIZED QUERY TABLE WIZARD (DB2 LUW)

A materialized query table (MQT) is a table based on the result of a query. An MQT contains information that is summarized from other tables and can save time when processing dynamic SQL queries. The Materialized Query Table Wizard lets you create a table without knowing the underlying commands.

#### To Open the Materialized Query Wizard

- 1 On the Explorer, find the datasource where you want to add the new materialized query table.
- 2 Expand the Schema branch, right-click Materialized Query Tables, and select New.

The table that follows describes the fields you may encounter as you complete the Materialized Query Table Wizard.

**NOTE:** These options are only available if the tablespace you selected is a database managed tablespace.

Required Information	Description
Who owns the table?	Choose the owner of the table you're creating from the drop-down list.
What is the name of the table?	Type the name of the materialized query table you are creating.
Select a tablespace on which to place the table:	OPTIONAL: No selection is the default. But you can select a tablespace that belongs to the new table's owner if you want.
Specify separate tablespaces for index and long data	OPTIONAL: Lets you separate indexes or long data from the table. Indexes Long data

Required Information	Description
Specify the query on which this table is based	Write the query you want to use to create the table. Note: Every select list element must have a name (use the AS clause for expressions)
Add the columns belonging to this table Add the columns belonging to the partition key	Click <b>Add</b> , <b>Insert</b> or <b>Edit</b> to add or modify table columns. Click <b>Drop</b> to delete a selected column.
Do you want the table replicated across database partitions?	The default is no, but check the box if you do want to replicate the table. Replicated materialized query tables can help you improve query performance by enabling collocation between tables. Replication is especially helpful when you have joins between large fact tables and small dimension tables. It's best if tables that are to be replicated are infrequently updated.
Definition Only	Lets you select definition options, Include Column Defaults and Include Identity Column Attributes. When you select the Definition Only option, the new table is treated as if it was a base table rather than a MQT. After you complete the wizard, Rapid SQL opens the Tables Editor.
Refreshable	Lets you select refresh options: Immediate: The table you are creating will be refreshed automatically when changes are made to the base table(s). Deferred: Static SQL will not be optimized. Changes to the base table(s) will not be reflected. Query Optimization: Enabled: Queries will be routed to the MQT. Disabled: This is the default. Maintained by: System: This is the default User After you complete the wizard, Rapid SQL opens the Material Query Tables Editor.
Would you like extra information regarding SQL changes to this table to be written to the log?	If you opted to replicate the table, you must make a selection here. Yes means you want to receive error messages issued by the DB2 replication programs. Include Longvar Columns means you want to receive error messages for these columns if long varchar data is a part of the table you're replicating. No
Would you like to skip logging changes made to this table by Insert... and Alter Table operations in the same unit of work in which this table is created?	Initially Not Logged: This is an option that can reduce logging and increase performance, but also means that you cannot recover the table when rolling forward. Logged Initially. This is the default.
What type of table lock would you like when it is being accessed?	Row: This is the default. During a table or index scan, DB2 locks each row that is scanned before determining whether that row is relevant to the query. Table: During a table scan, DB2 locks the table so no data can be added or altered while the query is executed.

Required Information	Description
What percent of free space to leave for load and reorganization?	-1 is the default.
Do you want data to append to the end of the table?	Yes No: This is the default
Do you want the access plan to this table to be based on existing statistics and optimization level?	Volatile: A volatile table's contents can vary from empty to huge at run time and can render collected statistics inaccurate. Not Volatile: This is the default.
Enter a comment	Optional

## NODEGROUP WIZARD (DB2 LUW)

A nodegroup is a named subset of one or more database partitions. The Node Group Wizard lets you create a node group without knowing the underlying commands. When you create a nodegroup, the wizard simply asks you to name the nodegroup and select the partitions to include in the nodegroup.

### To Open the Nodegroup Wizard

- 1 On the Explorer, find the datasource where you want to add the new Nodegroup.
- 2 Expand the Storage branch, right-click Nodegroups, and select New.

## PRIMARY KEY WIZARD (DB2 LUW)

A primary key is a column or group of columns that you can use to identify or access one or more specific rows in a table. A primary key is 'constrained' in that no values can be null and no two values are equal. You can only create one primary key for any table. The Create Primary Key Constraint dialog box lets you create a primary key without knowing the underlying commands.

When you create a primary key, specify the table owner and the table on which you want to place the primary key constraint. You name the constraint and select the column(s) you want to include.



**To create a new primary key using a wizard:**

- 1 [Open an Object Wizard](#) for a primary key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Primary Keys \(DB2 LUW\) - Properties](#).
  - **Columns** panel - for details, see [Primary Keys \(DB2 LUW\) - Columns](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**PRIMARY KEYS (DB2 LUW) - PROPERTIES**

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Schema</b> and <b>Table Name</b>	Choose the owner and name of the table in which the primary key is being created.
<b>Name</b>	Choose the owner and name of the primary key being created.

**PRIMARY KEYS (DB2 LUW) - COLUMNS**

From the **Column** dropdown, select a column for the primary key and specify a **Sort** option. To add more columns, click the New button and then follow the steps in the last instruction. Use the Delete button to drop columns.

**PROCEDURE WIZARD (DB2 LUW)**

Procedures are a reusable block of PL/SQL, stored in the database, that applications can call. Procedures streamline code development, debugging, and maintenance by being reusable. Procedures enhance database security by letting you write procedures granting users execution privileges to tables rather than letting them access tables directly.

**To create a new procedure using a wizard:**

- 1 [Open an Object Wizard](#) for a procedure.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Procedures \(DB2 LUW\) - Properties](#).
  - **Advanced** panel - for details, see [Procedures \(DB2 LUW\) - Advanced](#).
  - **Parameters** panel - for details, see [Procedures \(DB2 LUW\) - Parameters](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## PROCEDURES (DB2 LUW) - PROPERTIES

When creating or editing a procedure, this tab/panel lets you work with the following settings:

Setting	Description
<b>Schema</b>	Select the owner for the procedure.
<b>Name</b>	Provide the name of the function.
<b>Specific Name</b>	Optionally, provide the unique name of the procedure.
<b>Language</b>	Select among C, JAVA, COBOL, OLE, or SQL. The database manager will call the procedure accordingly.
<b>External Name</b>	Provide the external name of the procedure.
<b>SQL Access Level</b>	Select an option: MODIFIES SQL DATA - the procedure can support any SQL statement except those that cannot be supported in procedures. CONTAINS SQL DATA - only SQL statements that neither modify nor read SQL data can be executed in the procedure. READS SQL DATA - some SQL statements that don't modify SQL data can be included in the procedure.

## PROCEDURES (DB2 LUW) - ADVANCED

When creating or editing a procedure, this tab/panel lets you work with the following settings:

Setting	Description
<b>Results Sets</b>	Indicate the estimated upper bound of returned result sets. 0 is the default.
<b>External Action</b>	Select the External Action option.
<b>New Save Point</b>	Lets you specify a NEW SAVEPOINT LEVEL clause for the procedure.
<b>Threadsafe</b>	Specify whether the procedure is safe to run within the same process as other routines.

Setting	Description
<b>Fenced</b>	<p>If you select yes, you are saying you do not want the procedure to run in the manager operating system environment. This means the database management system will protect its internal resources from the procedure. This option can affect the procedure's operation.</p> <p>To run a procedure as not fenced, or a No selection, you must have SYSADMIN or DBADMIN privileges because of the potential to compromise data if the procedure has not been adequately tested.</p>
<b>Parameter Style</b>	<p>Lets you select an option: DB2DARI, DB2GENERAL, DB2SQL, GENERAL, GENERAL WITH NULLS, JAVA, and SQL.</p> <p>DB2GENERAL is for Java Language only.</p> <p>DB2SQL is for C, COBOL, or OLE Language only.</p> <p>GENERAL is for C Language only.</p> <p>GENERAL WITH NULLS is for C or COBOL Language only.</p> <p>JAVA is for Java Language only.</p> <p>SQL is for C, COBOL, or OLE Language only.</p>
<b>Program Type</b>	<p>MAIN: valid for C or COBOL Language and Parameter Style GENERAL, GENERAL WITH NULLS, SQL, or DB2SQL only. In this case, parameters will be passed as an argument counter or argument vector.</p> <p>SUBROUTINE: the procedure expects the parameters to be passed as separate arguments.</p>
<b>Inherit Special Registers</b>	Lets you specify this optional clause dictating that the procedure will inherit initial values from the environment of the invoking statement.
<b>DBINFO</b>	Specific information contains such information as the database name, application ID, database code page, and so on.
<b>Deterministic</b>	Enabling this feature specifies the procedure will always return the same result for given argument values. Disabling it means there are state values that affect the results and so the same result will not always be returned when you use identical inputs to invoke the procedure.
Parameter CCSID	Select an encoding scheme of ASCII, UNICODE, or NONE for character or graphic string parameters.

### PROCEDURES (DB2 LUW) - PARAMETERS

For each parameter for this function, use the New button to add a new parameter, provide a name for the parameter, and in the **Attributes** area, select a **Type**, specify a **Parameter Mode** of INPUT, OUTPUT, or INPUT\_OUTPUT, and if appropriate, the **Precision**, **Scale**, and **Size** options.

### SEQUENCE WIZARD (DB2 LUW)

A sequence allows the automatic generation of values, something well-suited to the generation of unique key values. Sequences are not tied to particular table columns. The Sequence Wizard lets you create a sequence without knowing the underlying commands. As you complete the Sequence Wizard, Rapid SQL constructs the necessary CREATE SEQUENCE statement from the information that you have supplied. When finished, you can instruct Rapid SQL to compile the sequence on the target Rapid SQL database or to write a script file containing the CREATE SEQUENCE statement.

The Sequence Wizard lets you:

- Specify the name and owner of the sequence.
- Set both the value of the sequence, and an interval and ranges for incrementing it.
- Cache the sequence, cycle the sequence when it reaches its minimum or maximum values, and guarantee that Rapid SQL generates sequence numbers in the order of request.

**NOTE:** To create a sequence, it must belong to your schema or you need CREATE SEQUENCE privileges.

### To Open the Sequence Wizard

- 1 On the Explorer, find the datasource where you want to add the new Sequence.
- 2 Expand the Schema branch, right-click Sequences, and select New.

The table that follows describes what you may encounter as you complete the Sequence Wizard:

Required Information	Description
Who owns the sequence?	You decide.
What is the sequence name?	Your choice.
What numeric datatype should the Sequence use?	Choose among BIGINT (big integer), decimal (choose width as well), integer, small integer.
What is the first sequence number to be generated?	Starting with 1 is the default.
What is the interval between sequence numbers?	Increment by 1 is the default.
What is the sequence's minimum value?	Choose none or set a value
What is the sequence's maximum value?	Choose none or set a value
Should DB2 preallocate sequence numbers and cache them for faster access?	Preallocating and storing values in the cache reduces synchronous I/O to the log when values are generated for the sequence. If Yes, give number of values No
Should the sequence continue to generate values after reaching either its maximum or minimum value?	Lets you make the sequence cycle and continue to generate numbers. Yes or No
Should the sequence numbers be generated in the order of request?	Select to generate sequence numbers in the order of request. The ORDER option is useful when you are using the sequence number as a timestamp. Yes or No

## SCHEMA WIZARD (DB2 LUW)

The Schema Wizard lets you create the structure of a database system including database objects.

### To Open the Schema Wizard

- 1 On the Explorer, find the datasource where you want to add the new Schema.
- 2 Right-click Schema and select New.

All you need to do when the single-panel wizard opens is to give a unique name to the schema you're creating.

## STRUCTURED TYPE WIZARD (DB2 LUW)

The Structured Type wizard lets you create the object type specification and define attributes for a structured type. You then use the Structured Type editor to define the methods and body for the type.

### To Open the Type Wizard

- 1 On the Explorer, find the datasource where you want to add the new Structured Type.
- 2 Expand the Schema branch, right-click Structured Type, and select Create.
- 3 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Structured Types \(DB2 LUW\) - Properties](#).
  - **Attributes** panel - for details, see [Structured Types \(DB2 LUW\) - Attributes](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 4 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

On completion, the Structured Type editor opens, letting you provide method and body specifications for the structured type. For details, see [Structured Types Editor \(IBM DB2 LUW\)](#).

## STRUCTURED TYPES (DB2 LUW) - PROPERTIES

When creating or editing a structured type, this tab/panel lets you work with the following settings:

Setting	Description
<b>Attributes</b>	<p><b>Instantiable</b> - if disabled, no constructor function can be generated, nor can a non-instantiable type be used as the type of a table or view.</p> <p><b>Final Type</b> - enabled, indicates that the structured type cannot be used as a supertype.</p> <p><b>With Function Access</b> - If enabled, all methods of the type you are creating, and those you will create, can be accessed using functional notation. Some applications prefer functional notation over method invocation.</p> <p><b>Without Comparisons</b> - If enabled, indicates that there are no comparison functions supported for instances of the structured type.</p> <p><b>Inline Length</b> - Specifies the maximum size of an instance of a structured type in a column. If the size of a structured type instance is less than the defined maximum, the data is stored inline with the other values in the row. If the size of the structured type is larger than the defined maximum, the structured type data is stored outside of the table (like LOBs).</p>
<b>Supertype</b>	<p><b>Supertype Schema</b> and <b>Supertype Name</b> - let you designate this structured type as a subtype by providing details of the owning supertype.</p>
<b>Reference</b>	<p><b>Cast (Source as Ref) With</b> and <b>Cast (Ref as Source) With</b> - Lets you name the cast function, although one will be created with the default name of the structured type you are creating. The cast function "casts" a value between the reference type and the representation type in both directions.</p> <p><b>Reference Using, Size, Precision, and Scale</b> - Define the built-in data type used as the underlying data type for the structured type you are creating and all its subtypes.</p>

## STRUCTURED TYPES (DB2 LUW) - ATTRIBUTES

For each attribute to be added to the structured type, click the New button, and provide a name for the attribute.

With an attribute selected, you can modify the **Datatype**. Depending on the datatype you choose you can provide the following datatype qualifiers:

- **Width** (decimal only)
- **Scale** (decimal only)
- **Size** (blob, character, clob, dbclob, graphic, varchar, vargraphic only)
- **Allow log** (blob, clob, dbclob only)
- **Allow compact** (blob, clob, dbclob only)
- **For bit data** (character, long varchar, varchar, only)

## TABLE WIZARD (DB2 LUW)

The Table Wizard lets you create a table without knowing the underlying commands.

**To create a new table using a wizard:**

- 1 [Open an Object Wizard](#) for an table.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Tables \(DB2 LUW\) - Properties](#).
  - **Columns** panel - for details, see [Tables \(DB2 LUW\) - Columns](#).
  - **Partition** panel - for details, see [Tables \(DB2 LUW\) - Partition](#).
  - **Tablespaces** panel - for details, see [Tables \(DB2 LUW\) - Tablespaces](#).
  - **Dimension** panel - for details, see [Tables \(DB2 LUW\) - Dimensions](#).
  - **Distribution Key Columns** panel - for details, see [Tables \(DB2 LUW\) - Distribution Key Columns](#).
  - **Indexes** panel - for details, see [Tables \(DB2 LUW\) - Indexes](#).
  - **Constraints** panel - for details, see [Tables \(DB2 LUW\) - Constraints](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **Permissions** panel - [Set Permissions for an Object](#)
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**TABLES (DB2 LUW) - PROPERTIES**

When creating or editing a table, this tab/panel lets you work with the following settings:

Setting	Description
<b>Schema</b>	Select the schema that is to own the table.
<b>Name</b>	Provide a name for the table.
<b>Percent Free</b>	
<b>Lock Size</b>	The table-level lock prevents concurrently operating applications from using or changing a table. When scanning a table for a query, a row-level lock locks the row when it is being assessed.
<b>Append Data</b>	Enable or Disable appending data to the end of the table.
<b>Volatile</b>	Enable this feature if a table contents may fluctuate from empty to very large. The access plan will not depend on the existing statistics for that table.
<b>Compress</b>	Enable or disable value compression
<b>Row Compression</b>	Enable or disable row compression.
<b>Security Policy</b>	Lets you add a security policy to a table.
<b>RestrictDrop</b>	Corresponds to the DB2 Restrict on Drop attribute.
<b>Log Index Build</b>	Enables this level of logging wne creating, recreating, or reorganizing an index.

Setting	Description
<b>CCSID</b>	Specify ASCII or UNICODE or leave unspecified. If specified, this is the encoding scheme for string data. If unspecified, CCSID default encoding is used.
<b>Tablespace, Index Tablespace, and Long Data Tablespace</b>	Select a tablespace, an index tablespace, and a tablespace for Long or LOB table columns.
<b>Do not initially log</b>	If enabled, all changes to the table will be flushed out at commit time. This also means that if a statement fails, the unit of work will rollback. If you are concerned about recoverability, disable this feature.
<b>Data Capture</b>	Specify additional information logged by selecting DATACAPTURE NONE, DATA CAPTURE CHANGES, or DATA CAPTURE CHANGES INCLUDE LONGVAR.

#### TABLES (DB2 LUW) - COLUMNS

For each column in the table, click the New button to create a column and provide a name for the column. Then, in the **Column Attributes** area, provide details for the column.

Use the Delete button to drop a selected column.

#### TABLES (DB2 LUW) - PARTITION

Under **Partition Columns**, for each partition column, click the New button and then choose a column name from the dropdown. To add a **Data Partition**, click the New button to open a dialog that lets you add a partition.

#### TABLES (DB2 LUW) - TABLESPACES

For each Data Tablespace or Long Tablespace, click the New button and then choose a tablespace from the dropdown. To specify an Index Tablespace, select a tablespace from the dropdown.

#### TABLES (DB2 LUW) - DIMENSIONS

For each column that is to make up a dimension, click the New button to open a dialog that lets you add a column.

#### TABLES (DB2 LUW) - DISTRIBUTION KEY COLUMNS

For each column that is to make up the distribution key, click the New button and then select a column from the dropdown

#### TABLES (DB2 LUW) - INDEXES

Click **Add** to open the Index wizard. For more information, see [Index Wizard \(DB2 LUW\)](#).



## TABLES (DB2 LUW) - CONSTRAINTS

Selecting a constraint type and clicking **Add** opens the object wizard for that object type. For details see:

- [Primary Key Wizard \(DB2 LUW\)](#)
- [Unique Key Wizard \(DB2 Z/OS\)](#)
- [Foreign Key Wizard \(DB2 Z/OS\)](#)

## TABLESPACE WIZARD (DB2 LUW)

Tablespaces establish connections between the physical storage devices of your database system and the logical containers or tables being use to store data. In essence, a tablespace is a storage structure that can hold tables, indexes, large objects, and long data. The Tablespace Wizard lets you create a tablespace without knowing the underlying commands.

**To create a new tablespace using a wizard:**

- 1 [Open an Object Wizard](#) for a tablespace.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Tablespaces \(DB2 LUW\) - Properties](#).
  - **Container** panel - for details, see [Tablespaces \(DB2 LUW\) - Container](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **Permissions** panel - [Set Permissions for an Object](#)
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## TABLESPACES (DB2 LUW) - PROPERTIES

When creating or editing a tablespace, this tab/panel lets you work with the following settings:

Setting	Description
<b>Tablespace properties</b>	<p><b>Type</b> - Select REGULAR, LARGE, TEMPORARY, or USER TEMPORARY.</p> <p><b>Use Automatic Storage</b> and <b>Managed By</b> let you specify whether storage is managed automatically, by the database, or by the system.</p> <p><b>Database Partition Group</b> - lets you select a database partition group.</p> <p><b>Buffer Pool</b> - lets you select a buffer pool.</p> <p><b>Drop Recovery</b> - For REGULAR type tablespaces, lets you enable/disable drop recovery.</p>
<b>Performance properties</b>	This group lets you specify or select the <b>Page Size</b> , <b>Extent Size</b> , <b>Prefetch Automatic</b> , <b>Prefetch Size</b> , <b>Overhead</b> , <b>Transfer Rate</b> , and <b>File System Caching</b> properties.

Setting	Description
<b>Automatic Storage properties</b>	This group lets you specify or select the <b>AutoResize</b> , <b>Initial Size</b> , <b>Increase Size</b> , <b>Max Size Unlimited</b> , and <b>Max Size</b> attributes.

### TABLESPACES (DB2 LUW) - CONTAINER

For each container in the tablespace, in the **Container Properties** area, provide the following container properties: **Database Partitions**, **Type** (FILE or DEVICE), **Name**, and **Size**, and then click the **New** button.

Use the **Delete** button to drop a selected container.

### TRIGGER WIZARD (DB2 LUW)

A trigger defines a set of actions that take place in conjunction with, or are triggered, for a particular base table, with an insert, update, or delete statement. Triggers are handy ways to validate input data, read from other tables for cross-referencing purposes, and other similar purposes. The Trigger Wizard lets you create a trigger without requiring you to know any of the underlying commands.

#### To create a new trigger using a wizard:

- 1 [Open an Object Wizard](#) for a trigger.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Triggers \(DB2 LUW\) - Properties](#).
  - **Column Selection** panel - for details, see [Triggers \(DB2 LUW\) - Column Selection](#).
  - **Definition** panel - for details, see [Triggers \(DB2 LUW\) - Definition](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

### TRIGGERS (DB2 LUW) - PROPERTIES

When creating or editing a trigger, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Schema and Table Name</b>	Choose the owner and name of the table for which the trigger is being created.
<b>Schema and Name</b>	Choose the owner and name of the trigger being created.

Setting	Description
<b>Trigger Timing</b>	<p><b>BEFORE:</b> These triggers serve as extensions to the constraint subsystem and are most often used to validate input data, generate values for newly inserted rows, and read from other tables for cross-reference purposes. Note: Before triggers must be created as a For Each Row.</p> <p><b>AFTER:</b> Such a trigger is run after the integrity constraint validations; they can be used to modify operations in the database or be used for activities beyond the database, like supporting an alert notification.</p>
<b>Trigger Events</b>	<p>An INSERT trigger must be associated with an INSERT statement. For example, if a data load operation doesn't include an INSERT statement, the trigger will not be invoked.</p> <p>An UPDATE trigger can be associated with specific columns of the base table and will only be activated if those columns are updated.</p>
<b>Trigger Type</b>	<p><b>STATEMENT:</b> (only fires once).</p> <p><b>ROW</b> (fires for each affected row): The trigger runs as many times as there are rows in the affected section. If the set of affected rows is empty, the trigger doesn't run.</p>
<b>Old Table Alias</b>	Type the name of a temporary table of rows as they exist before they're updated or deleted.
<b>New Table Alias</b>	Type a name for a temporary table of rows as they exist after they're inserted or updated.
<b>Old Row Alias</b>	Type a name for the rows as they are before they've been deleted or updated.
<b>New Row Alias</b>	Type a name for the rows as they are after they've been inserted or updated.

#### TRIGGERS (DB2 LUW) - COLUMN SELECTION

If you chose UPDATE as the **Trigger Event**, select the columns, select the check box beside each column that is to fire the trigger.

#### TRIGGERS (DB2 LUW) - DEFINITION

Complete the CREATE TRIGGER outline provided by typing or pasting the body of the trigger. Finally, use the **Execute** button to [Create the New Object](#).

#### UNIQUE KEY WIZARD (DB2 LUW)

A unique key constraint is a key for which no two of its values can be equal and no values can be null. A table can have a number of unique constraints, but it cannot have more than one unique constraint on the same set of columns. If you are creating a unique key constraint on a table that already exists (as opposed to creating a unique key at the time the table is first generated), a unique index must already exist on the columns of the unique key you want to constrain. If no unique index exists, the Index Wizard will open as you complete the Create Unique Key Constraint dialog box.

**To create a new unique key using a wizard:**

- 1 [Open an Object Wizard](#) for a unique key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Unique Keys \(DB2 LUW\) - Properties](#).
  - **Columns** panel - for details, see [Unique Keys \(DB2 LUW\) - Columns](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## UNIQUE KEYS (DB2 LUW) - PROPERTIES

When creating or editing a unique key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Schema</b> and <b>Table Name</b>	Choose the owner and name of the table in which the unique key is being created.
<b>Name</b>	Provide a name for the unique key being created.

## UNIQUE KEYS (DB2 LUW) - COLUMNS

From the **Column** dropdown, select a column for the primary key and specify a **Sort** option. To add more columns, click the New button and then follow the steps in the last instruction. Use the Delete button to drop columns.

## USER DATATYPE (DB2 LUW)

User defined datatypes allow column definitions to be consistent across a database. They let you associate frequently used datatype information to be associated with a specified function and take advantage of performance features available to built-in datatypes including indexing and parallel database queries.

**To create a new user datatype using a wizard:**

- 1 [Open an Object Wizard](#) for a user datatype.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [User Datatype \(DB2 LUW\) - Properties](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).

- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

### USER DATATYPE (DB2 LUW) - PROPERTIES

When creating or editing a user datatype, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the user datatype.
<b>Datatype</b>	Provide a name for the datatype.
<b>Type</b>	Select the base datatype.
<b>Size</b>	Provide the size of the datatype.
<b>Allow Bit Data</b>	The option is only available for certain datatypes. A check means you want to store the data in a bit format.

### USER WIZARD (DB2 LUW)

Users have authorization to use a database and its objects, and the User Wizard gives you an easy way to add new ones.

#### To create a new user using a wizard:

- 1 [Open an Object Wizard](#) for a user
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Users \(DB2 LUW\) - Properties](#).
  - **Object Permissions** and **System Permissions** panels - [Set Permissions for an Object](#)
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

### USERS (DB2 LUW) - PROPERTIES

When creating or editing a user datatype, this tab/panel lets you provide the user **Name**.

### VIEW WIZARD (DB2 LUW)

A view gives you an alternative way to look at data in one or more tables. You can customize how you see a table's contents.

**To create a new view using a wizard:**

- 1 [Open an Object Wizard](#) for a view.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Views \(DB2 LUW\) - Properties](#).
  - **Definition** panel - for details, see [Views \(DB2 LUW\) - Definition](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**VIEWS (DB2 LUW) - PROPERTIES**

When creating or editing a view, this tab/panel lets you work with the following settings:

Setting	Description
<b>Schema</b>	Select the owner of the view. The owner of the view must have SELECT privileges for the tables in the CREATE view statement or DBADM authority on the database that contains the table.
<b>Name</b>	Provide a name for the view.
<b>Check Type</b>	<p><b>CHECK_NONE</b> - No search conditions must be satisfied for insert or update operations.</p> <p><b>CHECK_LOCAL</b> - Update and insert operations on view must satisfy the search conditions of the view and underlying views that are defined with a check option. Furthermore, every updatable view that is directly or indirectly defined on view inherits those search conditions (the search conditions of view and all underlying views of that are defined with a check option) as a constraint on insert or update operations.</p> <p><b>CHECK_CASCADE</b> - Update and insert operations on the view must satisfy the search conditions of view and all underlying views, regardless of whether the underlying views were defined with a check option. Furthermore, every updatable view that is directly or indirectly defined on view inherits those search conditions (the search conditions of view and all underlying views) as a constraint on insert or update operations.</p>

**Views (DB2 LUW) - Definition**

Complete the CREATE VIEW statement by typing or pasting in the relevant query. Finally, use the **Execute button** to [Create the New Object](#).

**IBM DB2 FOR Z/OS OBJECT WIZARDS**

Rapid SQL lets you create DB2 OS390 objects using the following wizards:

- [Alias Wizard \(DB2 Z/OS\)](#)
- [Database Wizard \(DB2 Z/OS\)](#)
- [Foreign Key Wizard \(DB2 Z/OS\)](#)
- [Function Wizard \(DB2 Z/OS\)](#)

- [Index Wizard \(DB2 Z/OS\)](#)
- [Plan Wizard \(DB2 Z/OS\)](#)
- [Primary Key Wizard \(DB2 Z/OS\)](#)
- [Procedure Wizard \(DB2 Z/OS\)](#)
- [Stogroup Wizard \(DB2 Z/OS\)](#)
- [Synonym Wizard \(DB2 Z/OS\)](#)
- [Table Wizard \(DB2 Z/OS\)](#)
- [Tablespace Wizard \(DB2 Z/OS\)](#)
- [Trigger Wizard \(DB2 Z/OS\)](#)
- [Unique Key Wizard \(DB2 Z/OS\)](#)
- [User Datatype Wizard \(DB2 Z/OS\)](#)
- [User Wizard \(DB2 Z/OS\)](#)
- [View Wizard \(DB2 Z/OS\)](#)

In addition, see [Create Synonym](#).

## ALIAS WIZARD (DB2 Z/OS)

An alias offers you security and convenience so that you can refer to an object without revealing who owns it or what database it belongs to. You can create aliases for tables or views. The Alias Wizard lets you create an alias without knowing the underlying commands. As you complete the Alias Wizard process, Rapid SQL constructs the necessary CREATE ALIAS statement based on the information that you supply. To create an alias, you must have CREATE ALIAS privileges or sysadmin or sysctrl authority.

### To create a new alias using a wizard:

- 1 [Open an Object Wizard](#) for an alias.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Aliases \(DB2 z/OS\) - Properties](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## ALIASES (DB2 Z/OS) - PROPERTIES

When creating or editing a view, this tab/panel lets you work with the following settings:

Setting	Description
<b>Schema</b>	Select the schema that is to own the alias.
<b>Name</b>	Provide a name for the alias.
<b>Target Owner</b>	Select the owner of the object to which you are creating an alias.
<b>Target Type</b>	Select the type of object (TABLE, VIEW) to which you are creating an alias.
<b>Target Name</b>	Select the specific object to which you are creating an alias.

## DATABASE WIZARD (DB2 Z/OS)

The Database Wizard lets you create a database (a structured collection of data that can be updated or queried) without knowing the underlying commands. Databases can be simple, that is one file with many records and the same fields, or much more complicated with multiple files with different fields.

**To create a new database using a wizard:**

- 1 [Open an Object Wizard](#) for a database.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Databases \(DB2 z/OS\) - Properties](#).
  - **Permissions** panel - for details, see [Set Permissions for an Object](#)
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## DATABASES (DB2 Z/OS) - PROPERTIES

When creating or editing a database, this tab/panel lets you work with the following settings:

Setting	Description
<b>Name</b>	Provide a unique name for the database.
<b>Type</b>	Workfile: This option is only available if the server is configured in IBM DB2 for OS/390 and z/OS to allow sharing. For more information, contact your System administrator.  Temp: This option indicates the database is only for declared temporary tables. A temp database cannot be shared.
<b>Group Member</b>	Specifies the member for which this database is being created. Use this only in a shared environment.
<b>Tablespace Buffer Pool</b>	Select the default buffer pool to be used for tablespaces created within the database.



Setting	Description
<b>Index Buffer Pool</b>	Select the default buffer pool name to be used for indexes created within the database.
<b>Storage Group</b>	Select the default storage group to support the DASD space requirements for tablespaces and indexes within the database.
<b>Encoding Scheme</b>	Select an encoding schema of DEFAULT, ASCII, EBCDIC, or UNICODE.  NOTE: To change the encoding scheme for a database after it has been created to use a different coded character set identifier (CCSID) that supports the Euro symbol, all data must be unloaded and reloaded. For more information regarding the encoding scheme, contact your System administrator.

## FOREIGN KEY WIZARD (DB2 Z/OS)

Foreign keys are unique values that refer to specific columns of other tables. Thus, a foreign key links two tables together. Embarcadero Rapid SQL's Foreign Key Wizard makes it easy for you to create a relational link between two tables, thereby speeding queries and giving you faster access to data. The Foreign Key Wizard lets you create a foreign key without knowing the underlying commands.

### To create a new foreign key using a wizard:

- 1 [Open an Object Wizard](#) for a foreign key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Foreign Keys \(DB2 z/OS\) - Properties](#).
  - **Column Mapping** panel - for details, see [Foreign Keys \(DB2 z/OS\) - Column Mapping](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## FOREIGN KEYS (DB2 Z/OS) - PROPERTIES

When creating or editing a foreign key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Schema</b>	Select the owner of the referencing, or child, table.
<b>Table Name</b>	Select the name of the referencing, or child, table.
<b>Name</b>	If you do not want to use the system-generated name, provide a new one.
<b>Delete Rule</b>	Select the action to be taken (NO ACTION, RESTRICT, CASCADE, or SET NULL) when a row of the referenced, or parent, table is deleted.

## FOREIGN KEYS (DB2 z/OS) - COLUMN MAPPING

Under **Referenced Table**, choose the **Owner** and then the **Name** of the referenced, or parent, table.

Under the **Main Table**, select checkboxes corresponding to the columns that are to reference columns in the referenced table. Then, under **Referenced Table**, select the corresponding column check boxes.

## FUNCTION WIZARD (DB2 Z/OS)

To create a relationship between one set of values and another, Rapid SQL offers the Function Wizard. You can develop reusable subroutines so you can control, access, and manipulate the data that underlies an object. As you complete the Function Wizard process, Rapid SQL constructs the necessary CREATE FUNCTION statement based on the information that you supply.

**NOTE:** To create a user-defined function, you need CREATE ANY privileges or IMPLICIT\_SCHEMA authority on the database if the schema does not already exist.

**To create a new function using a wizard:**

- 1 [Open an Object Wizard](#) for a function.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Functions \(DB2 z/OS\) - Properties](#).
  - **Source** panel - (Only available for a Function Type of SOURCED.) for details, see [Functions \(DB2 z/OS\) - Source](#).
  - **Parameters** panel - for details, see [Functions \(DB2 z/OS\) - Parameters](#).
  - **Return Scalar** panel - for details, see [Functions \(DB2 z/OS\) - Return Scalar](#).
  - **Return Columns** panel - (Only available when you choose a Function Type of EXTERNAL TABLE.) for details, see [Functions \(DB2 z/OS\) - Return Columns](#).
  - **Body** panel - for details, see [Functions \(DB2 z/OS\) - Body](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## FUNCTIONS (DB2 Z/OS) - PROPERTIES

When creating or editing a function, this tab/panel lets you work with the following settings:

Setting	Description
<b>Schema, Name, and Specific Name</b>	Select the owner of the function, provide a name for the function, and provide the Specific name to be used by some SQL statements and DB2 commands for this function.

Setting	Description
<b>Function Type</b>	<p>External Scalar: This allows you to extend the function by adding your own or another party's definition for the function.</p> <p>External Table: Use this to create a function that is written in ASSEMBLE, C, COBOL, or PLI to return a table after it is deployed.</p> <p>Sourced: Here you are creating a function that is based on an existing scalar or table function with an application server.</p> <p>SQL: This type of function returns a single value when the function is invoked if the SQL statement that defines it is valid.</p>
<b>Language</b>	If you chose a Function Type of EXTERNAL SCALAR or EXTERNAL TABLE, specify a language of ASSEMBLE, C, COBOL, or PLI.
<b>Return Type</b>	Identifies the return type of the function.
<b>External Name</b>	Provide the External Name of the function.
<b>SQL Access Level</b>	<p>Indicates whether the function can execute SQL statements.</p> <p>CONTAINS SQL: Statements that don't read or modify SQL can be executed.</p> <p>NO SQL: No SQL statements can be executed.</p> <p>READS SQL: Statements that cannot modify SQL can be executed.</p>
<b>WLM Environment</b>	Specify a Workload Management Environment (Required if Language is JAVA/COMPJAVA/REXX, the Procedure contains a LOB parameter, Security is set to 'USER' or 'DEFINER', or program type is 'SUB').
<b>WLM For Nested</b>	Self-explanatory

## FUNCTIONS (DB2 Z/OS) - SOURCE

**NOTE:** Only available for a **Function Type** of SOURCED.

Select the **Schema**, **Name**, and **Specific Name** of the source function.

## FUNCTIONS (DB2 Z/OS) - PARAMETERS

For each parameter for this function, use the New button to add a new parameter, provide a name for the parameter, and in the **Attributes** area, select a **Type**, and if appropriate, the **Precision**, **Scale**, **Size**, and **As Locator** options.

## FUNCTIONS (DB2 Z/OS) - RETURN SCALAR

Under **Return Datatype**, select a Type and depending on your choice, provide or select **Precision**, **Scale**, **Size**, and **As Locator** options.

To make use of a CAST FROM clause, under **Cast Datatype** set **Enabled** to True, select a **Type**, and if appropriate, the **Scale**, **Size**, and **As Locator** options

## FUNCTIONS (DB2 Z/OS) - RETURN COLUMNS

**NOTE:** Only available when you choose a **Function Type** of EXTERNAL TABLE.

For each column returned by this function, use the **New** button to add a new parameter, provide a name for the parameter, and in the **Attributes** area, select a **Type**, and if appropriate, the **Precision**, **Scale**, **Size**, and **As Locator** options.

### Functions (DB2 z/OS) - Body

Enter the return statement for the function.

## INDEX WIZARD (DB2 Z/OS)

Like the index in a book, a database index makes it easier for you to access information quickly. An index lists the location of rows, sorted by the contents of one or more columns.

### To create a new index using a wizard:

- 1 [Open an Object Wizard](#) for an index.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Indexes \(DB2 z/OS\) - Properties](#)
  - **Columns** panel - for details, see [Indexes \(DB2 z/OS\) - Columns](#)
  - **Indexes** panel - for details, see [Indexes \(DB2 z/OS\) - Storage](#)
  - **Partitions** panel - for details, see [Indexes \(DB2 z/OS\) - Partitions](#)
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## INDEXES (DB2 Z/OS) - PROPERTIES

When creating or editing an index, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Schema and Table Name</b>	Choose the owner and name of the table in which the index is being created.
<b>Schema and Name</b>	Choose the owner and name of the index being created.
<b>Index Type</b>	<p>Unique: Prevents the selected table from having two or more rows with the same value of the index key. The uniqueness is enforced at the end of the SQL statement update. Also, null values are treated like any other values, so a column cannot contain more than one null. If you later elect to partition the index, the columns specified for the unique key must include all columns for the partitioning key.</p> <p>Non-Unique (default)</p>

Setting	Description
<b>Clustered</b>	Enable or disable clustering. Unless you specifically select the CLUSTER option when you create an index, the first index you create on a table will be bestowed with that distinction. Each table can only have one clustering index at a time.  The clustering index tells DB2 to insert rows in the table in the order of the clustering key values. Inserted rows will be stored contiguously in sequence when possible. Columns frequently searched with a range of values BETWEEN, less than, greater than, and LIKE, or GROUP BY, etc., are good candidates for a clustering index.
<b>Number of Partitions</b>	If you enabled Clustering, specify the number of partitions.
<b>Buffer Pool</b>	Provide the buffer pool in which this index should reside
<b>Defer, Close, Copy, and Define</b>	Enable or disable these DB2 options as required.
<b>Piece Size</b>	The maximum addressability of the dataset for a non-partitioned index.
<b>Compress</b>	Enabling this setting specifies a COMPRESS=YES parameter when creating the table, resulting in a compressed index. This feature is available if the size of the specified <b>Buffer Pool</b> is 8k, 16k, or 32k.
<b>Index on expression</b>	Setting the check box enables the Index on Expression feature. The key expression is then provided on the <b>Columns</b> tab/panel. For details, see <a href="#">Indexes (DB2 z/OS) - Columns</a> .
<b>Padded</b>	Setting this check box specifies that the table is to be created with the PADDED keyword. This feature must be enabled if you are using index key randomization. The RANDOM setting is applied on a column by column basis on the <b>Columns</b> tab/panel.

## INDEXES (DB2 z/OS) - COLUMNS

The actions you take on this tab/panel depend on whether you set the Index on Expression check box on the Properties panel. For more information, see [Indexes \(DB2 z/OS\) - Properties](#).

- If Index on Expression is disabled:

From the **Column** dropdown, select a column for the index and specify a **Sort** option.

**NOTE:** If you want to use the RANDOM sort option, you must first set the **Padded** check box on the **Properties** tab/panel.

- If Index on Expression is enabled:

Enter a valid key-expression in the **Expression** field and choose a **Sort** order.

To add more columns/key-expressions, click the New button and then follow the steps in the last instruction. Use the Delete button to drop columns.

## INDEXES (DB2 z/OS) - STORAGE

This tab/panel lets you perform the following tasks:

- Select a dataset management scheme
- Provide associated attribute values

**To select a data set management scheme:**

- 1 Click the **Edit** button. The **Data Set Management** dialog opens.
- 2 Set one of the following data set management options:
  - **DB2 will define and manage the data sets on a volume of the default storage group of the database**
  - **DB2 will define and manage the data sets on a volume of the specified storage group** - Select a storage group (a storage group is made up of disk device volumes): Each data set will be defined on a volume listed in the storage group you select.

Minimum primary space allocation: 12 kilobytes is the default.

Minimum secondary space allocation: 12 kilobytes is the default.

**NOTE:** NOTE: If the primary and (118 x) secondary space allocations combine to be greater than 2 gigabytes, more than one data set may eventually be used.

Erase data sets when index dropped? If choose this option, DB2 will overwrite all data with zeros before they are deleted as a security measure.

- **User will manage the data sets on a specified VCAT catalog-name** - Enter or select the VCAT. Do not select this option for an index on a declared temporary table.

- 3 Click **OK**.

In addition to the attributes specific to your data set management choice, this tab/panel also offers the following settings:

<b>Free Page</b>	One free page exists for every x pages. The x specifies how often to leave a page of free space when index entries are created from executing a DB2 utility or creating an index for a table with pre-existing rows. (0-255)
<b>Percent Free</b>	The percentage of free space you want to leave in every page when entries are added to an existing index. The default is 10%.
<b>GBP Cache</b>	This option is available only in a data-sharing environment. ALL: As pages are read, all of them will be cached in the group buffer pool. CHANGED: Updated pages are cached to the group buffer pool. NONE: No pages will be cached.

**INDEXES (DB2 z/OS) - PARTITIONS**

Displays the default settings for the number of partitions you specified on the **Properties** pane. Select a partition and click the **Edit** button to modify details for that partition.

## PLAN WIZARD (DB2 Z/OS)

A plan, also known as the application plan, is a control structure that is used to process the SQL statements DB2 encounters when it is executing those SQL statements. The Plan Wizard, really, the Bind Plan Wizard, creates the structure that is used in the bind process--the process by which the output from the SQL precompiler is converted into usable form. Some authorization checking is necessary.

### To Open the Bind Plan Wizard

- 1 On the Explorer, find the database where you want to add the new bind plan.
- 2 On the Plan branch, right-click and select Create.

The Bind Plan Wizard lets you set plan parameters, add packages, and set bind properties. The table below describes the options and functionality on the Bind Plan wizard:

Panel	Settings and tasks	
1	Plan Name	Lets you select the plan name.
	Qualifier	OPTIONAL: Lets you select a qualifier, the plan creator.
	Action	OPTIONAL: Lets you select an action.
	Sql Rules	OPTIONAL: Determines whether you can execute a type 2 CONNECT statement to an existing SQL connection, according to DB2 rules. Lets you select DB2 or STD.
	Cache Size	OPTIONAL: Lets you select or type the cachesize in bytes, the authorization cache acquired in the EDM pool for the plan. At run time, the authorization cache stores user IDs authorized to run. Consulting the cache can avoid a catalog lookup for checking authorization to run the plan.
	Plan Owner	OPTIONAL: Determines the authorization ID of the owner of the plan.
	Current Server	OPTIONAL: Determines the location to connect to before running the plan.
	Resource Acquire	OPTIONAL: Use - Acquires table space locks only when first used by a bound application program. Allocate - Acquires all table space locks when the plan is allocated. The value has no effect on dynamic SQL statements, which always use ACQUIRE(USE).
	Disconnect	OPTIONAL: Determines which remote connections to destroy during commit operations. The option applies to any application process that uses the plan and has remote connections of any type. Regardless of the value of this option, a commit operation destroys all connections in the release pending state.  Explicit - Destroy only connections in the release pending state. This value allows you maximum flexibility for controlling remote connections.  Automatic - Destroy all remote connections.  Conditional - Destroy all remote connections unless an open cursor defined as WITH HOLD is associated with the connection.
2	Lets you select the <b>Member Name</b> , <b>PDS Name</b> (partitioned data set) and click <b>Add</b> to enter each member and PDS name.	
3	Lets you select the <b>Location</b> to connect to, the <b>Collection</b> (location of the DBMS where the plan binds and where the description of the plan resides.) and a <b>Package</b> .	

Panel	Settings and tasks	
4	Isolation	Determines how far to isolate an application from the effects of other running applications.
	Keep Dynamic	Specifies that DB2 keeps dynamic SQL statements after commit points. The application does not need to prepare an SQL statement after every commit point. DB2 keeps the dynamic SQL statement until the application process ends, a rollback operation occurs or the application executes an explicit PREPARE statement with the same statement identifier.  If the prepared statement cache is active, DB2 keeps a copy of the prepared statement in the cache. If the prepared statement cache is not active, DB2 keeps only the SQL statement string past a commit point. DB2 then implicitly prepares the SQL statement if the application executes an OPEN, EXECUTE, or DESCRIBE operation for that statement.
	Current Data	Determines whether to require data currency for read-only and ambiguous cursors when the isolation level of cursor stability is in effect. It also determines whether block fetching can be used for distributed, ambiguous cursors.
	Degree	Determines whether to attempt to run a query using parallel processing to maximize performance. Lets you select an option.
	Dynamic Rules	Determines what values apply at run time for the following dynamic SQL attributes: The authorization ID that is used to check authorization The qualifier that is used for unqualified objects The source for application programming options that DB2 uses to parse and semantically verify dynamic SQL statements Whether dynamic SQL statements can include GRANT, REVOKE, ALTER, CREATE, DROP, and RENAME statements
	Release	Determines when to release resources that a program uses. Options are at each <b>Commit</b> point or <b>Deallocate</b> when the program terminates.
	Validate	Determines whether to recheck, at run time, errors found during bind. The option has no effect if all objects and needed privileges exist.  Bind - If not all objects or needed privileges exist at bind time, the wizard displays an error messages, and does not bind the package.  Run - If not all objects or privileges exist at bind time, the process issues warning messages, but the bind succeeds. DB2 checks existence and authorization again at run time for SQL statements that failed those checks during bind. The checks use the authorization ID of the plan owner.
5	Explain	Obtains information about how SQL statements in the member list of the plan, are to execute, and then inserts that information into the table owner.PLAN_TABLE, where owner is the authorization ID of the owner of the plan or package. This option does not obtain information for statements that access remote objects.
	Reopt(VARS)	Re-determines the access path at run time.
	Prepare	Prepares dynamic SQL statements that refer to remote objects.
	ImmedWrite	Immediate writes will be done for updates made to group buffer pool dependent pagesets or partitions.
	Opthint	Query optimization hints are used for static SQL.
	Encoding	Lets you select type of language for the package.
	Path	Lets you select a path that DB2 uses to resolve unqualified user-defined distinct types, functions, and stored procedure names (in CALL statements).
	Flag	Lets you select all message types or a specified subset to display: informational, warning, error, and completion messages.



Panel	Settings and tasks
6	<b>Enable</b> or <b>Disable</b> system connection types that can use the plan or package, select a <b>System</b> or <b>Cname</b> option.

## PRIMARY KEY WIZARD (DB2 Z/OS)

A primary key is a unique key that is part of a table's definition. There can be only one primary key for each table, and the columns contained in a primary key cannot have null values. A primary key constraint forbids duplicate values in one or more columns. A table with a primary key will be considered the parent table to a table with a foreign key, which becomes the dependent table.

**NOTE:** A nullable column cannot be a part of a primary key.

### To create a new primary key using a wizard:

- 1 [Open an Object Wizard](#) for a primary key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Primary Keys \(DB2 z/OS\) - Properties](#).
  - **Columns** panel - for details, see [Primary Keys \(DB2 z/OS\) - Columns](#).
  - **Storage** panel - for details, see [Primary Keys - Storage - Edit button \(manage datasets\)](#) and [Primary Keys \(DB2 z/OS\) - Storage - Attributes](#).
  - **Partitions** panel - for details, see [Primary Keys \(DB2 z/OS\) - Partitions](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## PRIMARY KEYS (DB2 Z/OS) - PROPERTIES

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Schema and Table Name</b>	Choose the owner and name of the table in which the index is being created.
<b>Schema and Name</b>	Choose the owner and name of the index being created.
<b>Clustered</b>	Enable or disable clustering.
<b>Number of Partitions</b>	If you enabled Clustering, specify the number of partitions.
<b>Buffer Pool</b>	Provide the buffer pool in which this index should reside

Setting	Description
<b>Defer, Close, Copy, and Define</b>	Enable or disable these DB2 options as required.
<b>Piece Size</b>	The maximum addressability of the dataset for a non-partitioned index.

## PRIMARY KEYS (DB2 Z/OS) - COLUMNS

From the **Column** dropdown, select a column for the primary key and specify a **Sort** option. To add more columns, click the New button and then follow the steps in the last instruction. Use the Delete button to drop columns.

## PRIMARY KEYS - STORAGE - EDIT BUTTON (MANAGE DATASETS)

**NOTE:** Availability differs according to the dataset management options you chose

Choose a data set management option:

**DB2 will define and manage the data sets on a volume of the default storage group of the database**

**DB2 will define and manage the data sets on a volume of the specified storage group**

Select a storage group (a storage group is made up of disk device volumes): Each data set will be defined on a volume listed in the storage group you select.

**Minimum primary space allocation:** 12 kilobytes is the default.

**Minimum secondary space allocation:** 12 kilobytes is the default.

**NOTE:** If the primary and (118 x) secondary space allocations combine to be greater than 2 gigabytes, more than one data set may eventually be used.

**Erase data sets when index dropped?** If you choose this option, DB2 will overwrite all data with zeros before they are deleted as a security measure.

**User will manage the data sets on a specified VCAT catalog-name** Enter or select the VCAT. Do not select this option for an index on a declared temporary table.

## PRIMARY KEYS (DB2 Z/OS) - STORAGE - ATTRIBUTES

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Storage Group, Primary Space Allocation, Secondary Space Allocation, Erase, and VCAT catalog</b>	The ability to set these options depends on the dataset management options you chose.

Setting	Description
<b>Free Page</b>	One free page exists for every x pages. The x specifies how often to leave a page of free space when index entries are created from executing a DB2 utility or creating an index for a table with pre-existing rows. (0-255)
<b>Percent Free</b>	The percentage of free space you want to leave in every page when entries are added to an existing index. The default is 10%.
<b>GBP Cache</b>	This option is available only in a data-sharing environment. ALL: As pages are read, all of them will be cached in the group buffer pool. CHANGED: Updated pages are cached to the group buffer pool. NONE: No pages will be cached.

### PRIMARY KEYS (DB2 Z/OS) - PARTITIONS

Displays the default settings for the number of partitions you specified on the **Properties** pane. Select a partition and click the **Edit** button to modify details for that partition.

### PROCEDURE WIZARD (DB2 Z/OS)

Procedures are a reusable block of PL/SQL, stored in the database, that applications can call. Procedures streamline code development, debugging, and maintenance by being reusable. Procedures enhance database security by letting you write procedures granting users execution privileges to tables rather than letting them access tables directly.

#### To create a new procedure using a wizard:

- 1 [Open an Object Wizard](#) for a procedure.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Procedures \(DB2 z/OS\) - Properties](#).
  - **Parameters** panel - for details, see [Procedures \(DB2 z/OS\) - Parameters](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

#### Important Notes

- If you are creating a SQL routine procedure, you must have the latest DB2 fixpack installed on your OS/390 Client. If you do not have the latest fixpack installed, the SQL routine procedure creation will fail.

## PROCEDURES (DB2 Z/OS) - PROPERTIES

When creating or editing a procedure, this tab/panel lets you work with the following settings:

Setting	Description
<b>Schema, Name, and Specific Name</b>	Select the schema that is to own the procedure, provide a name for the procedure, and provide the Specific Name for the procedure.
<b>Language</b>	<p>The database manager will call the procedure accordingly assuming the program is designed to run in the server's environment.</p> <p>Assemble: a stored procedure written in Assembler</p> <p>C: a stored procedure written in C or C++</p> <p>COBOL: a stored procedure written in COBOL</p> <p>CompJAVA: CompJAVA is no longer supported. Stored procedures should alternatively be written in JAVA.</p> <p>JAVA</p> <p>PLI: A stored procedure written in PL/I.</p> <p>REXX (Restructured Extended Executor Language) - Don't use this language when SQL parameter style is in effect. To specify REXX, the general parameter style or general with nulls.</p> <p>SQL</p>
<b>SQL Access Level</b>	<p>MODIFIES SQL DATA (Default): The procedure can support any SQL statement except those that cannot be supported in procedures.</p> <p>CONTAINS SQL DATA: Only SQL statements that neither modify nor read SQL data can be executed in the procedure.</p> <p>READS SQL DATA Some SQL statements that don't modify SQL data can be included in the procedure</p> <p>NO SQL: Only SQL statements with a data access classification of NO SQL can be executed. Don't select this option for a JAVA procedure that uses a .jar.</p>
<b>WLM Environment</b>	Specify a Workload Management Environment (Required if Language is JAVA/COMPJAVA/REXX, the Procedure contains a LOB parameter, or Security is set to 'USER' or 'DEFINER').
<b>WLM For Nested</b>	Self-explanatory

## PROCEDURES (DB2 Z/OS) - PARAMETERS

For each parameter for this function, use the New button to add a new parameter, provide a name for the parameter, and in the **Attributes** area, select a **Type**, specify a **Parameter Mode** of INPUT, OUTPUT, or INPUT\_OUTPUT, and if appropriate, the **Precision**, **Scale**, and **Size** options.

## STOGROUP WIZARD (DB2 Z/OS)

Stogroups are storage groups. You create them on the current server. Storage from the named sets of volumes you create can be allocated at a later date for tablespaces or index spaces.

**To Open the Stogroup Wizard**

- 1 On the Explorer, find the datasource where you want to add the new Storage Group.
- 2 Expand the Storage branch, right-click Stogroup, and select New.

The table below describes the fields you will encounter as you complete the Stogroup Wizard.

Required Information	Description
What is the name of the Stogroup?	Enter a name for the storage group.
VCAT	This is the integrated catalog facility catalog, or volume catalog. Name the catalog or choose one from the drop-down list if it is available.
Select the volumes in the Stogroup	Specify a set of volumes that may exist on the system but may not be in use by other storage groups.
Select All	Selects all listed volumes.
Unselect All	Unselects all listed volumes.
Add	Opens the <a href="#">Add Volume Dialog Box</a> .
Remove	Deletes all selected volumes from the list.

#### ADD VOLUME DIALOG BOX

The table below describes the options and functionality on the Add Volume dialog box:

Required Information	Description
Enter one or more volumes to add to the stogroup	Type the names of the volumes (separated by spaces) to add to the stogroup.
Check	Click to see if any additional information is available about the volumes you typed. Opens the <a href="#">Volumes Info Dialog Box</a> .
Or select volumes for the list	Lets you select volumes.

#### VOLUMES INFO DIALOG BOX

The table below describes the options and functionality on the Volume Info dialog box:

Required Information	Description
Volumes	Lets you select volumes.
OK	Click to add volumes to Stogroup Wizard.

## SYNONYM WIZARD (DB2 Z/OS)

A synonym is an alternative name you can create for a table or view. Using a synonym can make it easier for you to remember that table or view instead of having to think about the possibly cumbersome formal name (for example, a show dog may have a formal name that incorporates all his ancestors, but will answer to the name Spot around the house).

### To create a new synonym using a wizard:

- 1 [Open an Object Wizard](#) for a synonym.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Synonyms \(DB2 z/OS\) - Properties](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

### SYNONYMS (DB2 z/OS) - PROPERTIES

When creating or editing a synonym, this tab/panel lets you work with the following settings:

Setting	Description
<b>Schema</b>	Select the schema that is to own the synonym.
<b>Name</b>	Provide a name for the synonym.
<b>Referenced Object Owner</b>	Select the owner of the object to which you are creating a synonym.
<b>Referenced Object Type</b>	Select the type of object (TABLE, VIEW) to which you are creating a synonym.
<b>Referenced Object Name</b>	Select the specific object to which you are creating a synonym.

## TABLE WIZARD (DB2 Z/OS)

All data in a database is stored in a tabular format, that is a collection of rows and columns. Tables, therefore are fundamental to whatever database you are administering.

**To create a new table using a wizard:**

- 1 [Open an Object Wizard](#) for a table.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Tables \(DB2 z/OS\) - Properties](#).
  - **Columns** panel - for details, see [Tables \(DB2 z/OS\) - Columns](#).
  - **Indexes** panel - for details, see [Tables \(DB2 z/OS\) - Indexes](#).
  - **Constraints** panel - for details, see [Tables \(DB2 z/OS\) - Constraints](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **Permissions** panel - [Set Permissions for an Object](#)
  - **DDL View** panel - for details, see [Create the New Object](#).

The table that follows describes the fields you may encounter as you complete the Table Wizard.

**TABLES (DB2 z/OS) - PROPERTIES**

Select the **Schema**, provide a **Name**, and provide or select other table properties.

**TABLES (DB2 z/OS) - COLUMNS**

For each column in the table, click the **Add Column** button to create a column, provide a **Name** for the column and provide or select the remaining column attributes.

Use the **Delete** button to drop a selected column.

Use the arrow buttons to reposition the columns.

**Tables (DB2 z/OS) - Indexes**

Click Add to open the [Index Wizard \(DB2 Z/OS\)](#).

**TABLES (DB2 z/OS) - CONSTRAINTS**

Selecting a constraint type and clicking **Add** opens the object wizard for that object type. For details see:

- [Foreign Key Wizard \(DB2 Z/OS\)](#)
- [Primary Key Wizard \(DB2 Z/OS\)](#)
- [Unique Key Wizard \(DB2 Z/OS\)](#)
- [Create Synonym](#)

## TABSPACE WIZARD (DB2 Z/OS)

Tablespaces establish connections between the physical storage devices of your database system and the logical containers or tables being used to store data. In essence, a tablespace is a storage structure that can hold tables, indexes, large objects, and long data.

**NOTE:** To change the encoding scheme for a database after it is created to utilize a different coded character set identifier (CCSID) that supports the Euro symbol, all data must be unloaded and reloaded. For more information regarding the encoding scheme, contact your System administrator.

### To Open the Tablespace Wizard

- 1 On the Explorer, find the datasource where you want to add the new Tablespace.
- 2 Expand the Storage branch, right-click Tablespaces, and select New.
- 3 Use the following table as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Tablespaces \(DB2 z/OS\) - Properties](#).
  - **Permissions** panel - [Set Permissions for an Object](#)
  - **DDL View** panel - for details, see [Create the New Object](#).
- 4 Finally, use the **Execute** button to create the object.

### TABSPACES (DB2 z/OS) - PROPERTIES

When creating or editing a tablespace, this tab/panel lets you work with the following settings:

Group	Settings	Description
<b>Creation</b>	<b>Name and Database</b>	Provide a name and select an associated database.
	Database type	Lets you select a REGULAR or WORKFILE database type.
	Creation type	Lets you select among PARTITIONED, SEGMENTED, RANGE-PARTITIONED UNIVERSAL, PARTITION-BY-GROWTH UNIVERSAL, and LOB options.  For detailed information on these types and other required settings, see <a href="http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp">http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp</a> . Search on a term such as <b>universal tablespaces</b> .
<b>Tablespace management</b>	<b>Management type</b>	Lets you select how you want the tablespace managed: DEFAULT STORAGE GROUP, STORAGE GROUP, or VCAT CATALOG-NAME.
	<b>Storage group</b> , Minimum primary space allocation, Minimum secondary space allocation, and Erase rule	If you selected a <b>Management type</b> of STORAGE GROUP, select a storage group, specify minimum space allocations, and specify whether data sets are to be erased when the tablespace is dropped.
	VCatName	If you selected a <b>Management type</b> of VCAT catalog-name, select a catalog name.



Group	Settings	Description
<b>Bufferpool</b>	<b>Buffer Pool</b>	Select the buffer pool where the tablespace is to reside.
<b>Partitions and size</b>	<b>Number of partitions</b>	If you selected a <b>Creation Type</b> of PARTITIONED or RANGE-PARTITIONED UNIVERSAL, specify the number of partitions.
	Partition size (DSSIZE)	If you did not select a <b>Creation Type</b> of SEGMENTED, select a partition size.
	Segment size	If you selected a <b>Creation Type</b> of SEGMENTED, PARTITION-BY GROWTH, or RANGE-PARTITIONED UNIVERSAL, select the number of pages that are to be assigned to each segment of the tablespace.
	Max rows per page	Specify the number of rows (1-255) that can be placed on each data page.
<b>Space Management</b>	<b>Free space portion of each page (%)</b>	If you did not select a <b>Creation Type</b> of LOB, specify the percentage of each page to be left as free space when a tablespace is reloaded or reorganized.
	Free page frequency	If you did not select a <b>Creation Type</b> of LOB, specify how often to leave a free page when a tablespace is reloaded or reorganized.
<b>Other parameters</b>	<b>GBPCACHE</b>	Select a group buffer pool cache scheme of CHANGED, ALL, SYSTEM, or NONE.
	<b>Compress</b>	Enable this setting to allow compression.
	<b>Track modified pages</b>	Enable this setting to track modified pages in the space map pages.
	<b>Encoding scheme</b>	Select an encoding scheme of EBCDIC, UNICODE, ASCII, or NONE.
	<b>Log</b>	Enabling this check box specifies that changes to a LOB column are to be written to the log.
	<b>Define</b>	Enabling this check box specifies that tablespace data sets are defined when the tablespace is created. Otherwise, data sets are not created until data is inserted.
	<b>Member Cluster</b>	Enable this setting to manage space for inserts on a member-by-member basis.
	<b>Close rule</b>	Enable this setting to specify CLOSE YES, dictating the data set close priority.
	<b>Lock Size</b>	Specify a lock size of ANY, TABLESPACE, PAGE, or ROW.
	<b>Maximum locks</b>	Specify the number of locks that are allowed before escalating.

## TRIGGER WIZARD (DB2 Z/OS)

All data in a database is stored in a tabular format, that is a collection of rows and columns. Tables, therefore are fundamental to whatever database you are administering.

**To create a new trigger using a wizard:**

- 1 [Open an Object Wizard](#) for a trigger.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Triggers \(DB2 z/OS\) - Properties](#).
  - **Column Selection** panel - for details, see [Triggers \(DB2 z/OS\) - Column Selection](#).
  - **Definition** panel - for details, see [Triggers \(DB2 z/OS\) - Definition](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## TRIGGERS (DB2 Z/OS) - PROPERTIES

When creating or editing a trigger, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Schema and Table Name</b>	Choose the owner and name of the table for which the trigger is being created.
<b>Schema and Name</b>	Choose the owner and name of the trigger being created.
<b>Trigger Timing</b>	<p><b>BEFORE:</b> These triggers serve as extensions to the constraint subsystem and are most often used to validate input data, generate values for newly inserted rows, and read from other tables for cross-reference purposes. Note: Before triggers must be created as a For Each Row.</p> <p><b>AFTER:</b> Such a trigger is run after the integrity constraint validations; they can be used to modify operations in the database or be used for activities beyond the database, like supporting an alert notification.</p>
<b>Trigger Events</b>	<p>An INSERT trigger must be associated with an INSERT statement. For example, if a data load operation doesn't include an INSERT statement, the trigger will not be invoked.</p> <p>An UPDATE trigger can be associated with specific columns of the base table and will only be activated if those columns are updated.</p>
<b>Trigger Type</b>	<p><b>STATEMENT:</b> (only fires once).</p> <p><b>ROW</b> (fires for each affected row): The trigger runs as many times as there are rows in the affected section. If the set of affected rows is empty, the trigger doesn't run.</p>
<b>Old Table Alias</b>	Type the name of a temporary table of rows as they exist before they're updated or deleted.
<b>New Table Alias</b>	Type a name for a temporary table of rows as they exist after they're inserted or updated.
<b>Old Row Alias</b>	Type a name for the rows as they are before they've been deleted or updated.
<b>New Row Alias</b>	Type a name for the rows as they are after they've been inserted or updated.

## TRIGGERS (DB2 Z/OS) - COLUMN SELECTION

If you chose UPDATE as the **Trigger Event**, select the columns, select the check box beside each column that is to fire the trigger.

## TRIGGERS (DB2 z/OS) - DEFINITION

Complete the CREATE TRIGGER outline provided by typing or pasting the body of the trigger. Finally, use the **Execute** button to [Create the New Object](#).

## UNIQUE KEY WIZARD (DB2 Z/OS)

A unique key constraint is a key for which no two of its values can be equal and no values can be null. A table can have a number of unique constraints, but it cannot have more than one unique constraint on the same set of columns. If you are creating a unique key constraint on a table that already exists (as opposed to creating a unique key at the time the table is first generated), a unique index must already exist on the columns of the unique key you want to constrain. If no unique index exists, the Index Wizard will open as you complete the Create Unique Key Constraint dialog box.

### To create a new unique key using a wizard:

- 1 [Open an Object Wizard](#) for a unique key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Unique Keys \(DB2 z/OS\) - Properties](#).
  - **Columns** panel - for details, see [Unique Keys \(DB2 z/OS\) - Columns](#).
  - **Storage** panel - for details, see [Unique Keys \(DB2 z/OS\) - Storage - Edit button \(manage datasets\)](#) and [Unique Keys \(DB2 z/OS\) - Storage - Attributes](#).
  - **Partitions** panel - for details, see [Unique Keys \(DB2 z/OS\) - Partitions](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## UNIQUE KEYS (DB2 Z/OS) - PROPERTIES

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Schema and Table Name</b>	Choose the owner and name of the table in which the unique key is being created.
<b>Schema and Name</b>	Choose the owner and name of the index being created.
<b>Clustered</b>	Enable or disable clustering.
<b>Number of Partitions</b>	If you enabled Clustering, specify the number of partitions.
<b>Buffer Pool</b>	Provide the buffer pool in which this unique key should reside

Setting	Description
<b>Defer, Close, Copy, and Define</b>	Enable or disable these DB2 options as required.
<b>Piece Size</b>	The maximum addressability of the dataset for a non-partitioned index.

#### UNIQUE KEYS (DB2 Z/OS) - COLUMNS

From the **Column** dropdown, select a column for the unique key and specify a **Sort** option. To add more columns, click the New button and then follow the steps in the last instruction. Use the Delete button to drop columns.

#### UNIQUE KEYS (DB2 Z/OS) - STORAGE - **Edit** BUTTON (MANAGE DATASETS)

**NOTE:** Availability differs according to the dataset management options you chose

Choose a data set management option:

**DB2 will define and manage the data sets on a volume of the default storage group of the database**

**DB2 will define and manage the data sets on a volume of the specified storage group**

Select a storage group (a storage group is made up of disk device volumes): Each data set will be defined on a volume listed in the storage group you select.

**Minimum primary space allocation:** 12 kilobytes is the default.

**Minimum secondary space allocation:** 12 kilobytes is the default.

**NOTE:** If the primary and (118 x) secondary space allocations combine to be greater than 2 gigabytes, more than one data set may eventually be used.

**Erase data sets when index dropped?** If you choose this option, DB2 will overwrite all data with zeros before they are deleted as a security measure.

**User will manage the data sets on a specified VCAT catalog-name** Enter or select the VCAT. Do not select this option for an index on a declared temporary table.

#### UNIQUE KEYS (DB2 Z/OS) - STORAGE - ATTRIBUTES

**NOTE:** Availability differs according to the dataset management options you chose

When creating or editing a unique key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Storage Group, Primary Space Allocation, Secondary Space Allocation, Erase, and VCAT catalog</b>	The ability to set these options depends on the dataset management options you chose.
<b>Free Page</b>	One free page exists for every x pages. The x specifies how often to leave a page of free space when index entries are created from executing a DB2 utility or creating an index for a table with pre-existing rows. (0-255)
<b>Percent Free</b>	The percentage of free space you want to leave in every page when entries are added to an existing index. The default is 10%.
<b>GBP Cache</b>	This option is available only in a data-sharing environment. ALL: As pages are read, all of them will be cached in the group buffer pool. CHANGED: Updated pages are cached to the group buffer pool. NONE: No pages will be cached.

#### UNIQUE KEYS (DB2 Z/OS) - PARTITIONS

Displays the default settings for the number of partitions you specified on the **Properties** pane. Select a partition and click the **Edit** button to modify details for that partition.

#### USER DATATYPE WIZARD (DB2 Z/OS)

A datatype is a named set of valid values that can be manipulated by a set of operations. There are intrinsic datatypes, which are predefined and always available, and derived datatypes. A derived datatype is a user-defined datatype, which can include both intrinsic and previously derived datatypes. The User Datatype Wizard lets you create a derived datatype without knowing the underlying commands.

##### To create a new user datatype using a wizard:

- 1 [Open an Object Wizard](#) for a user datatype.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [User Datatypes \(DB2 z/OS\) - Properties](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## USER DATATYPES (DB2 Z/OS) - PROPERTIES

When creating or editing a user datatype, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the user datatype.
<b>Datatype</b>	Provide a name for the datatype.
<b>Type</b>	Select the base datatype.
<b>Size</b>	Provide the size of the datatype.
<b>For Data</b>	Select the MIXED, SBCS, or BIT option for the datatype.
<b>CCSID</b>	Select the NONE, ASCII, EBCDIC, or UNICODE option for the datatype.

## USER WIZARD (DB2 Z/OS)

Users have authorization to use a database and its objects, and the User Wizard gives you an easy way to add new ones.

**To create a new user using a wizard:**

- 1 [Open an Object Wizard](#) for a user
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Users \(DB2 z/OS\) - Properties](#).
  - **Object Permissions** and **System Permissions** panels - [Set Permissions for an Object](#)
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## USERS (DB2 Z/OS) - PROPERTIES

When creating or editing a user datatype, this tab/panel lets you provide the user **Name**.

## VIEW WIZARD (DB2 Z/OS)

A view gives you a new way of looking at data in a results table. Views behave like tables because you can query views and perform data manipulation operations on them. However, views do not actually store any data. Instead, they depend on data contained in their base tables. Columns added to the base table(s) after the view is created are not included in the result set. Views are thus handy tools for controlling access to a table. You can allow someone to see portions of data without allowing that user to see the table in its entirety. For example, you can create a view that will permit a user to see employee names in a table without allowing access to the Social Security numbers of that same table.

The wizard itself is a single panel. After you complete the wizard, the View Editor opens so you can complete the definition of the view, choose the columns to show in the view, the dependencies, and access privileges to the view.

### To create a new view using a wizard:

- 1 [Open an Object Wizard](#) for a view.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Views \(DB2 z/OS\) - Properties](#).
  - **Definition** panel - for details, see [Views \(DB2 z/OS\) - Definition](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## VIEWS (DB2 Z/OS) - PROPERTIES

When creating or editing a view, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the view. The owner of the view must have SELECT privileges for the tables in the CREATE view statement or DBADM authority on the database that contains the table.
<b>Name</b>	Provide a name for the view.
<b>Check Type</b>	<p><b>CHECK_NONE</b> - No search conditions must be satisfied for insert or update operations.</p> <p><b>CHECK_LOCAL</b> - Update and insert operations on view must satisfy the search conditions of the view and underlying views that are defined with a check option. Furthermore, every updatable view that is directly or indirectly defined on view inherits those search conditions (the search conditions of view and all underlying views of that are defined with a check option) as a constraint on insert or update operations.</p> <p><b>CHECK_CASCADE</b> - Update and insert operations on the view must satisfy the search conditions of view and all underlying views, regardless of whether the underlying views were defined with a check option. Furthermore, every updatable view that is directly or indirectly defined on view inherits those search conditions (the search conditions of view and all underlying views) as a constraint on insert or update operations.</p>

### Views (DB2 z/OS) - Definition

Complete the CREATE VIEW statement by typing or pasting in the relevant query. Finally, use the **Schedule** or **Execute** buttons to [Create the New Object](#).

## INTERBASE/FIREBIRD OBJECT WIZARDS

Rapid SQL provides wizards that let you create all supported InterBase/Firebird objects. For details, see the following topics:

ets you create InterBase/Firebird objects using the following wizards:

- [Blob Filters Wizard \(InterBase/Firebird\)](#)
- [Domains Wizard \(InterBase/Firebird\)](#)
- [Encryption Keys wizard \(InterBase/Firebird\)](#)
- [Exceptions Wizard \(InterBase/Firebird\)](#)
- [External Functions Wizard \(InterBase/Firebird\)](#)
- [Foreign Key Wizard \(IBM DB2 LUW\)](#)
- [Generators Wizard \(InterBase/Firebird\)](#)
- [Indexes Wizard \(InterBase/Firebird\)](#)
- [Primary Keys Wizard \(InterBase/Firebird\)](#)
- [Procedures Wizard \(InterBase/Firebird\)](#)
- [Roles Wizard \(InterBase/Firebird\)](#)
- [Shadows wizard \(InterBase/Firebird\)](#)
- [Tables Wizard \(InterBase/Firebird\)](#)
- [Triggers Wizard \(InterBase/Firebird\)](#)
- [Unique Keys Wizard \(InterBase/Firebird\)](#)
- [Users wizard \(InterBase/Firebird\)](#)
- [Views Wizard \(InterBase/Firebird\)](#)

### BLOB FILTERS WIZARD (INTERBASE/FIREBIRD)

The Blob Filter wizard lets you create and submit a DECLARE FILTER statement. Properties let you provide input and output subtypes, an entry point, and the module name.



**To create a new blob filter using a wizard:**

- 1 [Open an Object Wizard](#) for a blob filter.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Blob Filters \(InterBase/Firebird\) - Properties](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**BLOB FILTERS (INTERBASE/FIREBIRD) - PROPERTIES**

When creating or editing a Blob Filter, this tab/panel lets you work with the following settings:

Setting	Description
<b>Name</b>	A name that must be unique among filters on the database.
<b>Input subtype</b>	Lets you provide an INPUT_TYPE argument that specifies the Blob subtype from which data is to be converted.
<b>Output subtype</b>	Lets you provide an OUTPUT_TYPE argument that specifies the Blob subtype into which data is to be converted
<b>Entrypoint</b>	Lets you provide an ENTRY_POINT argument. Provide a quoted string specifying the name of the Blob filter as stored in a linked library.
<b>Module name</b>	Lets you provide a MODULE_NAME argument. Provide a quoted file specification identifying the object module in which the filter is stored.

**DOMAINS WIZARD (INTERBASE/FIREBIRD)**

The Domain wizard lets you provide datatype details and other basic clause/option values, in creating a domain.

**To create a new domain using a wizard:**

- 1 [Open an Object Wizard](#) for a domain.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Domains \(InterBase/Firebird\) - Properties](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).

## DOMAINS (INTERBASE/FIREBIRD) - PROPERTIES

When creating or editing a domain, this tab/panel lets you work with the following settings:

Group	Settings and descriptions	
<b>Creation</b>	Lets you provide a <b>Datatype</b> name.	
<b>Base Datatype</b>	Lets you select an InterBase/Firebird-supported <b>Type</b> . The following additional settings are available, depending on the type you choose:	
	<b>Allow Nulls</b>	Disabled, the domain is created with a NOT NULL clause.
	<b>Size, Precision, and Scale</b>	Provide InterBase/Firebird-standard datatype options.
	<b>Array</b>	Clicking the associated search button lets you select the upper and lower bounds for an array dimension.
<b>Character Options</b>	If you choose a character-based <b>Type</b> , the Character Set and Collation settings lets you select CHARACTER SET and COLLATION clause values used when creating the domain.	
<b>Lob Options</b>	If you choose a BLOB <b>Type</b> , you can specify the expected <b>LOB SEgment Length</b> , in bytes.	
<b>Domain Options</b>	<b>Default</b>	Specify a valid InterBase/Firebird DEFAULT argument value, to indicate the default column value that is entered when no other entry is made.
	<b>Check</b>	Specify a condition or requirement to be tested on data values at the time data is entered by applying a check constraint to a column.

## ENCRYPTION KEYS WIZARD (INTERBASE/FIREBIRD)

The Encryption Key wizard lets you create keys used to encrypt and decrypt databases or table columns. It lets you specify the credentials, basic algorithm details, and other encryption options.

**NOTE:** Before working with encryption keys, you should be familiar with InterBase/Firebird requirements regarding passwords and user permissions. See the InterBase documentation at <http://docs.embarcadero.com/products/interbase/>.

**To create a new encryption key using a wizard:**

- 1 [Open an Object Wizard](#) for an encryption key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Encryption Keys \(InterBase/Firebird\) - Properties](#).
  - **DDL View** panel - for details, see [Create the New Object](#).

## Encryption Keys (InterBase/Firebird) - Properties

When creating or editing an encryption key, this tab/panel lets you work with the following settings:

Property	Descriptions
<b>Name</b>	When creating an encryption key, provide a unique name. After creating the encryption key, the name cannot be edited.
<b>Algorithm</b>	Options represent choices of the Advanced Encryption Standard or Data Encryption Standard algorithm and a key length.
<b>Pad random</b>	Enabled, the encryption key is created with a PAD RANDOM option, causing equal values to have different ciphertext. Disabled, the encryption key is created with a PAD NULL option, specifying that random padding should not be performed.
<b>Init Vector random</b>	Enabled, the encryption key is created with an INIT_VECTOR RANDOM option, enabling Cipher Block Chaining (CBC) encryption, forcing equal values to take different ciphertext. Disabled, the encryption key is created with a INIT_VECTOR NULL option, enabling Electronic Cookbook (ECB).
<b>Password</b>	Optionally, provide a user password.
<b>Is Default</b>	When enabled, the encryption key is created with an AS DEFAULT option, making this key the database default when no explicit key is specified for database or column encryption.

## EXCEPTIONS WIZARD (INTERBASE/FIREBIRD)

The Exception wizard lets you provide the text of an exception and the name by which it is raised from a stored procedure.

### To create a new exception using a wizard:

- 1 [Open an Object Wizard](#) for an exception.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Exceptions \(InterBase/Firebird\) - Properties](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).

## EXCEPTIONS (INTERBASE/FIREBIRD) - PROPERTIES

When creating or editing an exception, this tab/panel lets you work with the following settings:

Group	Settings and descriptions
<b>Creation</b>	Lets you provide the <b>Name</b> by which the exception is referenced.
<b>General</b>	Lets you provide the <b>Text</b> of the exception.

## EXTERNAL FUNCTIONS WIZARD (INTERBASE/FIREBIRD)

The External Functions wizard lets you DECLARE an existing, user-defined function.

### To create a new external function using a wizard:

- 1 [Open an Object Wizard](#) for an external function.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [External Functions \(InterBase/Firebird\) - Properties](#).
  - **Parameters** panel - for details, see [External Functions \(InterBase/Firebird\) - Parameters](#).
  - **Return Type** panel - for details, see [External Functions \(InterBase/Firebird\) - Return Type](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).

### EXTERNAL FUNCTIONS (INTERBASE/FIREBIRD) - PROPERTIES

When creating or editing a generator, this tab/panel lets you work with the following settings:

Group	Settings and descriptions	
<b>Creation</b>	Lets you specify the <b>Name</b> of the function.	
<b>General</b>	<b>EntryPoint</b>	Type a quoted identifier as the ENTRY_POINT clause value for the declaration.
	<b>ModuleName</b>	Type a quoted identifier as the MODULE_NAME clause value for the declaration.

### EXTERNAL FUNCTIONS (INTERBASE/FIREBIRD) - PARAMETERS

#### For each parameter to be added to this function:

- 1 Use the New button to add a new parameter and type a name for the parameter in the space provided.
- 2 With the parameter selected, in the **Attributes** area select a **Type**, and if appropriate, provide or select **Precision**, **Scale**, and **Size** options.

Use the Delete button to drop a selected parameter and use the arrow keys to reorder the parameter list.

### EXTERNAL FUNCTIONS (INTERBASE/FIREBIRD) - RETURN TYPE

When creating an external function, this tab/panel lets you specify the datatype of the value returned by the external function. Select a **Type**, and if appropriate, provide or select **Precision**, **Scale**, and **Size** options.

You cannot modify the return type when editing an external function.

**For each parameter to be added to this function:**

- 1 Use the New button to add a new parameter and type a name for the parameter in the space provided.
- 2 With the parameter selected, in the **Attributes** area select a **Type**, and if appropriate, provide or select **Precision**, **Scale**, and **Size** options.

Use the Delete button to drop a selected parameter and use the arrow keys to reorder the parameter list.

## FOREIGN KEYS WIZARD (INTERBASE/FIREBIRD)

The Foreign Key Wizard lets you create a foreign key constraint, mapping the relevant columns and specifying delete and update integrity checks.

**To create a new foreign key using a wizard:**

- 1 [Open an Object Wizard](#) for a foreign key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Foreign keys \(InterBase/Firebird\) - Properties](#)
  - **Column Mapping** panel - for details, see [Foreign Keys \(InterBase/Firebird\) - Column Mapping](#)
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## FOREIGN KEYS (INTERBASE/FIREBIRD) - PROPERTIES

When creating or editing a foreign key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Owner</b>	The owner of the table where the foreign key is being created.
<b>Table Name</b>	The table where the foreign key link originates--the child table.
<b>Name</b>	Lets you provide a constraint name.
<b>Delete Rule and Update Rule</b>	Let you select ON UPDATE and ON DELETE argument values of NONE, CASCADE, SET DEFAULT, SET NULL, NO ACTION, or RESTRICT.

## FOREIGN KEYS (INTERBASE/FIREBIRD) - COLUMN MAPPING

- 1 Under **Referenced Table**, choose the **Owner** and then the **Name** of the referenced, or parent, table.
- 2 Under the **Main Table**, select check boxes corresponding to the columns that are to reference columns in the referenced table. Then, under **Referenced Table**, select the corresponding column check boxes.

## GENERATORS WIZARD (INTERBASE/FIREBIRD)

The Generator wizard lets you create a simple, named generator and set its initial value.

**To create a new generator using a wizard:**

- 1 [Open an Object Wizard](#) for a generator.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Generators \(InterBase/Firebird\) - Properties](#).
  - **DDL View** panel - for details, see [Create the New Object](#).

## GENERATORS (INTERBASE/FIREBIRD) - PROPERTIES

When creating or editing a generator, this tab/panel lets you work with the following settings:

Group	Settings and descriptions
<b>Creation</b>	Lets you provide the <b>Name</b> for the generator.
<b>Current Value</b>	Lets you provide a SET GENERATOR value, setting the value of the generator.

## INDEXES WIZARD (INTERBASE/FIREBIRD)

The Index Wizard lets you create a basic InterBase/Firebird index, specifying the sort order and ACTIVE/INACTIVE status for the index, and indicating whether duplicate values are allowed

**To create a new index using a wizard:**

- 1 [Open an Object Wizard](#) for an index.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Indexes \(InterBase/Firebird\) - Properties](#).
  - **Columns** panel - for details, see [Indexes \(InterBase/Firebird\) - Columns](#).
  - **DDL View** panel - for details, see [Create the New Object](#).

- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

#### INDEXES (INTERBASE/FIREBIRD) - PROPERTIES

When creating or editing an index, this tab/panel lets you work with the following settings:

Setting	Description
<b>Name</b>	The name of the index.
<b>Table Name</b>	The owning table.
<b>Unique</b>	When enabled, the index is created with the UNIQUE keyword, disallowing duplicate values.
<b>Enabled</b>	When enabled, the index is created or altered with an ACTIVE status. When disabled, the index is created or altered with an INACTIVE status.
<b>Descending</b>	Specifies the sort order. When enabled, the index is created with the DESCENDING keyword. When disabled, the index is created with the ASCENDING keyword.

#### INDEXES (INTERBASE/FIREBIRD) - COLUMNS

When adding or editing an index, this tab/panel lets you manage the columns that make up the index.

##### To add a column to the index:

Click the New button and select a column from the **Column** dropdown.

##### To delete a column from the index:

Select the column and click the Remove button.

##### To change the position of a column in the list:

Select the column and use the arrow keys to reorder the column list.

#### PRIMARY KEYS WIZARD (INTERBASE/FIREBIRD)

The Primary Key Wizard lets you add a basic InterBase/Firebird primary key to a table.

**To create a new primary key using a wizard:**

- 1 [Open an Object Wizard](#) for a primary key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Primary Keys \(InterBase/Firebird\) - Properties](#).
  - **Columns** panel - [Primary Keys \(InterBase/Firebird\) - Columns](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**PRIMARY KEYS (INTERBASE/FIREBIRD) - PROPERTIES**

When creating or editing an index, this tab/panel lets you work with the following settings:

Setting	Description
<b>Name</b>	The name of the primary key.
<b>Table Name</b>	The owning table.

**PRIMARY KEYS (INTERBASE/FIREBIRD) - COLUMNS**

When adding or editing a primary key, this tab/panel lets you manage the columns that make up the primary key.

**NOTE:** Only columns defined with the **Allow Nulls** property disabled can be added to a primary key.

**To add a column to the primary key:**

Click the New button and select a column from the **Column** dropdown.

**To delete a column from the primary key:**

Select the column and click the Remove button.

**To change the position of a column in the list:**

Select the column and use the arrow keys to reorder the column list.

**PROCEDURES WIZARD (INTERBASE/FIREBIRD)**

The Procedures Wizard lets you create a procedure definition that includes parameter specifications, the body of the procedure, and execution permissions.



**To create a new procedure using a wizard:**

- 1 [Open an Object Wizard](#) for a procedure.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Procedures \(InterBase/Firebird\) - Properties](#).
  - **Parameters** panel - [Procedures \(InterBase/Firebird\) - Parameters](#).
  - **Body** panel - [Procedures \(InterBase/Firebird\) - Body](#).
  - **Permissions** panel - [Set Permissions for an Object](#)
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**PROCEDURES (INTERBASE/FIREBIRD) - PROPERTIES**

When creating a procedure, this panel/tab lets you provide a **Name** for the procedure. After creation, you cannot change its name.

**PROCEDURES (INTERBASE/FIREBIRD) - PARAMETERS**

This tab lets you add and modify the parameters for a procedure.

**To add a parameter:**

- 1 Click the New button and type a name for the parameter in the space provided.
- 2 In the **Attributes** area, select a **Type** and if appropriate, provide **Precision**, **Scale**, and **Size** values.
- 3 Select a **Parameter Mode** of INPUT or OUTPUT.

**To edit a parameter:**

- Select the parameter and modify the values in the **Attributes** area.

**To delete a parameter:**

- Select the parameter and click the Delete button.

**To change the ordering position of a parameter:**

- Select the parameter and use the arrow buttons to move the parameter up or down in the list.

## PROCEDURES (INTERBASE/FIREBIRD) - BODY

This tab provides a simple editor that lets you add or modify the body of the procedure.

## ROLES WIZARD (INTERBASE/FIREBIRD)

The Roles Wizard lets you create a basic InterBase/Firebird role.

**To create a new role using a wizard:**

- 1 [Open an Object Wizard](#) for a role.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Roles \(InterBase/Firebird\) - Properties](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

The Roles editor opens, letting you immediately assign users to the new roles and set up a set of permissions for the role. For more information, see [Roles editor \(InterBase/Firebird\)](#).

## ROLES (INTERBASE/FIREBIRD) - PROPERTIES

When creating a role, this tab/panel lets you provide a **Name** for the role. When editing a role, this tab/panel displays the **Name** and **Authorization Owner** for the role.

## SHADOWS WIZARD (INTERBASE/FIREBIRD)

The Shadows Wizard lets you create a shadow, consisting of one or more in-sync copies of the database stored on secondary storage devices.

**To create a new shadow using a wizard:**

- 1 [Open an Object Wizard](#) for a shadow.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Shadows \(InterBase/Firebird\) - Properties](#).
  - **Storage** panel - [Shadows \(InterBase/Firebird\) - Storage](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## SHADOWS (INTERBASE/FIREBIRD) - PROPERTIES

When creating or editing a shadow, this tab/panel lets you work with the following settings:

Setting	Description
<b>Shadow Set</b>	Lets you specify a set number, designating a shadow set to which files specified on the <b>Storage</b> panel belong. Once the shadow set is created, this value cannot be edited.
<b>Conditional</b>	Lets you include the <b>CONDITIONAL</b> argument with the <b>CREATE SHADOW</b> statement, when creating a new shadow. This allows shadowing to continue if the primary shadow becomes unavailable or if the shadow replaces the database due to disk failure
<b>Behavior</b>	Lets you include <b>AUTO</b> or <b>MANUAL</b> arguments with the <b>CREATE SHADOW</b> statement, specifying default access behavior when no shadow is available. <b>AUTO</b> specifies that all attachments and accesses succeed, all references to the shadow are deleted, and the shadow file is detached. <b>MANUAL</b> specifies that database attachments and accesses fail until a shadow becomes available or until all shadow references are removed from the database.

## SHADOWS (INTERBASE/FIREBIRD) - STORAGE

When creating or editing a shadow, this tab/panel lets you manage a primary and one or more secondary files making up the shadow set. Note the following when managing shadow files:

- When creating a shadow, a primary file entry is automatically created when you select this tab. You must provide primary file property values before proceeding.
- A secondary file cannot be created if there is no primary file entry.

**To add a primary or secondary file to the shadow:**

- 1 Click the **New** button.
- 2 Use the following table as a guide to providing values in the **Property/Value** list for the file entry.

Setting	Description
<b>Physical File Name</b>	Provide the path and file name for the shadow file, which must reside on the local file system.
<b>Starting Page</b>	Starting page number of a secondary shadow file.
<b>Size</b>	The size, in pages.

**To delete a primary or secondary file entry:**

Select the file entry in the primary and secondary file list and click the **Delete** button.

## TABLES WIZARD (INTERBASE/FIREBIRD)

The Tables wizard lets you create a table definition, providing column descriptions, indexes, permissions, and constraints for the table.

**To create a new table using a wizard:**

- 1 [Open an Object Wizard](#) for a table.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Tables \(InterBase/Firebird\) - Properties](#).
  - **Columns** panel - for details, see [Tables \(InterBase/Firebird\) - Columns](#).
  - **Indexes** panel - for details, see [Tables \(InterBase/Firebird\) - Indexes](#).
  - **Constraints** panel - for details, see [Tables \(InterBase/Firebird\) - Constraints](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **Permissions** panel - [Set Permissions for an Object](#)
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## TABLES (INTERBASE/FIREBIRD) - PROPERTIES

When creating a table, this tab/panel lets you provide a **Name** for the table. When editing a table, the **Name** is for display only.

## TABLES (INTERBASE/FIREBIRD) - COLUMNS

When creating or editing a table, this tab/panel lets you manage the columns for the table:

**To add a column to the table:**

- 1 Use the **Add Column** dropdown to add a new column to the bottom of the column list or to insert a new column above the currently selected column in the column list.
- 2 Provide a **Name** for the column.
- 3 Either select a **Type** for the column or select **Computed** and provide a valid InterBase/Firebird **Computed Expression**.
- 4 Use the following table as a guide to providing additional property values, noting that availability of a property differs by data type and other property selections.

Property	Description
<b>Array Dimensions</b>	Lets you type a valid InterBase/Firebird array_dim specifier or use the ... button to open a dialog that lets you build the specifier.
<b>Size</b>	Lets you provide a size for character datatypes.
<b>Allow Nulls</b>	Disabled, submits a NOT NULL argument.
<b>Default Character Set</b>	Submits a CHARACTER SET argument value.
<b>Default Collation</b>	Submits a COLLATE argument value.

Property	Description
Default Value	Submits a literal COLLATE argument value.

**To edit a column:**

- Select the column from the list and edit property values as per the descriptions above.

**To drop a column:**

- Select the column from the list and click the **Delete** button.

**To change the position of a column in the list:**

- Select the column from the list and use the arrow buttons to change its position.

## TABLES (INTERBASE/FIREBIRD) - INDEXES

When creating or editing tables, this tab/panel lets you manage indexes for the table:

- Click **Add** to open a wizard that lets you add a new index to the table. For information on using the wizard, see [Indexes Wizard \(InterBase/Firebird\)](#).
- Select an index and click **Edit** to open an editor on the index. For information on using the editor, see [Indexes editor \(InterBase/Firebird\)](#).
- Select an index and click **Drop** to drop an index from the table.

## TABLES (INTERBASE/FIREBIRD) - CONSTRAINTS

When creating or editing tables, this tab/panel lets you manage foreign key, primary key, unique key and check constraints for the table:

- Select a constraint type and click **Add** to open a wizard that lets you add a new constraint of that type to the table.
- Select an existing constraint and click **Edit** to open a wizard that lets you add modify that constraint.

**NOTE:** The wizards used to add and modify constraints offer functionality identical to the editors and wizards used when creating those object types. For details, see the following topics:

[Foreign Key Wizard \(IBM DB2 LUW\)](#) and [Foreign Keys editor \(InterBase/Firebird\)](#)

[Primary Keys Wizard \(InterBase/Firebird\)](#) and [Primary Keys editor \(InterBase/Firebird\)](#)

[Unique Keys Wizard \(InterBase/Firebird\)](#) and [Unique Keys editor \(InterBase/Firebird\)](#)

[Add or Modify Constraint](#)

- Select a constraint and click **Drop** to drop a constraint from the table.

## TRIGGERS WIZARD (INTERBASE/FIREBIRD)

The Triggers wizard lets you create a basic InterBase trigger, providing a trigger body and specifying key properties such as the associated event type and the trigger timing.

### To create a new trigger using a wizard:

- 1 [Open an Object Wizard](#) for a trigger.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Triggers \(InterBase/Firebird\) - Properties](#).
  - **Body** panel - for details, see [Triggers \(InterBase/Firebird\) - Body](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

### TRIGGERS (INTERBASE/FIREBIRD) - PROPERTIES

When creating a trigger, this tab/panel lets you provide values for the following trigger properties:

Property	Description
<b>Parent Type</b>	Select TABLE or VIEW as the Parent type.
<b>Parent Name</b>	Select the table or view that causes the trigger to fire when the specified operation occurs.
<b>Name</b>	Provide a name for the trigger.
<b>Trigger Timing</b>	Select whether the trigger fires BEFORE or AFTER the associated operation.
<b>Trigger Events</b>	Specifies the table operation (DELETE, INSERT, or UPDATE) that causes the trigger to fire.
<b>Enabled</b>	Enabled, this options submits an ACTIVE argument. Otherwise, an INACTIVE argument is submitted.
<b>Position</b>	Specify a POSITION number (0 - 32,767) to control firing order for triggers before or after the same action.

When editing a trigger, the values cannot be modified.

### TRIGGERS (INTERBASE/FIREBIRD) - BODY

When creating or editing a trigger, this tab/panel previews the DDL reflecting your property values and lets you provide or modify the trigger body.

## UNIQUE KEYS WIZARD (INTERBASE/FIREBIRD)

The Unique Key Wizard lets you add a basic InterBase/Firebird unique key to a table.

### To create a new unique key using a wizard:

- 1 [Open an Object Wizard](#) for a unique key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Unique Keys \(InterBase/Firebird\) - Properties](#).
  - **Columns** panel - [Unique Keys \(InterBase/Firebird\) - Columns](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

### UNIQUE KEYS (INTERBASE/FIREBIRD) - PROPERTIES

When creating a unique key, this tab/panel lets you provide a **Name** for the unique key and select the owning **Table Name**. These values cannot be changed when editing the unique key.

### UNIQUE KEYS (INTERBASE/FIREBIRD) - COLUMNS

When adding or editing a unique key, this tab/panel lets you manage the columns that make up the unique key.

**NOTE:** Only columns defined with the **Allow Nulls** property disabled can be added to a unique key.

### To add a column to the unique key:

Click the New button and select a column from the **Column** dropdown.

### To delete a column from the unique key:

Select the column and click the Remove button.

### To change the position of a column in the list:

Select the column and use the arrow keys to reorder the column list.

## USERS WIZARD (INTERBASE/FIREBIRD)

The Users wizard lets you create a basic InterBase/Firebird user, providing basic properties, assigning roles, and granting relevant permissions.

**NOTE:** The Users node is only available for databases with the Embedded User Authentication option enabled. Before using Rapid SQL against InterBase/Firebird users, you should be familiar with this feature and with InterBase/Firebird permissions with regard to this object type. For access to InterBase documentation, see <http://docs.embarcadero.com/products/interbase/>.

### To create a new user:

- 1 [Open an Object Wizard](#) for a user.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Users \(InterBase/Firebird\) - Properties](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **Roles** panel - for details, see [Users \(InterBase/Firebird\) - Roles](#).
  - **Object Permissions** panel - [Set Permissions for an Object](#)
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

The User editor opens. For more information, see [Users editor \(InterBase/Firebird\)](#).

### USERS (INTERBASE/FIREBIRD) - PROPERTIES

When creating or editing a user, this tab/panel lets you provide values for the following properties:

Property	Description
<b>Name</b>	The login name for this user.
<b>Password</b>	The password for this user.
<b>Group</b>	Lets you specify a group name for the user. If unspecified, the user is created with the NO GROUP NAME option.
<b>Default Role</b>	Lets you select a default role for the user. If unspecified, the user is created with the NO DEFAULT ROLE option.  Additional roles can be assigned for this user on the Roles tab/panel. For details, see <a href="#">Users (InterBase/Firebird) - Roles</a> .
<b>First Name, Middle Name, and Last Name</b>	Let you provide the user's name, as opposed to the login name.
<b>Active</b>	This control is not enabled when creating a user and by default the user is created with the ACTIVE option. When editing a user, selecting this check box allows the user to log into the database. When unselected, the user is prevented from logging in to the database.



Property	Description
<b>System User Name</b>	Lets you specify a systemuser name for the user. If unspecified, the user is created with the NO SYSTEM USER NAME option.
<b>Description</b>	Lets you provide a short description for this user
<b>GID</b>	Lets you specify a numeric group ID for this user.
<b>UID</b>	Lets you specify a numeric user ID for this user.

### USERS (INTERBASE/FIREBIRD) - ROLES

When creating or editing a user, this tab/panel lets you assign the valid, non-default roles that a user can specify when logging in to the database.

**NOTE:** A default login role can be assigned for this user on the Properties tab/panel. For details, see [Users \(InterBase/Firebird\) - Roles](#).

#### To assign roles to the user:

- Select the associated check box.

### VIEWS WIZARD (INTERBASE/FIREBIRD)

The Views wizard lets you create a basic InterBase/Firebird view

#### To create a new view using a wizard:

- 1 [Open an Object Wizard](#) for a view.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Views \(InterBase/Firebird\) - Properties](#).
  - **Comment** panel - for details, see [Add a Comment to an Object](#).
  - **Definition** panel - for details, see [Views \(InterBase/Firebird\) - Definition](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

### VIEWS (INTERBASE/FIREBIRD) - PROPERTIES

When creating or editing a view, this tab/panel lets you provide values for the following properties:

Property	Description
<b>Name</b>	Provide a name for the view.

Property	Description
<b>Check Option</b>	When enabled, a WITH CHECK OPTION argument is included with the DDL generated to create the view. This prevents INSERT or UPDATE operations on an updatable view if that operation violates the search condition specified in the associated WHERE clause.

When editing a view, these values cannot be modified.

#### VIEWS (INTERBASE/FIREBIRD) - DEFINITION

When creating or editing a view, this tab/panel previews the DDL reflecting your property values and lets you provide or edit the relevant query.

## MICROSOFT SQL SERVER OBJECT WIZARDS

Rapid SQL lets you create SQL Server objects using the following wizards:

- [Database Wizard \(SQL Server\)](#)
- [Default Wizard \(SQL Server\)](#)
- [Extended Procedure Wizard \(SQL Server\)](#)
- [Foreign Key Wizard \(SQL Server\)](#)
- [Function Wizard \(SQL Server\)](#)
- [Index Wizard \(SQL Server\)](#)
- [Login Wizard \(SQL Server\)](#)
- [Primary Key Wizard \(SQL Server\)](#)
- [Procedure Wizard \(SQL Server\)](#)
- [Rule Wizard \(SQL Server\)](#)
- [Schema Wizard \(SQL Server\)](#)
- [Table Wizard \(SQL Server\)](#)
- [Trigger Wizard \(SQL Server\)](#)
- [Unique Key Wizard \(SQL Server\)](#)
- [User Message Wizard \(SQL Server\)](#)
- [User Wizard \(SQL Server\)](#)
- [User Datatype Wizard \(SQL Server\)](#)
- [View Wizard \(SQL Server\)](#)

## DATABASE WIZARD (SQL SERVER)

The Database Wizard presents you with a different set of options based on your server version to create the database accurately on each platform.

**TIP:** Microsoft SQL Server recommends that you do not create any user objects, such as tables, views, stored procedures, or triggers, in the master database. The master database includes the system tables that store the system information used by SQL Server, such as configuration option settings.

### To create a new database using a wizard:

- 1 [Open an Object Wizard](#) for a database.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Options** panel - for details, see [Databases \(SQL Server\) - Options](#).
  - **Placement** panel - for details, see [Databases \(SQL Server\) - Placement](#).
  - **Transaction Log** panel - for details, see [Databases \(SQL Server\) - Transaction Log](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## DATABASES (SQL SERVER) - OPTIONS

When creating or editing a database, this tab/panel lets you work with the following settings:

Setting	Description
<b>Name</b>	Provide a name for the database.
<b>Attach existing OS files</b>	To create a database from an existing set of operating system files, there must be a <filespec> entry for the first PRIMARY file. The PRIMARY filegroup contains all the database system tables. Primary files have a .mdf extension.
<b>Compatible Level</b>	Select a version compatibility level.
<b>Properties group</b>	Select the following settings: <b>ANSI null default, ANSI nulls, ANSI padding, ANSI warnings, auto create statistics, auto update statistics, autoclose, autoshrink, concat null yields null, cursor close on commit, arithabort, db chaining, dbo use only, default to local cursor, merge publish, numeric roundabout, offline, published, quoted identifier, read only, recursive triggers, select into/bulkcopy/pllsort, single user, subscribed, torn page detection, and trunc log on chkpt.</b>

## DATABASES (SQL SERVER) - PLACEMENT

Indicate the file where you want the database to live. For example, a new Books database could include author and title filegroups.

By default, when you open the Wizard and click the **Placement** tab, a filegroup definition, using the name you provided for the database and default settings, is displayed. For each filegroup to be added, click the New button, provide a **Device File Name** for the filegroup, and use the **File Group Properties** and **Device File Properties** groups to provide the attributes of the filegroup.

Use the Delete button to delete a selected filegroup.

#### DATABASES (SQL SERVER) - TRANSACTION LOG

The transaction log file is a required file for each database. This file holds the log information to recover the database. There can be multiple log files for a database, but there has to be at least one. Traditionally the logfile extension has been .ldf.

By default, when you open the Wizard and click the **Transaction Log** tab, a transaction log file definition, using the name derived from the name you provided for the database and with default settings, is displayed. For each file to be added, click the New button, provide a **Device File Name**, and use the **Log Device Properties** group to provide the attributes of the file.

Use the Delete button to delete a selected file.

**NOTE:** As you complete the wizard, be aware that the Primary file contains startup information for the database and is also used to store data. The transaction log files hold the information used to recover the database.

#### DEFAULT WIZARD (SQL SERVER)

When bound to a column or user-defined object, a default ensures that a specific value will be inserted into the column where the object will be bound if no explicit value is given.

The Default Wizard lets you name the default and specify its value.

#### To create a new default using a wizard:

- 1 [Open an Object Wizard](#) for a default.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Defaults \(SQL Server\) - Properties](#).
  - **Dependencies** panel - [Defaults \(SQL Server\) - Dependencies](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

### DEFAULTS (SQL SERVER) - PROPERTIES

When creating or editing a default, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the schema that is to own the default.
<b>Name</b>	Provide a name for the default.
<b>Value</b>	Provide the value of the default.

### DEFAULTS (SQL SERVER) - DEPENDENCIES

From the **Type** dropdown, choose Column or Datatype, and if you chose **Column**, choose a Table from the **Table** dropdown. The list on the left is populated with candidate columns or datatypes. To move a candidate from the list on the left to the dependencies column on the right, select the candidate and click **Add**. Remove columns or datatypes from the dependencies list on the right by selecting the column or datatype and clicking **Remove**.

### EXTENDED PROCEDURE WIZARD (SQL SERVER)

Extended Procedures are dynamic link libraries that can be used to load and execute application routines written in other programming languages, such as C or Visual Basic. Extended Procedures function and appear in the same manner as SQL Server stored procedures in that you can pass parameters to them and obtain results.

Extended stored procedures provide a method for calling procedural language functions from within the Adaptive Server.

**NOTE:** Extended procedures can only be created in the Master database.

#### To create a new Extended procedure using a wizard:

- 1 [Open an Object Wizard](#) for an extended procedure.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Extended Procedures \(SQL Server\) - Properties](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

### EXTENDED PROCEDURES (SQL SERVER) - PROPERTIES

When creating or editing an extended procedure, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the extended procedure.

Setting	Description
<b>Name</b>	Provide a name for the extended procedure.
<b>Library</b>	Provide the name of the DLL containing the extended procedure.

## FOREIGN KEY WIZARD (SQL SERVER)

Foreign keys are unique values that refer to specific columns of other tables. Thus, a foreign key links two tables together. Embarcadero Rapid SQL's Foreign Key Wizard makes it easy for you to create a relational link between two tables, thereby speeding queries and giving you faster access to data. The column in the initial table, the parent table, is called the primary key. The corresponding column in the (child) table that references the primary key, is the foreign key. Foreign keys can also refer to columns within the same table.

### To create a new Foreign Key using a wizard:

- 1 [Open an Object Wizard](#) for a foreign key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Foreign Keys \(SQL Server\) - Properties](#).
  - **Column Mapping** panel - [Foreign Keys \(SQL Server\) - Column Mapping](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## FOREIGN KEYS (SQL SERVER) - PROPERTIES

When creating or editing a foreign key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Owner and Table Name</b>	Select the owner and name of the table for which the foreign key is being created.
<b>Name</b>	Provide a name for the foreign key.
<b>Enabled</b>	Enables or disables the foreign key.
<b>Not For Replication</b>	Replication copies and distributes data and database objects from one database to another and then synchronizes information between databases for consistency.
<b>Delete Rule</b>	If you choose the CASCADE option, all rows containing data involved with the foreign key will be deleted after a delete operation.
<b>Update Rule</b>	If you choose the CASCADE option, all rows containing data involved with the foreign key will be deleted after an update operation.

## FOREIGN KEYS (SQL SERVER) - COLUMN MAPPING

Under **Referenced Table**, choose the **Owner** and then the **Name** of the referenced, or parent, table.

Under the **Main Table**, select checkboxes corresponding to the columns that are to reference columns in the referenced table. Then, under **Referenced Table**, select the corresponding column check boxes.

## FUNCTION WIZARD (SQL SERVER)

Functions are subroutines that you define so you can reuse code without having to reinvent the wheel each time. You can use functions to determine the best methods for controlling access and manipulation of the underlying data contained in an object. A function returns a value, unlike a stored procedure, which does not.

- To create a user-defined function, you need CREATE ANY privileges or IMPLICIT\_SCHEMA authority on the database if the schema does not already exist.

### To create a new function using a wizard:

- 1 [Open an Object Wizard](#) for a function.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Functions \(SQL Server\) - Properties](#).
  - **Definition** panel - for details, see [Functions \(SQL Server\) - Definition](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## FUNCTIONS (SQL SERVER) - PROPERTIES

When creating or editing a function, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the function.
<b>Name</b>	Provide a name for the function.
<b>Schema Binding</b>	Choose whether the function is bound to database objects that it references.
<b>Encryption</b>	Choose whether SQL Server encrypts table columns that contain the text of the CREATE FUNCTION statement.

## Functions (SQL Server) - Definition

Complete the CREATE FUNCTION outline provided by typing or pasting the body of the function. Finally, use the **Schedule** or **Execute** buttons to [Create the New Object](#).

## INDEX WIZARD (SQL SERVER)

Like an index in a book, a table index helps you get at the data you want without having to read through the whole table. Indexes can exist on single column or on multiple columns. Indexes appear in the form of B-trees. And, as for books, you can have multiple indexes for a single table. You can also create indexes for a view.

### To create a new index using a wizard:

- 1 [Open an Object Wizard](#) for an extended procedure.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Indexes \(SQL Server\) - Properties](#).
  - **Columns** panel - [Indexes \(SQL Server\) - Columns](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## INDEXES (SQL SERVER) - PROPERTIES

When creating or editing an index, this tab/panel lets you work with the following settings:

Setting	Description
<b>Parent Type</b>	Select TABLE or VIEW.
<b>Parent Owner</b>	Select the owner of the table or view.
<b>Parent Name</b>	Select the specific table or view containing the columns you want to index.
<b>Name</b>	Provide a name for the index.
<b>Build Online</b>	Enabling this feature specifies that the ONLINE=ON clause is used when creating this object and can subsequently be used when rebuilding or dropping this object.
<b>Max degree of parallelism</b>	Lets you specify a MAXDOP index operation value, limiting the number of processors used in parallel plan execution.
<b>Index Type</b>	Select UNIQUE or NONUNIQUE. An index is unique when no two rows are permitted to have the same index value. (Note: A clustered index on a view must be unique.) If an INSERT or UPDATE statement creates a duplicate value, the operation may fail.
<b>Clustered</b>	<p>A clustered index is one in which the physical order of rows matches the order of indexed rows. A table or view can only have one clustered index at a time.</p> <p>In a nonclustered index, the physical order of rows is independent of the indexed order of rows.</p> <p>For an indexed view, you can only create a nonclustered index if there is already a clustered index extant.</p> <p>With this check box enabled, subsequent rebuild and drop operations offer an online option.</p>



Setting	Description
<b>Ignore Duplicate Key</b>	This option controls what happens when an attempt is made to insert a duplicate key value into a column that is part of a unique clustered index.  If the option is selected and an INSERT statement that creates a duplicate key is executed, SQL Server issues a warning and ignores the duplicate row.  If not selected, SQL Server issues an error message and rolls back the entire INSERT statement.
<b>Statistics Recompute</b>	Enabling this feature means queries involving the table run at the optimal level as distribution statistics are updated automatically when the index is created. If you disable this option, you can compromise query performance.
<b>File Group</b>	
<b>Fill Factor</b>	Fillfactor specifies a percentage (0-100) that indicates how full each leaf level of each index page during index creation should be. When an index page fills up, time is needed to split the index page to make room for more rows. This can affect performance, so choose the fillfactor value carefully. When a fillfactor is given, SQL server rounds up the number of rows on each index page. If no value is given, the default is zero.
<b>Pad Index</b>	Enable or disable padding of index pages.
<b>Sort in TempDB</b>	Select to store the intermediate index sort results in tempdb. This option may reduce the time needed to create an index if tempdb is on a different set of disks than the user database, but it increases the amount of disk space used to create an index. In addition to the space required in the user database to create the index, tempdb must have about the same amount of additional space to hold the intermediate sort results.
<b>Allow Row Locks and Allow Page Locks</b>	Lets you enable locking granularity at the page and row level (Default=TRUE)  <b>NOTE:</b> You cannot reorganize an index (primary key, or unique key) that has an <b>Allow Page Locks</b> property set to FALSE. For information on reorganizing indexes, see <a href="#">Reorganizing SQL Server indexes, primary keys, and unique keys</a> .

## INDEXES (SQL SERVER) - COLUMNS

From the **Column** dropdown, select a column for the index and specify a **Sort** option. To add more columns, click the New button and then follow the steps in the last instruction. Use the Delete button to drop columns.

## LOGIN WIZARD (SQL SERVER)

Your login account controls access to the server and all of the databases within it. Only the System Administrator or System Security Officer can create logins. Once you can log into a server, you need additional privileges to access user databases. Specifically, each database owner adds the login as a user or alias to the database.

### To Open the Login Wizard

- 1 On the Explorer, open the database where you want to add a login.
- 2 Expand the Security node.
- 3 Right-click Logins and then click New.

The table that follows describes the fields you may encounter as you complete the wizard:

Required Information	Description
What is the name of the login?	You can use either the name of an individual, a role, or whatever you please.
Specify whether this login account will be for a SQL Server Login or a Windows NT User or NT Group	Self-explanatory
What should be the login's default database? (Master is not recommended)	Self-explanatory. Master is not recommended to prevent users from creating objects there by mistake.
What is the login's default language? The server's default is used if you don't make a selection	Self-explanatory
Select databases to create users Database roles	Self-explanatory. Note: The options on the panel vary by version of MS SQL Server.

## PRIMARY KEY WIZARD (SQL SERVER)

Primary key constraints make sure that no duplicate values or NULLS are entered in the columns you specify. You can use primary key constraints to enforce uniqueness and referential integrity. A table can only have a single primary key constraint.

The dialog box lets you specify the owner and table on which you want to place the primary key constraint.

### To create a new primary key using a wizard:

- 1 [Open an Object Wizard](#) for a primary key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Primary Keys \(SQL Server\) - Properties](#).
  - **Columns** panel - [Primary Keys \(SQL Server\) - Properties](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## PRIMARY KEYS (SQL SERVER) - PROPERTIES

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Owner and Table Name</b>	Choose the owner and name of the table in which the primary key is being created.
<b>Name</b>	Provide a name of the primary key being created.
<b>Build Online</b>	Enabling this feature specifies that the ONLINE=ON clause is used when creating this object and can subsequently be used when rebuilding or dropping this object.
<b>Max degree of parallelism</b>	Lets you specify a MAXDOP index operation value, limiting the number of processors used in parallel plan execution.
<b>Clustered</b>	Enable or disable clustering. With this check box enabled, subsequent rebuild and drop operations offer an online option.
<b>File Group</b>	If you do not specify a filegroup, Microsoft SQL Server creates the index in the default filegroup.
<b>Fill Factor</b>	This specifies how full each index page that's storing data should be. The fill factor is a percentage value between 0 and 100.

## Primary Keys (SQL Server) - Columns

From the **Column** dropdown, select a column for the primary key and specify a **Sort** option. To add more columns, click the New button and then follow the steps in the last instruction. Use the Delete button to drop columns.

## PROCEDURE WIZARD (SQL SERVER)

Procedures are a reusable block of PL/SQL, stored in the database, that applications can call. Procedures streamline code development, debugging, and maintenance by being reusable. Procedures enhance database security by letting you write procedures granting users execution privileges to tables rather than letting them access tables directly.

The Procedure Wizard lets you:

- Name the procedure and specify its body.
- Specify any execution options and you can encrypt the stored procedure text in syscomments.

### To create a new procedure using a wizard:

- 1 [Open an Object Wizard](#) for a procedure.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Procedures \(SQL Server\) - Properties](#).
  - **Definition** panel - for details, see [Procedures \(SQL Server\) - Definition](#).

- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## PROCEDURES (SQL SERVER) - PROPERTIES

When creating or editing a procedures, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the procedure.
<b>Name</b>	Provide a name for the procedure
<b>Procedure Number</b>	Optionally, provide a procedure number. By using a number, you can group procedures of the same name together. This also enables you to drop them using only one DROP PROCEDURE statement. So, the procedures bill;1, bill;2, bill;3, etc. will be dropped simultaneously when the time comes.
<b>Replication</b>	This option creates a procedure that is used as a stored procedure filter and is executed only during replication.
<b>Recompile</b>	The plan for this procedure will not be cached and the procedure is recompiled when it is run. This option is appropriate when you're using atypical or temporary values and you don't want to override the execution plan cached in memory.
<b>Encryption</b>	If you select this option, SQL Server will encrypt the syscomments table entry containing the text of the CREATE PROCEDURE statement. It keeps the procedure from being published as part of replication.

## PROCEDURES (SQL SERVER) - DEFINITION

Complete the CREATE PROCEDURE outline provided by typing or pasting the body of the procedure. Finally, use the **Execute** button to [Create the New Object](#).

## RULE WIZARD (SQL SERVER)

Rules promote data integrity by allowing you to validate the values supplied to a table column. They are reusable objects that you can bind to table columns or user datatypes. Check constraints are similar to rules, and are in fact the preferred way of restricting data. A column or user-defined data type can have only one rule bound to it, but a column can have both a rule and one or more check constraints associated with it. Not that a rule cannot apply to data already existing in the database at the time you're creating the rule and can't be bound to a system-created data type. If you create a new rule when one already exists, the new rule will override the previous one.

**To create a new rule using a wizard:**

- 1 [Open an Object Wizard](#) for a rule.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Rules \(SQL Server\) - Properties](#).
  - **Dependencies** panel - [Rules \(SQL Server\) - Dependencies](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**RULES (SQL SERVER) - PROPERTIES**

When creating or editing a rule, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the Owner of the rule.
<b>Name</b>	Provide a name for the rule.
<b>Restriction</b>	Type the condition. The rule restriction is the condition that defines the rule and can be any expression valid in a WHERE clause and can include such elements as arithmetic operators, relational operators, and predicates (for example, IN, LIKE, BETWEEN).

**RULES (SQL SERVER) - DEPENDENCIES**

From the **Type** dropdown, choose Column or Datatype, and if you chose **Column**, choose a Table from the **Table** dropdown. The list on the left is populated with candidate columns or datatypes. To move a candidate from the list on the left to the dependencies column on the right, select the candidate and click **Add**. Remove columns or datatypes from the dependencies list on the right by selecting the column or datatype and clicking **Remove**.

**SCHEMA WIZARD (SQL SERVER)**

Rapid SQL lets you create a schema on Microsoft SQL Server.

**To create a new schema using a wizard:**

- 1 [Open an Object Wizard](#) for a schema.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Schema \(SQL Server\) - Properties](#).
  - **Permissions** panel - [Schema \(SQL Server\) - Permissions](#).
  - **DDL View** panel - for details, see [Create the New Object](#).

- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

#### SCHEMA (SQL SERVER) - PROPERTIES

When creating or editing a schema, this tab/panel lets you provide a **Name** and select the **Owner** of the schema.

#### SCHEMA (SQL SERVER) - PERMISSIONS

For each specific permission to be granted, select the cell corresponding to the name and specific permission, and click the **Grant** button. To revoke a privilege, select a cell showing a Granted permission and click **Revoke**.

#### TABLE WIZARD (SQL SERVER)

A table is a column-based arrangement of data in which the content of one column has a bearing on the other column(s). So, for example, a table might have a column for authors, another column for the books each author has written, and a third for the number of copies each title by a given author has sold. The data moves across the columns in rows.

You must have CREATE TABLE permissions to generate a new table.

#### To create a new table using a wizard:

- 1 [Open an Object Wizard](#) for a table.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Dependencies** panel - [Tables \(SQL Server\) - Properties](#).
  - **Columns** panel - [Tables \(SQL Server\) - Columns](#).
  - **Indexes** panel - [Tables \(SQL Server\) - Indexes](#).
  - **Constraints** panel - [Tables \(SQL Server\) - Constraints](#).
  - **Permissions** panel - [Set Permissions for an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

#### TABLES (SQL SERVER) - PROPERTIES

When creating or editing a table, this tab/panel lets you work with the following settings:

Setting	Description
<b>Schema</b>	Select the owner of the table.

Setting	Description
<b>Name</b>	Provide a name for the table
<b>ANSI_NULLS option</b>	By setting this option, you are setting ANSI_NULLS, ANSI_PADDING, ANSI_WARNINGS, and QUOTED_IDENTIFIER to on so the table can be used in an Indexed View.
<b>Filegroup</b>	A filegroup categorizes the OS files containing data from a single SQL Server database to simplify some tasks, like backing up. A single database can contain more than one filegroup but a filegroup can't contain more than the OS files for a single database. Pick the filegroup you want to use from the drop-down list. If you make no selection, the default is used.
<b>Text Image Filegroup</b>	Unless you make a selection, the text and image columns will be stored on the same filegroup as the table. Because these datatypes may contain lots of data, putting them on a different filegroup can help performance.

### TABLES (SQL SERVER) - COLUMNS

Use the **Add Column** button to add the columns for the table. After providing a **Name** for a new column, you can modify column properties in the **Property/Value** list. Available properties depend on the datatype you choose as well as on the property values you select:

**Computed** and **Computed Expression** - Let you define a column as a computed column and provide the computed column expression.

**Type** - Lets you select a datatype (depending on the type, additional properties such as **Size**, **Width**, and **Scale** may be available).

**Identity Column** - Select this check box to define the column as an identity column.

**Allow Nulls** - Select this check box to allow nulls in this column.

**Default Collation** - available for text/character datatypes, lets you specify a default collation.

**Default Value** - Lets you type a constant value or select a function returning a constant value to serve as the default for the column

**Default Binding** and **Rule Binding** - Let you bind a rule or default to a column.

**Is Sparse** - Available to columns that allow NULL values, optimizes storage of the column for null values. This property does apply to the following data types: text, ntext, image, timestamp, user-defined data type, geometry, or geography. Columns with default values, default or rule bindings, cannot be defined as sparse. Computed columns cannot be defined as Sparse, but the columns in the computed expression can be Sparse columns.

Optionally, you can select a column and modify its values or select a column and **Delete** it.

**NOTE:** For SQL Server 2000, if you create a table with a Column datatype = text., you can set the storage and image values on the Storage Tab of the Tables Editor Storage Tab. When you have a text datatype, the Storage Tab displays a Text In Row box where you can specify the maximum size to be stored.

**NOTE:** Because the smalldatetime datatype stores dates and time with less precision than the datetime datatype, before outputting you use the CAST or CONVERT functions to convert any boxes with the smalldatetime datatype to either VARCHAR or datetime datatypes. For more information, see SQL Server Books Online, Transact-SQL Reference.

## TABLES (SQL SERVER) - INDEXES

Click Add to open the [Index Wizard \(DB2 Z/OS\)](#).

## TABLES (SQL SERVER) - CONSTRAINTS

Selecting a constraint type and clicking **Add** opens the object wizard for that object type. For details see:

- [Primary Key Wizard \(SQL Server\)](#)
- [Unique Key Wizard \(SQL Server\)](#)
- [Foreign Key Wizard \(SQL Server\)](#)
- [Create Synonym](#)

## TRIGGER WIZARD (SQL SERVER)

Triggers are a special type of procedure that automatically fire when defined data modification operations (insert, update, or delete) occur on a target table or view. Triggers fire after an insert, update or delete, but belong to the same transaction as the data modification operation. Triggers can be implemented to enforce business rules or referential data integrity.

**Important Notes**

- For more information on the syntax for Trigger bodies, consult the Microsoft SQL Server Transact-SQL Documentation.

**To create a new trigger using a wizard:**

- 1 [Open an Object Wizard](#) for a trigger.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Triggers \(SQL Server\) - Properties](#).
  - **Definition** panel - for details, see [Triggers \(SQL Server\) - Definition](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**NOTE:** You can use the Trigger Editor opens to create dependencies or alter the trigger statement.



## TRIGGERS (SQL SERVER) - PROPERTIES

When creating or editing a trigger, this tab/panel lets you work with the following settings:

Setting	Description
<b>Parent Type</b>	Select the TABLE or VIEW on which the trigger is to be created.
<b>Parent Schema</b>	Select the owner of the table or view on which the trigger is to be created.
<b>Parent Name</b>	Select the specific table or view in which the trigger is to be created.
<b>Name</b>	Provide a name for the trigger.
<b>Trigger Timing</b>	<p><b>INSTEAD OF:</b> This is the only option for a View trigger. An INSTEAD OF trigger fires in place of the triggering statement and will not make changes to the data unless the conditions of the INSTEAD OF statement are met first. So, your UPDATE execution statement is replaced by an INSTEAD OF UPDATE statement as a way to enforce particular business rules you establish.</p> <p><b>AFTER:</b> An AFTER trigger fires following the successful completion of the triggering action. So, for example, the trigger would fire after an UPDATE statement has executed and after constraints have been checked and verified.</p>
<b>Fire On Insert</b>	An INSERT trigger must be associated with an INSERT statement. For example, if a data load operation doesn't include an INSERT statement, the trigger won't be invoked.
<b>Fire On Update</b>	An UPDATE trigger can be associated with specific columns of the base table and will only be activated if those columns are updated.
<b>Fire On Delete</b>	A <b>DELETE</b> trigger is associated with a DELETE operation.
<b>Encrypted</b>	If you choose to encrypt the trigger, the trigger can't be published as part of SQL Server replication.

- 4 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## TRIGGERS (SQL SERVER) - DEFINITION

Complete the CREATE TRIGGER outline provided by typing or pasting the body of the trigger. Finally, use the **Schedule** or **Execute** buttons to [Create the New Object](#).

## UNIQUE KEY WIZARD (SQL SERVER)

Unique keys can enforce logical keys that are not chosen as the primary key. In other words, you can use a unique key to ensure no duplicate values are entered in specific columns that are not a part of the primary key. Although you can only attach one primary key to a table, you can attach multiple unique keys. Also, you can use unique keys on columns that allow null values.

**To create a new unique key using a wizard:**

- 1 [Open an Object Wizard](#) for a unique key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Unique Keys \(SQL Server\) - Properties](#).
  - **Columns** panel - [Unique Keys \(SQL Server\) - Columns](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**UNIQUE KEYS (SQL SERVER) - PROPERTIES**

When creating or editing a unique key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Owner</b> and <b>Table Name</b>	Choose the owner and name of the table in which the unique key is being created.
<b>Name</b>	Provide a name of the unique key being created.
Build Online	Enabling this feature specifies that the ONLINE=ON clause is used when creating this object and can subsequently be used when rebuilding or dropping this object.
Max degree of parallelism	Lets you specify a MAXDOP index operation value, limiting the number of processors used in parallel plan execution.
<b>Clustered</b>	Enable or disable clustering. With this check box enabled, subsequent rebuild and drop operations offer an online option.
<b>File Group</b>	If you do not specify a filegroup, Microsoft SQL Server creates the unique key in the default filegroup.
<b>Fill Factor</b>	This specifies how full each index page that's storing data should be. The fill factor is a percentage value between 0 and 100.

**UNIQUE KEYS (SQL SERVER) - COLUMNS**

From the **Column** dropdown, select a column for the primary key and specify a **Sort** option. To add more columns, click the New button and then follow the steps in the last instruction. Use the Delete button to drop columns.

**USER MESSAGE WIZARD (SQL SERVER)**

The User Message wizard lets you create multiple-language versions of user messages, such as errors and warnings, associated with a single message number. Key properties include the severity level and whether the message is automatically written to the NT Event Log.

**NOTE:** The user messages node only displays under the master database.

**To create a new user message using a wizard:**

- 1 Open a creation wizard for a user message. For details, see [Open an Object Wizard](#).
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [User Messages \(SQL Server\) - Properties](#).
  - **Information** panel - for details, see [User Messages \(SQL Server\) - Information](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object.

**USER MESSAGES (SQL SERVER) - PROPERTIES**

When creating or editing a user message, this tab/panel lets you work with the following settings:

Setting	Description
<b>Message Number</b>	Lets you specify a message number. The number must be greater than 50,000. The wizard automatically loads a value one higher than the currently highest used message number, but you can override the value.
<b>Severity</b>	Lets you select a predefined severity level between 001 and 025. The levels and their meanings are indicated in the User Message editor.
<b>Write to NT Event Log</b>	Lets you specify that the message is always written to the Windows NT Event Log.

**USER MESSAGES (SQL SERVER) - INFORMATION**

Lets you create the different language versions of the text for the message. The first version must be created in us\_english. Click the **Add new text for the user message** button, and in the dialog box that opens, select a **Language** of us\_english and provide the **Message Text**.

You can subsequently use the same process to create each different language version of the us\_english message.

**NOTE:** You cannot create two versions of a message for the same language.

This tab/panel also lets you edit and delete messages.

**USER WIZARD (SQL SERVER)**

The User Wizard lets you create a user who will then have access to the database where you are registering him or her. You can also identify the appropriate user group and the system privileges you want to assign to the new user.

**To create a new user using a wizard:**

- 1 [Open an Object Wizard](#) for a user
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Users \(SQL Server\) - Properties](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**USERS (SQL SERVER) - PROPERTIES**

When creating or editing a user, this tab/panel lets you work with the following settings:

Setting	Description
<b>Login Name</b>	Select the server login associated with this user.
<b>Name</b>	Provide the user name.

**USER DATATYPE WIZARD (SQL SERVER)**

User datatypes promote domain consistency by streamlining the definition of commonly used table columns in a database. You can build a customized datatype from system datatypes and bind defaults and rules to it to enhance integrity. When you reference the user datatype in a column, the column assumes all of the properties of the user datatype.

**To create a new user datatype using a wizard:**

- 1 [Open an Object Wizard](#) for a user datatype.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [User Datatypes \(SQL Server\) - Properties](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**USER DATATYPES (SQL SERVER) - PROPERTIES**

When creating or editing a user datatype, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the user datatype.
<b>Datatype</b>	Provide a name for the datatype.

Setting	Description
<b>Type</b>	Select the base datatype.
<b>Size</b>	Provide the size of the datatype.
<b>Allow Nulls</b>	Null has no explicitly assigned value. Null is not equivalent to zero or blank. A value of null is not considered to be greater than, less than, or equivalent to any other value, including another value of null.
<b>Default Binding</b>	Defaults promote data integrity by supplying a default value to a column if the user does not explicitly provide one. They are reusable objects that you can bind to user datatypes.
<b>Rule Binding</b>	Rules promote data integrity by allowing you to validate the values supplied to a column. They are reusable objects that you can bind to user datatypes.

## VIEW WIZARD (SQL SERVER)

Views are SQL queries stored in the system catalog that customize the display of data contained in one or more tables. Views behave like tables because you can query views and perform data manipulation operations on them. However, views do not actually store any data. Instead, they depend on data contained in their base tables.

### To create a new view using a wizard:

- 1 [Open an Object Wizard](#) for a view.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Views \(SQL Server\) - Properties](#).
  - **Definition** panel - for details, see [Views \(SQL Server\) - Definition](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## VIEWS (SQL SERVER) - PROPERTIES

When creating or editing a view, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the view. The owner of the view must have SELECT privileges for the tables in the CREATE view statement or DBADM authority on the database that contains the table.
<b>Name</b>	Provide a name for the view.
<b>Encryption</b>	If you select this option, the view will not be published with SQL Server replication.
<b>Schema Binding</b>	When you specify this option, the base table or tables cannot be modified in a way that would affect the view definition. The view definition itself must first be modified or dropped to remove dependencies on the table that is to be modified
<b>Check Condition</b>	When a row is modified through a view, this option makes sure the data remains visible through the view after the modification is committed.

Setting	Description
<b>Owner</b>	Select the owner of the view. The owner of the view must have SELECT privileges for the tables in the CREATE view statement or DBADM authority on the database that contains the table.

## VIEWS (SQL SERVER) - DEFINITION

Complete the CREATE VIEW statement by typing or pasting in the relevant query.

## ORACLE OBJECT WIZARDS

Rapid SQL lets you create Oracle objects using the following wizards:

- [Cluster Wizard \(Oracle\)](#)
- [Database Link Wizard \(Oracle\)](#)
- [Directory Wizard \(Oracle\)](#)
- [Foreign Key Wizard \(Oracle\)](#)
- [Function Wizard \(Oracle\)](#)
- [Index Wizard \(Oracle\)](#)
- [Library Wizard \(Oracle\)](#)
- [Materialized View Wizard \(Oracle\)](#)
- [Materialized View Log Wizard \(Oracle\)](#)
- [Outline Wizard \(Oracle\)](#)
- [Package Wizard \(Oracle\)](#)
- [Primary Key Wizard \(Oracle\)](#)
- [Procedure Wizard \(Oracle\)](#)
- [Profile Wizard \(Oracle\)](#)
- [Role Wizard \(Oracle\)](#)
- [Rollback Segment Wizard \(Oracle\)](#)
- [Sequence Wizard \(Oracle\)](#)
- [Synonym Wizard \(Oracle\)](#)
- [Table Wizard \(Oracle\)](#)
- [Tablespace Wizard \(Oracle\)](#)
- [Trigger Wizard \(Oracle\)](#)

- [Object Type Wizard \(Oracle\)](#)
- [Unique Key Wizard \(Oracle\)](#)
- [User Wizard \(Oracle\)](#)
- [View Wizard \(Oracle\)](#)

## CLUSTER WIZARD (ORACLE)

The Cluster Wizard lets you create a cluster. A cluster is a schema object that has one or more tables that all have one or more common columns. Rows of one or more tables that share the same value in these common columns are stored together in the database. The related columns of tables stored in a cluster are known as the cluster key.

### Important Notes

- To create a cluster, you need the CREATE CLUSTER or CREATE ANY CLUSTER system privilege.

#### TO OPEN THE CLUSTER WIZARD

- 1 On the Explorer, find the schema where you want to add the new cluster.
- 2 On the Cluster branch, right-click and select New.

OR

- 1 On the main toolbar, click Datasource and scroll to Objects
- 2 Click Clusters and then click New from the toolbar.

The tables below describe the fields you may encounter as you complete the wizard.

Required Information	Description
Who owns the cluster?	Pick an owner
What is the name of the cluster?	Type a unique name
On which tablespace do you want to create the cluster?	Self-explanatory
Add columns that are in this cluster These are the columns that are common between the tables you are "clustering"	<b>Add</b> or <b>Edit</b> Button - For more information, see <a href="#">Add or Modify Cluster Column..</a> <b>Drop</b> Button - Drops the column.
What is the size of this cluster?	This is the estimated number of bytes/KB/MB required by an average cluster key and its associated rows. Do not exceed the size of a data block.

Required Information	Description
What is the cluster type?	<p>Index: Rows having the same cluster key value are stored together. Each separate cluster key is stored only once in each data block. An indexed cluster is helpful if your clustered tables might grow unpredictably.</p> <p>Hash: Rows with the same hash key value are stored together. This is helpful if the tables are static.</p>
If this is a hash cluster, what is the number of hash keys?	Type the number of hash keys. Oracle will round the value up to the nearest prime number.
If this is a hash cluster, what is the hash function?	<p>Oracle uses a hash function to generate a distribution of numeric values, called hash values, which are based on specific cluster key values. The key of a hash cluster, like the key of an index cluster, can be a single column or composite key (multiple column key). To find or store a row in a hash cluster, Oracle applies the hash function to the row's cluster key value. The resulting hash value corresponds to a data block in the cluster, which Oracle then reads or writes on behalf of the issued statement.</p> <p>Default is the Oracle internal hash function, otherwise specify the hash expression you want to use.</p>
How many transaction entries are allowed for each data block in the cluster?	<p>Each transaction that updates a data block requires a transaction entry.</p> <p>Initial (1-255): The initial parameter ensures that a minimum number of concurrent transactions can update a data block, avoiding the overhead of allocating a transaction entry dynamically.</p> <p>Maximum (1-255): The maximum parameter limits concurrency on a data block.</p>
What is the percent of space reserved for future updates?	Percent Free (0-99): This sets the percentage of a data block to be reserved for possible row updates that are included in the block. The value you set is the percent kept free.
What is the minimum percentage of used space that Oracle maintains for each data block?	<p>The storage parameter lets you tune performance by minimizing the occurrence of row migration and chaining caused by update operations that extend the length of rows stored on the data block.</p> <p>Percent Used (1-99)</p>
How large are the cluster's extents?	<p>The unit of space allocated to an object whenever the object needs more space.</p> <p>Initial Extent - The initial space extent (in bytes) allocated to the object.</p> <p>Next Extent - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required.</p>
Specify the number of free lists	<p>Free lists are lists of data blocks that have space available for inserting rows. Identifying multiple free lists can reduce contention for free lists when concurrent inserts take place and potentially improve the performance of the cluster.</p> <p>Free Lists: The default and minimum value is 1; this option should be set higher if multiple processes access the same data block.</p>
Specify the number of free list groups (specify only if you are using the parallel server option)	This is the number of groups of free lists for the database objects being created.
Define a default buffer pool for this cluster	<p>Default - Select to retain the default.</p> <p>Keep - Select to retain the object in memory to avoid I/O conflicts.</p> <p>ORACLE 8i ONLY: Recycle - Select to rid data blocks from memory as soon as they are no longer in use.</p>



Required Information	Description
Oracle's parallel query option	The parallel server query option lets you process queries using many query server processes running against multiple CPUs, which provides substantial performance gains such as reducing the query completion time.
Choosing Cache	Cache: This keeps the data block in memory by placing it at the most recently used end. This option is useful for small lookup tables. No Cache

#### ADDING OR MODIFYING A CLUSTER COLUMN

The Add or Modify Cluster Column dialog box lets you manage cluster columns. You can open the dialog box in the Oracle Cluster Wizard. For details, see [Cluster Wizard \(Oracle\)](#).

The table below describes the options and functionality on the Add or Modify Cluster Column dialog box:

Option	Description
Column Name	Lets you type the column name.
Datatype	Lets you select the datatype for the cluster. If you select CHAR, RAW or VARCHAR2, in the Width box, type the width value. If you select NUMBER, in the Width box, type the width value and in the Scale box, type the scale value.

#### Completing the Add or Modify Cluster Column Dialog Box

To complete this dialog box, do the following:

- 1 In the *Add Cluster Column* dialog box, in the *Column Name* box, type the column name.
- 2 Click the *Datatype* list, click the datatype for the cluster.
  - If you clicked *CHAR*, *RAW* or *VARCHAR2*, in the *Width* box, type the width value.
  - If you clicked *NUMBER*, in the *Width* box, type the width value and in the *Scale* box, type the scale value.
- 3 Click the Add button.
- 4 To continue adding columns to the cluster, repeat steps 1-3.
- 5 When you finish adding columns, click Close.

Rapid SQL closes the Add Cluster Column dialog box.

For more information, see [Adding or modifying a cluster column](#).

## DATABASE LINK WIZARD (ORACLE)

A database link specifies a communication path from one database to another. If you're creating a link to a remote database, a database session is established in the remote database on behalf of the local application request. By creating either a public or private database link, you can determine which schema on the remote database the link will establish connections to by creating fixed, current, and connected database links. By creating a link you can reuse connectivity instructions each time you connect to the remote database.

### To Open the Database Link Wizard

- 1 On the Explorer, find the schema where you want to add the new database link.
- 2 On the Database Links branch, right-click and select New.

OR

- 1 On the main toolbar, click Datasource and scroll to Objects
- 2 Click Database Links and then click New from the toolbar.

The table that follows describes the fields you will encounter as you complete the wizard.

**NOTE:** To create a public database link, you need CREATE PUBLIC DATABASE LINK privileges.

Required Information	Description
What is the name of the database link?	Create a unique name
Should the database link be made public?	A public link is a link on a local database that's accessible to all users on that database.  If you keep the database link private by selecting No, a link is created in a specific schema of the local database and only the owner of the link can use it to access database objects in the corresponding remote database.
What is the name of the remote user?	Self-explanatory
What is the remote user's password?	Create a password for the remote user.
What is the connection string?	Self-explanatory.

## DIRECTORY WIZARD (ORACLE)

A directory object specifies an alias for a directory on the server file system where external binary file LOBs and external table data are located. The wizard completes a CREATE DIRECTORY statement from the information you supply. The Directory Wizard prompts you to name the directory and provide the full-qualified directory path.

### To Open the Directory Wizard

- 1 On the Explorer, find the Datasource where you want to create a directory and expand the Storage node.
- 2 Right-click Directories and select New.

The table that follows describes the fields you will encounter as you complete the wizard.

Required Information	Description
What is the name of the directory?	Type a meaningful name for the directory.
What is the directory path?	Type the full path name of the outside operating system directory that you want to alias in the directory (for example, /Video/Library/G_Rated). NOTE: Oracle trusts that directory you're specifying exists. The onus is on you to make sure it's valid and in the correct format (as required by your operating system).

After you have created the Directory, you can give other users Read or Write privileges by opening the new Directory on the Directories node, and making changes at the Privileges tab.

### FOREIGN KEY WIZARD (ORACLE)

A foreign key value in one table (child table) refers to a primary key value in another table (parent table). For example, the Author Name column in a publisher's database may be the primary key for a table of addresses that includes the Author Name column. If an author isn't included in the parent table, you can't add the address to the dependent address table. So foreign keys enforce referential integrity between tables by verifying the existence of foreign key values in the parent table before letting you insert or update foreign key values in the child table. In other words, a foreign key is an integrity constraint that requires each value in one table's column to match a value in a related table's data.

#### To create a new foreign key using a wizard:

- 1 [Open an Object Wizard](#) for a foreign key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Foreign Keys \(Oracle\) - Properties](#).
  - **Column Mapping** panel - [Foreign Keys \(Oracle\) - Column Mapping](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## FOREIGN KEYS (ORACLE) - PROPERTIES

When creating or editing a foreign key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Owner</b>	The owner of the table where the foreign key is being created.
<b>Table Name</b>	This is the table where the foreign key link originates--the child table.
<b>Name</b>	Lets you select a constraint name. System Generated Name - DB2 automatically generates a name. User Specified Constraint Name - You type the name.
<b>Enabled</b>	Enable or disable the foreign key. Enabled ensures that all data modifications to a given table (or tables) satisfy the conditions of the constraints. When disabled, the constraint is temporarily not operational.
<b>Delete Rule</b>	Select an action: <b>NO ACTION</b> - ensures that referenced values cannot be updated or deleted if to do so would violate referential integrity. <b>CASCADE</b> permits a referenced row in a child table to be deleted/updated if it is deleted/updated in the parent table. A row in the child table is SET NULL when rows in the parent table are deleted/updated.

## FOREIGN KEYS (ORACLE) - COLUMN MAPPING

Under **Referenced Table**, choose the **Owner** and then the **Name** of the referenced, or parent, table.

Under the **Main Table**, select checkboxes corresponding to the columns that are to reference columns in the referenced table. Then, under **Referenced Table**, select the corresponding column check boxes.

## FUNCTION WIZARD (ORACLE)

Functions are subroutines that you define and are useful for reusing application logic. You can use functions to determine the best methods for controlling access and manipulation of the underlying data contained in an object. A function returns a value, unlike a procedure, which does not.

**To create a new function using a wizard:**

- 1 [Open an Object Wizard](#) for a function.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Functions \(Oracle\) - Properties](#).
  - **Definition** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## FUNCTIONS (ORACLE) - PROPERTIES

When creating or editing a function, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the function.
<b>Name</b>	Provide a name for the function.

## FUNCTIONS (ORACLE) - DEFINITION

Complete the CREATE FUNCTION outline provided by typing or pasting the body of the function. Finally, use the **Execute** button to [Create the New Object](#).

## INDEX WIZARD (ORACLE)

Indexes are optional structures associated with tables. You can create indexes specifically to speed SQL statement execution on a table. When properly used, indexes are the primary means of reducing disk I/O. Indexes are logically and physically independent of the data in the associated table. Unique Indexes guarantee that no two rows of a table have duplicate values in the columns that define the index.

**NOTE:** The Index Wizard varies slightly in content based on the version of Oracle to which you are connected.

**NOTE:** To create indexes in your own schema, you need INDEX privileges on the target table. To create indexes in other schema, you need CREATE ANY INDEX privileges.

**NOTE:** For Oracle 8i or later, you can place a unique key constraint on an Index-Organized table.

**TIP:** Index-organized tables take up less storage space and quickly access table rows. Index-organized tables stores rows in primary key order reducing the amount of storage space needed.

**TIP:** An advantage of using index-organized tables is that the tables use less memory because key columns are not duplicated in the table and index. Rapid SQL stores the remaining non-key columns in the index structure.

**To create a new index using a wizard:**

- 1 [Open an Object Wizard](#) for an index.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Indexes \(Oracle\) - Properties](#).
  - **Columns** panel - [Indexes \(Oracle\) - Columns](#).
  - **Storage** panel - [Indexes \(Oracle\) - Storage](#).
  - **Partition** panel - [Indexes \(Oracle\) - Partition](#).
  - **Definition** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**INDEXES (ORACLE) - PROPERTIES**

When creating or editing an index, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Owner and Table Name</b>	Choose the owner and name of the table in which the index is being created.
<b>Owner and Name</b>	Choose the owner and provide the name of the index being created.
<b>Index Type</b>	<p>NONUNIQUE - In a non-unique index, the ROWID is treated as part of the key. Oracle treats a constraint as deferrable.</p> <p>UNIQUE - Select if the index is a unique constraint. The values in the indexed columns must be distinct.</p> <p>BITMAP - Widely used in data warehousing environments. The environments typically have large amounts of data and ad hoc queries, but a low level of concurrent DML transactions.</p>
<b>No Sort</b>	Enable this feature if the rows in the table already stored in ascending order. This increases the speed of the index creation process. Oracle does not sort the rows.
<b>Logging</b>	Enabling logs this operation to the redo file.
<b>Reverse</b>	Enabling this feature stores the bytes of the index block in reverse order and excludes the ROWID. The ROWID is a globally unique identifier for a row in a database. It is created at the time the row is inserted into a table, and destroyed when it is removed from a table.
<b>Function-Based</b>	Permits the results of known queries to be returned much more quickly. When you select this option, you are asked for the expression that governs the function-based index you are creating.
<b>No Parallel Execution</b>	<p>The parallel server query option lets you process queries, using many query server processes, running against multiple CPUs. This option provides substantial performance gains such as reduction of the query completion time.</p> <p>After creation, ALTER INDEX for NOPARALLEL execution - when you use multiple query servers and you select this option, the parallel query option remains in place, but parallel processing will be removed. If, for example, multiple users on numerous nodes are modifying the same small set of data, the cost of synchronization from the parallel processing may have an unnecessarily large drag on throughput.</p>

Setting	Description
<b>Parallel Degree</b>	The value you select indicates the number of query server processes that should be used in the operation.
<b>Parallel Instances</b>	The value you select indicates how you want the parallel query partitioned between the Parallel Servers.

### INDEXES (ORACLE) - COLUMNS

From the **Column** dropdown, select a column for the index and specify a **Sort** option. To add more columns, click the New button and then follow the steps in the last instruction. Use the Delete button to drop columns.

### INDEXES (ORACLE) - STORAGE

When creating or editing an index, this tab/panel lets you work with the following settings:

Setting	Description
<b>Data Block Storage</b> group	<p>Select the <b>DEFAULT Tablespace</b> only if you are creating a local partitioned index and want the partitions in the same tablespace as the partitions in the underlying table. (Each partition of a local index is associated with one partition of the table. Oracle can then keep the index partitions in synch with table partitions.)</p> <p>A transaction entry is needed for each INSERT, UPDATE, DELETE, etc. statement that accesses one or more rows in the block. Transaction entries in many operating systems require approx. 23 bytes.</p> <p><b>Percent Free</b> identifies how much space you want to allocate for new rows or updates to existing rows.</p> <p><b>Initial Transactions</b> ensures that a minimum number of concurrent transactions can update an index block, avoiding the overhead of allocating a transaction entry dynamically.</p> <p><b>Maximum Transactions</b> limits concurrency on an index block.</p>
<b>Extents</b> group	<p>An extent is the unit of space allocated to an object whenever the object needs more space.</p> <p><b>Initial Extent</b> - The initial space extent (in bytes) allocated to the object.</p> <p><b>Next Extent</b> - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required.</p> <p><b>Percentage Increase</b> - Lets you type the percentage.</p> <p>NOTE: You should be careful when setting Percent Increase because it magnifies how an object grows and, therefore, can materially affect available free space in a tablespace.</p> <p><b>Minimum Extents</b> - For a dictionary managed tablespace, this is the total number of extents to be allocated when the index is first created. For a locally managed tablespace, this is simply the initial amount of space allocated.</p> <p><b>Maximum Extents</b> - For a dictionary managed tablespace, this is the total number of extents that can ever be allocated to the index. In a locally managed tablespace, the database will automatically manage the extents.</p>

Setting	Description
<b>Freelists group</b>	<p><b>Free lists</b> let you manage the allocation of data blocks when concurrent processes are issued against the index. You can potentially improve the performance of the index by identifying multiple free lists, which can reduce contention for free lists when concurrent inserts take place.</p> <p>The default and minimum value is 1. You should increase this number if multiple processes access the same data block.</p> <p><b>Free List Groups</b> is the number of groups of free lists.</p> <p>NOTE: This option is only applicable for the parallel server option.</p>
<b>Buffer Pool</b>	<p>DEFAULT - Choose this if you want to use the default bufferpool.</p> <p>KEEP - Use this to retain the object in memory to avoid I/O conflicts. This type of bufferpool stores frequently referenced data blocks in a separate cache.</p> <p>RECYCLE - Select this option to save cache space by ridding data blocks from memory as soon as they are no longer in use.</p>

## INDEXES (ORACLE) - PARTITION

Clicking Create Partition opens the [Add, Insert, or Modify Partition](#) wizard.

## LIBRARY WIZARD (ORACLE)

Libraries are an object type introduced in Oracle8 that represent a call to an operating system shared library cache. After the call is made, libraries can be used by SQL or PL/SQL to link to external procedures or functions. Libraries are only to be used on operating systems that support shared libraries and dynamic linking. Libraries serve as pointers or aliases to physical operating system shared library files and do not exist as a physical object; rather they rely on the physical existence of the files in the external operating system library to which they refer. To access the function or procedures stored in the library, you need execute privileges at the operating system level where the shared library resides.

- To create a library in your own schema, you need CREATE ANY LIBRARY privileges. To use the functions or procedures stored in the library, you need object EXECUTE privileges on the library.

### To Open the Library Wizard

- 1 On the Explorer, find the datasource where you want to create a Library and expand the Schema node.
- 2 Right-click the Libraries node, and select New.

The table that follows describes the fields you may encounter using this wizard:

Required Information	Description
Who owns the library?	Self-explanatory.
What is the name of the library?	Self-explanatory.



Required Information	Description
What is the file specification?	Lets you type the file name and location. You must type the complete location (for example, D:\Embarcadero\ETLIB21D.DLL).

## MATERIALIZED VIEW WIZARD (ORACLE)

A materialized view gives you indirect access to table data by storing a query's results in a separate schema object. Essentially, a materialized view is a database object that contains the results of a query.

The Materialized View Wizard lets you:

- Specify the materialized view owner and to name the materialized view.
- Specify the materialized view's refresh configuration.
- Place the materialized view on a tablespace and specify the query that should be used to populate the materialized view.
- Specify how Oracle should allocate data blocks to store the materialized view.
- Specify how Oracle should manage the growth of the materialized view.
- Specify if Oracle updates the materialized view, register an existing table, and specify how to populate a materialized view.
- Specify if the data for the materialized view is cached, if you want the updates logged, and to specify a number of threads for a parallel operation.
- Specify rollback segments, and enable query rewrites.

### To Open the Materialized View Wizard

- 1 On the Explorer, find the datasource where you want to create a Materialized View and expand the Schema node.
- 2 Right-click the Materialized Views node, and select New.

The table that follows describes the fields you may encounter as you complete the wizard:

Required Information	Description
Who owns the materialized view?	Self-explanatory
What is the name of the materialized view?	Self-explanatory.

Required Information	Description
How should the materialized view be refreshed?	<p>Fast - Using the information logged on the materialized view logs or a partition maintenance operation, the refresh applies incremental changes. See <a href="#">Fast Refresh Requirements</a> for more information.</p> <p>Complete - This refresh recalculates the materialized view's defining query.</p> <p>Force - Applies fast refresh when feasible, otherwise uses a complete refresh.</p> <p>Never - Materialized view will not be refreshed with the optional refresh mechanisms.</p>
Choose a refresh mechanism	<p>On Demand - This option requires that all refreshes be manually executed.</p> <p>On Commit - Select to refresh the materialized view whenever Oracle processes a transaction. Only select this option for materialized views on single table aggregates and materialized views containing joins.</p> <p>Automatically - Select to refresh the materialized view automatically. In the On this date: boxes select a time and date, and then select a refresh amount and a unit of time.</p>
Where do you want to place the materialized view?	Select the tablespace where you want the materialized view placed.
What is the materialized view query?	Type the SQL query to be used to populate and to refresh the materialized view.
Select a refresh method	<p>Primary Key - A primary key's values uniquely identify the rows in a table. Changes are propagated according to row changes as identified by the primary key value of the row. Only one primary key can be defined for each table. The primary key of the master table is the basis for this refresh option, which is the default option.</p> <p>ROWID - A globally unique identifier for a row in a database based on the physical row identifiers. A RowID is created at the time the row is inserted into a table, and destroyed when it is removed from a table. ROWID materialized views cannot contain distinct or aggregate functions or GROUP BY subqueries, joins and set operations.</p>
How many transaction entries are allowed for each datablock in the materialized view?	<p>A transaction is a logical unit of work that contains one or more SQL statements. Each transaction that updates a data block requires a transaction entry.</p> <p>Initial (1-255) - Ensures that a minimum number of concurrent transactions can update a data block, avoiding the overhead of allocating a transaction entry dynamically.</p> <p>Maximum (1-255) - Limits concurrency on a data block.</p>
What is the percent of space reserved for future updates?	Percent Free (0-99) - This sets the percentage of a data block to be reserved for possible row updates that are included in the block. The value you set is the percent kept free.
What is the minimum percentage of used space that Oracle maintains for each datablock?	<p>Percent Used (0-99) - Set the amount of space to be used for each datablock.</p> <p>NOTE: The sum of percent free and the percent used cannot exceed 100.</p>
How large are the materialized views extents?	<p>The unit of space allocated to an object whenever the object needs more space. An extent is a specific number of contiguous data blocks set aside for storing a specific type of information.</p> <p>Initial Extent (KB) - The initial space extent (in bytes) allocated to the object.</p> <p>Next Extent - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required.</p>
How many extents should be allocated to the materialized view?	<p>Minimum Extents - The appropriate minimum extents value for the object.</p> <p>Maximum Extents - The appropriate maximum extents value for the object.</p>

Required Information	Description
What is the growth rate for sizing additional materialized views?	Percent Increase - Magnifies how an object grows and can materially affect available free space in a tablespace. Select a value in the corresponding box.
Can the materialized view be updated?	Yes/No
Do you want to register a prebuilt table to the view?	Yes/No. This option is particularly useful for registering large materialized views in a data warehousing environment.
Should the materialized view be immediately filled?	Yes/No: Select Yes if you want the materialized view populated immediately or during the next refresh operation.
Should data for the materialized view be cached?	Yes/No: Select if you want Oracle to put data you access frequently at the most recently used end of the list in the buffer cache when a full table scan is performed. This option is useful for small lookup tables.
Do you want updates to be logged?	Yes/No.
Do you want to specify the number of threads used in a parallel operation?	Parallel processes means that multiple processes work simultaneously to run a single statement. This can cut the amount of time it takes to get a response. Specify the degree of parallelism if you so desire. If you leave the default at 1, the operation will not be "parallelized."
Would you like to specify rollback segments to be used for the materialized view refresh?	A rollback segment temporarily stores old data that has changed in a SQL statement transaction until it is committed. The "before" image of the database, as it were. Local Rollback Segment - Default indicates that Oracle will select the rollback segment to use on the local machine. Master Rollback Segment - Specify the remote rollback segment used at the remote master site for the individual materialized view.
Is the materialized view eligible for query rewrite?	Select to enable the materialized view for query rewrite. Only enable query rewrite if expressions in the statement are repeatable.
Do you want to partition this materialized view?	Yes/No Partitioning methods available are: Range: Data is mapped to partitions based on ranges of column values. This is the default. Composite: Based on the range method of partitioning, you can create subpartitions within each partition. Hash: Data is distributed evenly over a specified number of partitions. Data need not fit into a logical range. List: You control explicitly how rows map to partitions. List partitions allow you to group and organize unrelated sets of data.
Do you want to enable Row Movement?	Yes/No Enabling row movement allows you to specify whether Oracle can move a table row when you are compressing a table or performing an update on partitioned data.
Select the partitioning columns	Self-explanatory
Select the subpartitioning method	Self-explanatory
Select the subpartitioning columns	Self-explanatory

Required Information	Description
Hash Partitioning methods	None Partition Definition: Specify number of partitions and (optionally) tablespaces Specify individual partitions by name and (optionally) tablespaces
Create list/ordered list of partitions	Self-explanatory The <b>Add Partition</b> dialog box may open. For details, see .
Specify number of subpartitions	Self-explanatory Click Add, Insert, or Edit to open a dialog box that lets you work with subpartition properties.
Select the default tablespaces to contain the subpartitions (optional)	Self-explanatory

### Fast Refresh Requirements

	When the Materialized View has:		
	Only Joins	Joins and Aggregates	Aggregate on a Single Table
Detail tables only	X	X	X
Single table only			X
Table Appears only once in the FROM list	X	X	X
No non-repeating expressions like SYSDATE and ROWNUM	X	X	X
No references to RAW or LONG RAW	X	X	X
No GROUP BY	X		
Rowids of all the detail tables must appear in the SELECT list of the query	X		
Expressions are allowed in the GROUP BY and SELECT clauses provided they are the same		X	X
Aggregates allowed but cannot be nested		X	X
AVG with COUNT		X	X
SUM with COUNT			X
	Only Joins	Joins and Aggregates	Aggregate on a Single Table
VARIANCE with COUNT and SUM		X	X

	When the Materialized View has:		
STDDEV with COUNT and SUM		X	X
WHERE clause includes join predicates which can be ANDed but not ORed.	X	X	
No WHERE clause			X
No HAVING or CONNECT BY	X	X	X
No subqueries, inline views, or set functions like UNION or MINUS	X	X	X
COUNT(*) must be present			X
No MIN and MAX allowed			X
If outer joins, then unique constraints must exist on the join columns of the inner join table	X		
Materialized View logs must exist and contain all columns referenced in the materialized view and have been created with the LOG NEW VALUES clause			X
Materialized View Logs must exist with rowids of all the detail tables	X		
Non-aggregate expression in SELECT and GROUP BY must be straight columns			X
DML to detail table	X		X
Direct path data load	X	X	X
ON COMMIT	X		X
ON DEMAND	X	X	X

## MATERIALIZED VIEW LOG WIZARD (ORACLE)

Materialized view logs are tables that maintain a history of modifications to the master table, and they are used to refresh simple materialized views. When you create a materialized view log, Oracle automatically creates a log table to track data changes in the master table and a log trigger to maintain the data in the log table. A log can refresh the materialized view incrementally and is therefore often less time-consuming than a complete refresh.

The Materialized View Log Wizard lets you:

- Specify the materialized view log owner and master table.
- Select refresh types and select column filters.
- Specify how Oracle should allocate data blocks to store the materialized view log.

- Specify how Oracle should manage the growth of the materialized view.
- Specify if you want the data for the materialized view log cached, if you want updates logged, and to enable parallel query.
- Specify if you want the log to hold new values.

### To Open the Materialized View Log Wizard

- 1 On the Explorer, find the datasource where you want to create a Materialized View Log and expand the Schema node.
- 2 Right-click the Materialized View Log node, and select New.

The table that follows describes the fields you may encounter as you complete the wizard:

Required Information	Description
Who owns the materialized view log's master table?	Self-explanatory
Which table will serve as the materialized view log's master table?	Self-explanatory
On which tablespace do you want to place the log?	Self-explanatory
Which refresh types would you like to use?	<p>Primary Key - The log records changes to the master table based on the primary key of affected rows. A primary key's values uniquely identify the rows in a table. Changes are propagated according to row changes as identified by the primary key value of the row. Only one primary key can be defined for each table. The primary key of the master table is the basis for this refresh option, which is the default option.</p> <p>ROWID - The log records changes to the master table based on the ROWID of the affected rows. A ROWID is a globally unique identifier for a row in a database based on the physical row identifiers. A ROWID is created at the time the row is inserted into a table, and destroyed when it is removed from a table. ROWID materialized views cannot contain distinct or aggregate functions or GROUP BY subqueries, joins and set operations.</p>
Optional: Select any filter column(s) to be recorded in the materialized view log.	A filter column is a column whose values you want to be recorded in the materialized view log for any rows that are changed. You can specify only one primary key, one ROWID, and one filter column list per materialized view log.
How many transaction entries are allowed for each data block in the materialized view log?	<p>A transaction is a logical unit of work that contains one or more SQL statements. Each transaction that updates a data block requires a transaction entry.</p> <p>Initial (1-255) - Ensures that a minimum number of concurrent transactions can update a data block, avoiding the overhead of allocating a transaction entry dynamically.</p> <p>Maximum (1-255) - Limits concurrency on a data block.</p>
What is the percent of space reserved for future updates?	Percent Free (0-99) - This sets the percentage of a data block to be reserved for possible row updates that are included in the block. The value you set is the percent kept free.

Required Information	Description
What is the minimum percentage of used space that Oracle maintains for each data block?	Percent Used (0-99) - Set the amount of space to be used for each datablock. NOTE: The sum of percent free and the percent used cannot exceed 100.
How large are the materialized view log's extents?	Initial Extent (KB) -The default is the value specified for the tablespace where the materialized view log resides.
Should the data for a materialized view log be cached?	Yes/No - Select Yes if you want Oracle to put data you access frequently at the most recently used end of the least recently used list in the buffer cache when a full table scan is performed. This option is useful for small lookup tables.  No indicates that your most frequently accessed data blocks are put at the least recently used end of the least recently used list of the buffer cache.
Do you want updates to be logged?	Yes/No
Do you want to enable parallel query for the log?	Degree. The integer is number of parallel threads used in the parallel operation.  The Parallel server query option lets you process queries using many query server processes running against multiple CPUs. This option provides substantial performance gains such as reduction of the query completion time.
Should the log hold new values?	Yes/No - Yes indicates both old and new values should be saved in the materialized view log.  No disables recording of new values in the log. This is the default.

## OUTLINE WIZARD (ORACLE)

Outlines are a set of results for the execution plan generation of a particular SQL statement. When you create an outline, plan stability examines the optimization results using the same data used to generate the execution plan. That is, Oracle uses the input to the execution plan to generate an outline, and not the execution plan itself.

**NOTE:** To create an outline, you must have CREATE ANY OUTLINE system privileges.

### To create a new outline using a wizard:

- 1 [Open an Object Wizard](#) for an outline.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Outlines \(Oracle\) - Properties](#).
  - **Definition** panel - [Outlines \(Oracle\) - Definition](#).
  - **DDL** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## OUTLINES (ORACLE) - PROPERTIES

When creating or editing an outline, this tab/panel lets you work with the following settings:

Setting	Description
<b>Stored Outline Name</b>	Provide a unique name for the outline.
<b>Category</b>	The category is a name you want to use to group stored outlines. (You can type over the word Default that appears automatically.)

## OUTLINES (ORACLE) - DEFINITION

Type the SQL statement you want to store as an outline.

**NOTE:** The only SQL statements possible with stored outlines are SELECT, DELETE, UPDATE, INSERT...SELECT, and CREATE TABLE...AS SELECT.

## PACKAGE WIZARD (ORACLE)

A package is a collection of related program objects stored together in the database. A package specification or package header, declares variables, constants, etc., that are visible outside the package's immediate scope. The package body defines objects declared by the specification but are not visible to applications outside the package. Packages contain all the information needed to process SQL statements from a single source file. You can use packages to process and call batches of SQL.

**To Open the Package Wizard**

- 1 On the Explorer, find the datasource where you want to create a package and expand the Schema node.
- 2 Right-click the Packages node, and select New.

You're asked to name an owner for the package and give it a name. When you click Finish, the Packages editor opens to the header tab where you indicate any package specifications and create the package body. For more information, see [Package Bodies Editor \(Oracle\)](#).

## PRIMARY KEY WIZARD (ORACLE)

Primary key (constraint)s are a set of table columns that can uniquely identify every row of a table. No fields that are a part of the primary key can have null values, and each table can have only one primary key.



**To create a new primary key using a wizard:**

- 1 [Open an Object Wizard](#) for a primary key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Primary Keys \(Oracle\) - Properties](#).
  - **Columns** panel - [Primary Keys \(Oracle\) - Columns](#).
  - **Storage** panel - [Primary Keys \(Oracle\) - Storage](#).
  - **Partition** panel - [Primary Keys \(Oracle\) - Partition](#).
  - **DDL** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**PRIMARY KEYS (ORACLE) - PROPERTIES**

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Owner and Table Name</b>	Choose the owner and name of the table in which the primary key is being created.
<b>Name</b>	Provide a name for the primary key being created.
<b>No Sort</b>	Enable this feature if the rows in the table already stored in ascending order. This increases the speed of the index creation process. Oracle does not sort the rows.
<b>Logging</b>	Enabling logs this operation to the redo file.
<b>Reverse</b>	Enabling this feature stores the bytes of the index block in reverse order and excludes the ROWID. The ROWID is a globally unique identifier for a row in a database. It is created at the time the row is inserted into a table, and destroyed when it is removed from a table.
<b>Validate</b>	Enabling this option indicates that existing data is checked against the constraint when the primary key is enabled. Leaving it disabled indicates that only new data is to be checked against the constraint.
<b>Deferrable</b>	Dictates whether constraint checking can be deferred until the end of a transaction.
<b>Deferred</b>	This option is enabled only if you enabled the <b>Deferrable</b> option. Select IMMEDIATE to have the constraint checked at the end of every DDL statement. Select DEFERRED to have the constraint checked only at the end of a transaction.
<b>Enabled</b>	Enables or disables the primary key.

**PRIMARY KEYS (ORACLE) - COLUMNS**

From the **Column** dropdown, select a column for the index and specify a **Sort** option. To add more columns, click the New button and then follow the steps in the last instruction. Use the Delete button to drop columns.

## PRIMARY KEYS (ORACLE) - STORAGE

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Data Block Storage</b> group	<p>Select the <b>DEFAULT Tablespace</b> only if you are creating a local partitioned index and want the partitions in the same tablespace as the partitions in the underlying table. (Each partition of a local index is associated with one partition of the table. Oracle can then keep the index partitions in synch with table partitions.)</p> <p>A transaction entry is needed for each INSERT, UPDATE, DELETE, etc. statement that accesses one or more rows in the block. Transaction entries in many operating systems require approx. 23 bytes.</p> <p><b>Percent Free</b> identifies how much space you want to allocate for new rows or updates to existing rows.</p> <p><b>Initial Transactions</b> ensures that a minimum number of concurrent transactions can update an index block, avoiding the overhead of allocating a transaction entry dynamically.</p> <p><b>Maximum Transactions</b> limits concurrency on an index block.</p>
<b>Extents</b> group	<p>An extent is the unit of space allocated to an object whenever the object needs more space.</p> <p><b>Initial Extent</b> - The initial space extent (in bytes) allocated to the object.</p> <p><b>Next Extent</b> - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required.</p> <p><b>Percentage Increase</b> - Lets you type the percentage.</p> <p>NOTE: You should be careful when setting Percent Increase because it magnifies how an object grows and, therefore, can materially affect available free space in a tablespace.</p> <p><b>Minimum Extents</b> - For a dictionary managed tablespace, this is the total number of extents to be allocated when the index is first created. For a locally managed tablespace, this is simply the initial amount of space allocated.</p> <p><b>Maximum Extents</b> - For a dictionary managed tablespace, this is the total number of extents that can ever be allocated to the index. In a locally managed tablespace, the database will automatically manage the extents.</p>
<b>Freelists</b> group	<p><b>Free lists</b> let you manage the allocation of data blocks when concurrent processes are issued against the index. You can potentially improve the performance of the index by identifying multiple free lists, which can reduce contention for free lists when concurrent inserts take place.</p> <p>The default and minimum value is 1. You should increase this number if multiple processes access the same data block.</p> <p><b>Free List Groups</b> is the number of groups of free lists.</p> <p>NOTE: This option is only applicable for the parallel server option.</p>
<b>Buffer Pool</b>	<p><b>DEFAULT</b> - Choose this if you want to use the default bufferpool.</p> <p><b>KEEP</b> - Use this to retain the object in memory to avoid I/O conflicts. This type of bufferpool stores frequently referenced data blocks in a separate cache.</p> <p><b>RECYCLE</b> - Select this option to save cache space by ridding data blocks from memory as soon as they are no longer in use.</p>

## PRIMARY KEYS (ORACLE) - PARTITION

Clicking **Create Partition** opens the [Add, Insert, or Modify Partition](#) wizard.

## PROCEDURE WIZARD (ORACLE)

Procedures are a reusable block of PL/SQL, stored in the database, that applications can call. Procedures streamline code development, debugging, and maintenance by being reusable. Procedures enhance database security by letting you write procedures granting users execution privileges to tables rather than letting them access tables directly. Procedures do not return values, unlike functions.

**NOTE:** To create a procedure in your own schema, you need CREATE PROCEDURE privileges. To create a procedure in someone else's schema, you need CREATE ANY PROCEDURE privileges.

### To create a new procedure using a wizard:

- 1 [Open an Object Wizard](#) for a procedure.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Procedures \(Oracle\) - Properties](#).
  - **Definition** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## PROCEDURES (ORACLE) - PROPERTIES

When creating or editing a procedure, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the procedure.
<b>Name</b>	Provide a name for the procedure

## PROCEDURES (ORACLE) - DEFINITION

Complete the CREATE PROCEDURE outline provided by typing or pasting the body of the procedure. Finally, use the **Execute** button to [Create the New Object](#).

## PROFILE WIZARD (ORACLE)

Profiles are a mechanism for allocating system and database resources to users. In essence, a profile is a set of limits with a name attached to them. If the profile is active, Oracle will limit use and resources to the limits defined by the profile.

The Profile Wizard lets you:

- Name the profile.
- Set composite limit.

- Set session limits for SGA shared pool.
- Set limits on total connection time per session and Idle time per session.
- Set limits on concurrent sessions per user, CPU time per session, and data blocks read per session.
- Set limits on CPU time per call, and number of data blocks read for a call to process an SQL statement.
- Set the number of failed login attempts, and the days an account locks.

**NOTE:** To create a profile, you need the CREATE PROFILE system privilege.

The Default option is subject to the limit specified for that particular resource. The default profile initially permits unlimited resources. Limits to the default profile can be made using an Alter statement.

The Unlimited option allows the user with that profile to take use unlimited amounts of that resource.

The Other options take on the resource limits that you indicate.

### To Open the Profile Wizard

- 1 On the Explorer, find the datasource where you want to create a procedure and expand the Security node.
- 2 Right-click the Profile node, and select New.

The table that follows describes the fields you may encounter as you complete the wizard.

Required Information	Description
What is the name of the profile?	Write a name that's 30 characters or less
What is the composite limit on resources per session?	You can set a single composite limit for all resource limits in a profile in addition to setting specific resource limits for a profile. Explicit and composite limits can peaceably coexist. The limit that is reached first stops the session's activity. Service units are a weighted sum of CPU per session, connect time, logical reads per session, and private SGA. Default/Unlimited/Other (service units)
What is the limit on the amount of private space a session can allocate in the shared pool of the SGA?	This limit only applies if you are using Shared Server architecture. Default/Unlimited/Other (KB)
What is the limit on the total connection time per session?	The total elapsed time limit for a session Default/Unlimited/Other (Minutes)
What is the limit on idle time per session?	Permitted periods of continuous inactive time during a session, expressed in minutes. Long-running queries and other operations are not subject to this limit. Default/Unlimited/Other (Minutes)

Required Information	Description
What is the limit on concurrent sessions per user?	Default/Unlimited/Other
What is the limit on CPU time per session?	Default/Unlimited/Other (Hundredths of a Second)
What is the limit on data blocks read per session?	Default/Unlimited/Other
What is the limit on CPU time per call?	The limit on a parse, execute, or fetch. Default/Unlimited/Other (Hundredths of a Second)
What is the limit on the number of data blocks read for a call to process a SQL statement?	This is the number of logical reads per call. Default/Unlimited/Other
How many failed login attempts will be allowed before an account is locked?	Default/Unlimited/Other
How long will an account be locked after the specified number of failed login attempts?	Default/Unlimited/Other (Days)
What is the lifetime of the password?	Default/Unlimited/Other (Days)
How many days must pass before a password can be reused?	Default/Unlimited/Other (Days)
How many password changes are required before the current password can be reused?	Default/Unlimited/Other
What is the grace period allowed for a password to be changed without expiring?	Default/Unlimited/Other (Days)
What is the name of the password complexity verification routine?	Default Null - Specifies no password verification is performed. Function Name

## ROLE WIZARD (ORACLE)

As you complete the Role Wizard, Rapid SQL constructs the necessary CREATE ROLE statement from the information that you have supplied. The Role Wizard lets you specify a name for the role and whether or not the role should be identified.

**NOTE:** To create a role, you need the CREATE ROLE system privilege.

### To Open the Role Wizard

- 1 On the Explorer, find the datasource where you want to create a role and expand the Security node.
- 2 Right-click the Roles node, and select New.

The table that follows describes the fields you may encounter as you complete the wizard.

Required Information	Description
What is the name of the role?	Self-explanatory.
How should the role be identified?	<p>Not identified - Selecting this means the role you are creating will be enabled immediately.</p> <p>Identified: If this is your choice, you're indicating one of the following authorization methods will be followed:</p> <p>Globally - Select to indicate that Oracle permits access to the user by obtaining user name and password information from the security domain central authority.</p> <p>Externally - Select to indicate that Oracle should verify the database user name against an existing operating system user name.</p> <p>Password - Select to indicate that Oracle should identify the role with the password you provide. In the Password box, type the password for the user.</p>

## ROLLBACK SEGMENT WIZARD (ORACLE)

Rollback segments manage all transactions in your Oracle databases. A transaction is a read, modify, and write cycle for an Oracle database. A rollback entry is made for all transactions unless a particular clause is specified. So, a rollback segment is a transaction log that consists of a pre-update image value and transaction status, among other things. The rollback segments maintain read consistency among concurrent users in a database and if the transaction fails for any reason, the old image is taken from the rollback segment. By maintaining a history of data changes, rollback segments can rollback uncommitted transactions so that data is rolled back to the prior state. SYS owns all rollback segments no matter who created them and are not accessible to users, just Oracle.

Oracle, it should be mentioned, strongly recommends that you use automatic undo management to simplify managing databases. Tuning rollback segments is a manual process that has largely been deprecated by Oracle, but it is supported for backward compatibility reasons.

The Rollback Segment Wizard lets you:

- Name the rollback segment and to place it online or off-line.
- Place the rollback segment on a tablespace.
- Specify the initial next and optimal extent size as well as the minimum and maximum number of extents that should be allocated to the rollback segment.

**NOTE:** This wizard is not available if auto-UNDO management is enabled.

**TIP:** Make sure enough rollback segments exist on a database to handle the imposed workload. One rule of thumb is to create one rollback segment for every four concurrent users.

## To Open the Rollback Segment Wizard

- 1 On the Explorer, find the datasource where you want to create a rollback segment and expand the Storage node.
- 2 Right-click the Roleback Segments node, and select New.

The table that follows describes the fields you may encounter as you complete the wizard.

Required Information	Description
What is the name of the rollback segment?	Self-explanatory.
Should this rollback segment be made public?	Yes - A public rollback segment can be brought online by any instance in a parallel server. Public rollback segments form a pool of rollback segments that can be used by any instance that needs one. No - This is the default. A private rollback segment can only be acquired by the instance specifying the segment in its initialization file.
Do you want to place the rollback segment to be online following its creation?	Online - To be useful, a rollback segment must be online. Offline - You may want to take rollback segments offline if you want to take a tablespace offline and it contains rollback segments that you want to keep from being used.
On which tablespace do you want to place this rollback segment?	Self-explanatory. Oracle suggests that you create one or more tablespaces specifically to hold all rollback segments. This way, the data contained in the rollback segments is held apart from other data types.
What extent sizes do you want to assign to this rollback segment?	Initial size (KB) Next size (KB) Optimal size (KB) Null/Default
What are the minimum and maximum number of extents to allocate to the rollback segment?	Minimum Maximum

## SEQUENCE WIZARD (ORACLE)

Sequences are programmable database objects that provide numbers in sequence for input to a table. A sequence can be used to automatically generate primary key values for tables. Once defined, a sequence can be made available to many users. When you create a sequence, you can define its initial value, increment interval, and maximum value.

The Sequence Wizard lets you:

- Specify the name and owner of the sequence.
- Set both the value of the sequence, and an interval and ranges for incrementing it.

- Cache the sequence, cycle the sequence when it reaches its minimum or maximum values, and guarantee that Oracle generates sequence numbers in the order of request.

**NOTE:** To create a sequence, it must belong to your schema or you need CREATE SEQUENCE privilege.

### To Open the Sequence Wizard

- 1 On the Explorer, find the datasource where you want to create a rollback segment and expand the Schema node.
- 2 Right-click the Sequences node, and select New.

The table that follows describes the fields you may encounter as you complete the wizard.

Required Information	Description
Who owns the sequence?	Self-explanatory.
What is the sequence name?	Self-explanatory.
What is the first sequence number to be generated?	Start with: Pick an integer.
What is the interval between sequence numbers?	Increment by: Positive numbers will generate ascending values, and a negative number will generate descending numbers
What is the sequence's minimum value?	The default is 1 for ascending sequences; This integer value can have 28 or fewer digits. None Minimum value - Identify how low the sequence can go.
What is the sequence's maximum value?	For descending values, the default is 1. Lets you specify the maximum value the sequence can generate. This integer value can have 28 or fewer digits. None Maximum value - Indicate the highest sequence value that will be allowed.
Should Oracle preallocate sequence numbers and cache them for faster access?	This is the number of sequence values you want to specify in the SGA buffers. This speeds access, but cached numbers are erased when the database is shut down. The default value is 20. Yes Number of values No
Should the sequence continue to generate values after reaching either its maximum or minimum value?	If you say no, the sequences will automatically recycle to the minimum value when you've hit the maximum for ascending sequences and vice versa for descending sequence values. Yes No



Required Information	Description
Should the sequence numbers be generated in the order of request?	This may be required when sequences are required for timestamping. But generally, because sequences are naturally ordered, this is only necessary for if you use Oracle RAC clusters for parallel mode.  Yes  No

## SNAPSHOT AND SNAPSHOT LOG WIZARDS (ORACLE)

Oracle has replaced the snapshot functionality with materialized views. Refer to the [Materialized View Wizard \(Oracle\)](#) and [Materialized View Log Wizard \(Oracle\)](#).

## SYNONYM WIZARD (ORACLE)

Synonyms are alternate names for database objects to be used as a reference by users or applications. A synonym offers you security and convenience so that you can refer to an object without revealing who owns it or what database it belongs to. Synonyms Depending on the platform, you can define synonyms on tables, views, sequences, procedures, functions, packages, and materialized views. If an underlying object needs to be renamed or moved, it's easy enough to redefine the synonym without having to modify any applications based on the synonym.

**NOTE:** To create a private synonym, you need CREATE SYNONYM privileges. To create a public synonym, you need CREATE PUBLIC SYNONYM privileges.

### To create a new synonym using a wizard:

- 1 [Open an Object Wizard](#) for a synonym.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - [Synonyms \(Oracle\) - Properties](#).
  - **Definition** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## SYNONYMS (ORACLE) - PROPERTIES

When creating or editing a synonym, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner and Name</b>	Choose the owner and provide a name for the synonym being created.
<b>Referenced Object Owner</b>	Select the owner of the object to which the synonym is to refer.

Setting	Description
<b>Referenced Object Type</b>	Select the type of the object to which the synonym is to refer.
<b>Referenced Object Name</b>	Select the specific object to which the synonym is to refer.
<b>Database Link</b>	If the object resides on a remote database, select a database link.

## TABLE WIZARD (ORACLE)

Tables are the most basic data storage units for Oracle. As you might expect, data is stored in rows and columns. By completing the Table Wizard, Rapid SQL constructs the necessary CREATE TABLE statement from the information that you supply. The Table Wizard varies slightly in content based on the version of Oracle on the target datasource. But in all cases, you name columns and determine column width or precision and scale depending on the column's data type. A row collects the column information that corresponds to a single record.

You can set rules for the columns to live by, and these are called integrity constraints. For example, if you select NOT NULL, that column will have to have a value in each row.

Also, before beginning, consider what kind of table you want to create as the wizard will ask you to choose:

Heap organized table	This is a basic table where data is stored as an unordered collection, i.e., heap.
Index-organized table	A B-tree index structure stores data, sorted by primary key. Nonkey column values are stored too.
Partitioned table	Data is broken down into smaller, more manageable pieces called partitions or subpartitions. Each partition can be managed individually and operate independent of the other partitions.
Clustered table	This is a table that, once created, is part of a cluster. A cluster is a group of tables that share some data blocks and columns and are often used together. To create a cluster, use the <a href="#">Cluster Wizard (Oracle)</a> .

**NOTE:** The table wizard panels differ depending on what options you select.

**NOTE:** To simplify the process of creating a table, the Table Wizard focuses on creating the basic table definition with a primary key constraint. After you create the basic table definition you can add unique and foreign keys to the table on the Constraints Tab of the Tables Editor.

**To create a new table using a wizard:**

- 1 [Open an Object Wizard](#) for a table.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Tables \(Oracle\) - Properties](#).
  - **Columns** panel - for details, see [Tables \(Oracle\) - Columns](#).
  - **Indexes** panel (not available with a **Row Organization** of **EXTERNAL**) - for details, see [Tables \(Oracle\) - Columns](#).
  - **Constraints** panel (not available with a **Row Organization** of **EXTERNAL**) - for details, see [Tables \(Oracle\) - Constraints](#).
  - **Storage** panel (not available with a **Row Organization** of **EXTERNAL**) - for details, see [Tables \(Oracle\) - Storage](#).
  - **IOT Properties** panel (not available with a **Row Organization** of **EXTERNAL**) - for details, see [Tables \(Oracle\) - IOT Properties](#).
  - **Partition** (not available with a **Row Organization** of **EXTERNAL**) - for details, see [Tables \(Oracle\) - Partition](#).
  - **Comment** panel - - for details, see [Add a Comment to an Object](#).
  - **Permissions** panel - for details, see [Set Permissions for an Object](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**TABLES (ORACLE) - PROPERTIES**

When creating or editing a table, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the table.
<b>Name</b>	Provide a name for the table
<b>Cache</b>	Enabling this feature keeps a block in memory by placing it at the most recently used end. This option is useful for small lookup tables.
<b>Row Movement</b>	Enabling this option permits the migration of a row to a new partition if its key is updated.
<b>Parallel Degree</b>	A value indicating the number of query server processes that should be used in the operation.
<b>Parallel Instances</b>	A value indicating how you want the parallel query partitioned between the Parallel Servers.
<b>Physical group</b>	Choose a Row Organization of INDEX, HEAP, or <b>EXTERNAL</b> . If you chose INDEX or HEAP, enable or disable <b>Logging</b> . If you chose EXTERNAL, provide an <b>External Type</b> , <b>Default Directory</b> , <b>Access Parameters</b> , <b>Reject Limit</b> , and <b>Location</b> .
<b>Logging</b>	Redo logs minimize the loss of data in the event that an uncontrolled shutdown happens.

## TABLES (ORACLE) - COLUMNS

For each column in the table, click the New button to create a column and provide a name for the column. Then, in the **Column Attributes** area, provide details for the column.

Use the Delete button to drop a selected column.

## Tables (Oracle) - Indexes

**NOTE:** This tab/panel is not available with a **Row Organization** of **EXTERNAL**.

Click **Add** to open the [Index Wizard \(Oracle\)](#).

## TABLES (ORACLE) - CONSTRAINTS

**NOTE:** This tab/panel is not available with a **Row Organization** of **EXTERNAL**.

Selecting a constraint type and clicking **Add** opens the object wizard for that object type. For details see:

- [Primary Key Wizard \(DB2 LUW\)](#)
- [Unique Key Wizard \(DB2 LUW\)](#)
- [Foreign Key Wizard \(IBM DB2 LUW\)](#)

## TABLES (ORACLE) - STORAGE

**NOTE:** This tab/panel is not available with a **Row Organization** of **EXTERNAL**.

When creating or editing a table, this tab/panel has the following settings:

Settings	Description
<b>Data Block Storage</b> group	<p>Select the DEFAULT <b>Tablespace</b> only if you want the partitions in the same tablespace as the partitions in the underlying table.</p> <p><b>Percent Free</b> identifies how much space you want to allocate for new rows or updates to existing rows.</p> <p><b>Initial Transactions</b> ensures that a minimum number of concurrent transactions can update a primary key block, avoiding the overhead of allocating a transaction entry dynamically.</p> <p><b>Maximum Transactions</b> limits concurrency on a primary key block.</p>

Settings	Description
<b>Extents</b> group	<p>An extent is the unit of space allocated to an object whenever the object needs more space.</p> <p><b>Initial Extent</b> - The initial space extent (in bytes) allocated to the object.</p> <p><b>Next Extent</b> - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required.</p> <p><b>Percentage Increase</b> - Lets you type the percentage.</p> <p>NOTE: You should be careful when setting Percent Increase because it magnifies how an object grows and, therefore, can materially affect available free space in a tablespace.</p> <p><b>Minimum Extents</b> - For a dictionary managed tablespace, this is the total number of extents to be allocated when the index is first created. For a locally managed tablespace, this is simply the initial amount of space allocated.</p> <p><b>Maximum Extents</b> - For a dictionary managed tablespace, this is the total number of extents that can ever be allocated to the index. In a locally managed tablespace, the database will automatically manage the extents.</p>
<b>Freelists</b> group	<p><b>Free lists</b> let you manage the allocation of data blocks when concurrent processes are issued against the primary key. You can potentially improve the performance of the primary key by identifying multiple free lists, which can reduce contention for free lists when concurrent inserts take place.</p> <p>The default and minimum value is 1. You should increase this number if multiple processes access the same data block.</p> <p><b>Free List Groups</b> is the number of groups of free lists.</p> <p>NOTE: This option is only applicable for the parallel server option.</p>
<b>Buffer Pool</b>	<p>DEFAULT - Choose this if you want to use the default bufferpool.</p> <p>KEEP - Use this to retain the object in memory to avoid I/O conflicts. This type of bufferpool stores frequently referenced data blocks in a separate cache.</p> <p>RECYCLE - Select this option to save cache space by ridding data blocks from memory as soon as they are no longer in use.</p>

## TABLES (ORACLE) - IOT PROPERTIES

**NOTE:** This tab/panel is not available with a **Row Organization** of **EXTERNAL**.

Provide compression and space details for an index-organized table.

## TABLES (ORACLE) - PARTITION

**NOTE:** This tab/panel is not available with a **Row Organization** of **EXTERNAL**.

Prior to working with partitions, you should be familiar with the material in [Rapid SQL and Oracle Partitioning](#).

Click **Create Partition** to [Partition a table](#).

## RAPID SQL AND ORACLE PARTITIONING

Partitioning your tables lets you get around the problem of supporting large tables. Partitioning lets you break large tables into smaller pieces, which are called partitions. Partitions make the data in your table easier to manage and analyze. Your SQL statements can access the partitions rather than the entire table. Partitions are most useful in data warehouse applications, which store large amounts of data.

The table below describes the types of partitions in Oracle:

Partition Type	Description
Range	Use range partitioning to map rows to partitions based on ranges of column values. This type of partitioning is useful when dealing with data that has logical ranges into which it can be distributed; for example, months of the year. Performance is best when the data evenly distributes across the range. If partitioning by range causes partitions to vary dramatically in size because of unequal distribution, you may want to consider one of the other methods of partitioning.
Hash	Use hash partitioning if your data does not easily lend itself to range partitioning, but you would like to partition for performance and manageability reasons. Hash partitioning provides a method of evenly distributing data across a specified number of partitions. Rows are mapped into partitions based on a hash value of the partitioning key. Creating and using hash partitions gives you a highly tunable method of data placement, because you can influence availability and performance by spreading these evenly sized partitions across I/O devices (striping).
Composite	In Oracle 8i, Oracle introduced both hash and composite partitioning. Hash partitions partition the table according to a hash function. Composite partitions use both range and hash types, first partitioning the data by a range of values, and then further dividing the partitions into subpartitions by way of a hash function. This option is not available for index-organized tables.
List	Use list partitioning when you require explicit control over how rows map to partitions. You can specify a list of discrete values for the partitioning column in the description for each partition. This is different from range partitioning, where a range of values is associated with a partition, and from hash partitioning, where the user has no control of the row to partition mapping.

## PARTITION A TABLE

The Add partition wizard lets you set up partitions for a table. Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:

Step	Settings and tasks
<b>Properties</b>	Select a <b>Partition Type</b> and optionally, a Subpartition Type.
Columns	For each column, click the New button and select a name from the Column dropdown. Use the Delete button to drop a selected column.

Step	Settings and tasks
<b>Subpartition Columns</b> (only available with a <b>Partition Type</b> of <b>RANGE</b> and <b>Subpartition Type</b> of <b>HASH</b> or <b>LIST</b> )	For each column, click the New button and select a name from the Column dropdown. Use the Delete button to drop a selected column.
<b>Subpartitions</b> (only available with a <b>Subpartition Type</b> of <b>HASH</b> )	Specify a <b>Default number of partitions</b> . For each partition, click the New button and then select a tablespace from the dropdown.
Range Definitions (only available with a <b>Partition Type</b> of <b>RANGE</b> )	Click the New button to <a href="#">Add a partition definition</a> .
Partition Definition (only available with a <b>Partition Type</b> of <b>HASH</b> )	To specify a partition method other than <b>None</b> , take one of the following actions: (1) Select the <b>Number Of Partitions</b> radio box, specify the <b>Number Of Partitions</b> , and for each partition, click the New button and choose a tablespace from the dropdown, or (2) select the <b>By Partition Name</b> radio box and for each partition, click the New button provide a name and then choose a tablespace from the dropdown.
List Definitions (only available with a <b>Partition Type</b> of <b>List</b> )	Click New to <a href="#">Add a partition definition</a> .

#### ADD A PARTITION DEFINITION

Use the following topics as a guide in completing the settings in this wizard:

Step	Settings and tasks	
<b>Partition Definition</b>	<b>Name</b>	Provide a name.
	<b>Tablespace</b>	Select a tablespace from the dropdown
	<b>Logging</b>	Enable or disable logging.
<b>Subpartitions</b>	To specify a partition method other than <b>None</b> , select the <b>By Subpartition Name</b> radio box and for each partition, click the New button to open a dialog that lets you provide subpartition values.  <b>NOTE:</b> When you split a range-list partition, you cannot specify the new partitions' subpartition information.	
<b>Storage</b>	Provide or select Data Block Storage, Extents, Freelists, and Buffer Pool values.	

#### TABLESPACE WIZARD (ORACLE)

Tablespaces are logical storage structures that act as partitions for the database. Each tablespace consists of one or more datafiles which are the physical structures holding the data. You can create a tablespace to store table data and other objects related to table performance such as indexes or large object data. Tablespaces are used to manage large complex databases. Once you have created a tablespace, you can place objects on it.

The Tablespace Wizard lets you:

- Name the tablespace, and specify space management.
- Specify what types of objects are stored on the tablespace, and place the tablespace online or offline.
- Add the datafiles that comprise the tablespace and specify the parameters for the datafiles.
- Specify how Oracle should manage the growth of the tablespace.

### Important Notes

- For auto-UNDO management to be in effect, set init.ora parameter to undo\_management. When set to MANUAL (the default), it disables auto-UNDO management. When set to AUTO, auto-UNDO management is enabled.
- To determine if the undo\_management parameter is set to AUTO, use the following query:

```
SELECT VALUE
FROM   SYS.V_$PARAMETER
WHERE  NAME = 'undo_management'
```

**NOTE:** This parameter cannot be set dynamically via the ALTER SYSTEM or ALTER SESSION.

For users using a version earlier than Oracle 8i and locally managed tablespaces, there are manual methods you can use to assist in the fight against tablespace fragmentation. They include:

- Setting PCTINCREASE to zero for all tablespaces and objects to promote same-sized extents.
- Specifying equal-sized allotments for your INITIAL and NEXT object storage parameters.
- Grouping objects with like growth and storage needs together in their own tablespaces.

One of the best ways to avoid fragmentation in a tablespace is to pre-allocate the space that your objects will use. If possible, plan for one to two years' growth for each object and allocate your space accordingly. Having initial empty objects will not affect table scan times as Oracle only scans up to the high-water mark (the last used block) in a table.

Of all your tablespaces, you want to avoid fragmentation problems in your SYSTEM tablespace the most as this is the major hotbed tablespace for Oracle activities. The easiest way to avoid this is to not allow any user (even the default DBA ID's SYS and SYSTEM) to have access to it. There are three ways to do this:

- Ensure no user has a DEFAULT or TEMPORARY tablespace assignment of SYSTEM.
- Ensure no user has a quota set for SYSTEM.
- Ensure no user has been granted the UNLIMITED TABLESPACE privilege.



## To Open the Tablespace Wizard

- 1 On the Explorer, find the datasource where you want to create a tablespace and expand the Storage node.
- 2 Right-click the Tablespace node, and select New.

The table that follows describes the fields you may encounter as you complete the wizard.

Required Information	Description
What is the name of the tablespace	Self-explanatory
What type of tablespace would you like to create?	<p>Permanent: Contains schema objects you want to keep on an ongoing basis. The objects are stored in datafiles.</p> <p>Temporary: Schema objects will last only as long as your session continues.</p> <p>Undo: This is a kind of permanent tablespace that holds undo data if your database is operating in automatic undo mode. Oracle recommends the automatic undo mode as the wiser choice than using rollback segments to undo.</p>
Do you want space management to be performed through the data dictionary or managed locally?	<p>Dictionary Managed: The dictionary tables in the SYS schema will track allocated and free extents. Oracle does not recommend this option because it can be slow.</p> <p>Locally Managed: Each tablespace manages its own free and used space so locally managed tablespaces are generally much more efficient.</p>
Will the extents be allocated by the system or do you want to choose a uniform extent size?	<p>Autoallocate extents: Users cannot specify an extent size because the tablespace is system managed. This option is not available for temporary tablespaces.</p> <p>Uniform extents: The default extent size is 1 megabyte. Extents for temporary tablespaces are uniform in size, but if you want to select a size other than the default, you can. Uniform is not an option for undo tablespaces.</p>
If you want to set a size for uniform extents, then enter it below, otherwise the default size will be used	The default extent size is 1 megabyte.
Should the tablespace be placed online following its creation?	<p>NOTE: This option is not available for temporary tablespaces.</p> <p>Online: This specifies the tablespace will be available to users immediately after it is created and while the database is up. This is the default.</p> <p>Offline: This is a tablespace that's unavailable to users even when the database is open. The datafiles in the tablespace are offline too.</p>
Do you want redo logging when Schema/Data is modified?	Self-explanatory. Logging is the default. The option is only available for permanent tablespaces.
Add Datafiles for the tablespace	See <a href="#">Adding or modifying a tablespace datafile</a>
What type of segment management should be used?	<p>This is an option available only to permanent, locally-managed tablespaces.</p> <p>Automatic (Oracle bitmaps): Free space segments will be managed by bitmap. Subsequent storage specifications for objects in the tablespace will be ignored if you make this selection (which Oracle recommends).</p> <p>Manual (Oracle free lists): The database will manage the segment free space using free lists.</p>

Required Information	Description
What block size should the tablespace use?	With this option, you can specify a nonstandard block size. The integer must correspond with the setting of one DB-nK_CACHE_SIZE parameter setting. This isn't available if you are creating a temporary tablespace.
Specify the default storage parameter for any objects created on the tablespace	Initial extent: The default value is the size of 5 data blocks. For manually managed space segments, the minimum value is 2 data blocks plus 1 data block for each free list group you specify.  Next extent: This is the size of the subsequent extent to be allocated after the initial extent is filled.  Minimum extent: The total number of extents to allocate at the object's creation. The default is 1, which is equivalent to the initial extent.
How many extents should be allocated to objects placed on the tablespace	Minimum extents: The total number of extents to allocate at the object's creation. The default is 1, which is equivalent to the initial extent.  Maximum extents: This is the total number of extents, including the first, that can be allocated to an object.
What should be the growth rate for sizing additional extents	Percent increase: The percent by which the third and following extents grow beyond the prior extent. Thus, if give a value of 20, the next extent will be 20% larger than the one that preceded it.

#### ADDING OR MODIFYING A TABLESPACE DATAFILE

**NOTE:** This functionality is available for all platforms.

The Add/Modify Datafile dialog box lets you add a datafile to a tablespace, or edit an existing one. You can open the dialog box in the Oracle Tablespace Wizard. For more information, see [Tablespace Wizard \(Oracle\)](#).

The table below describes the options and functionality on the Add/Modify Datafile dialog box:

Option	Description
Datafile Name	Type the name of the datafile in the box.
Datafile Size	Type the datafile size and then click the list to indicate MB or KB.
Datafile Status	Select ONLINE or OFFLINE availability for the datafile
Reuse Existing File?	The default is No.
Autoextend	Select this check box if you want to autoextend the datafile when more extents are required. This enables the options below.
Disk Space to Allocate to the Datafile When More Extents are Required	Type the disk space to allocate to the datafile when more extents are required by typing it and then click the list to indicate MB or KB.
Maximum Disk Space Allowed for Allocation to the Datafile	Specify the maximum disk space allowed for allocation to the datafile. Select Unlimited or Other. In the Other box, type the amount and then click the list to indicate MB or KB.

## Completing the Add/Modify Datafile Dialog Box

To complete this dialog box, do the following:

- 1 In the *Datafile Name* box, type the datafile name.
- 2 In the *Datafile Size* box, type the datafile size and then click the list to indicate MB or KB.
- 3 Reuse Existing File?
  - To reuse the existing file, click the Yes option button.
  - If you do not want to reuse the existing file, click the No option button.
- 4 Select the *Autoextend* check box.

Rapid SQL auto extends the datafile.

- In the *Disk Space to Allocate to the Datafile When More Extents are Required* box, type the disk space to allocate to the datafile when more extents are required by typing it and then click the list to indicate MB or KB.
- Specify the maximum disk space allowed for allocation to the datafile by clicking the *Unlimited* option button or the *Other* option button.

**NOTE:** If you click Other, in the Other box, type the amount and then click the list to indicate MB or KB.

- 5 Click OK.

Rapid SQL adds the datafile to the tablespace definition.

- 6 Click Close.

Rapid SQL closes the Add Datafile dialog box.

For more information, see:

[Adding or modifying a tablespace datafile](#)

## TRIGGER WIZARD (ORACLE)

A trigger is a special type of procedure that automatically fires when defined data modification operations (insert, update, or delete) occur on a designated table or view. Triggers fire after an insert, update or delete, but belong to the same transaction as the data modification operation. A stored procedure, on the other hand, is activated explicitly by the user or application (or trigger). Triggers should only be used for centralized, global operations that should be initiated by the triggering statement and should not be based on whether a particular user or database application issues the statement.

### Important Notes

- To create triggers in your own schema, you need CREATE TRIGGER privileges. To create triggers in other schemas, you need CREATE ANY TRIGGER privileges.

**To create a new trigger using a wizard:**

- 1 [Open an Object Wizard](#) for a trigger.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Triggers \(Oracle\) - Properties](#).
  - **Column Selection** panel - for details, see [Triggers \(Oracle\) - Column Selection](#).
  - **Action** panel - for details, see [Triggers \(Oracle\) - Action](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**TRIGGERS (ORACLE) - PROPERTIES**

When creating or editing a trigger, this tab/panel lets you work with the following settings:

Setting	Description
<b>Parent Type, Parent Owner, and Parent Name</b>	Select the type (TABLE or VIEW), owner and specific object for which the trigger is being created.
<b>Name</b>	Provide a name for the trigger being created.
<b>Enabled</b>	Enable or disable the trigger.
<b>Trigger Timing</b>	<p><b>BEFORE:</b> These triggers serve as extensions to the constraint subsystem and are most often used to validate input data, generate values for newly inserted rows, and read from other tables for cross-reference purposes. Note: Before triggers must be created as a For Each Row.</p> <p><b>AFTER:</b> Such a trigger is run after the integrity constraint validations; they can be used to modify operations in the database or be used for activities beyond the database, like supporting an alert notification.</p>
<b>Fire On Insert, Fire On Update, and Fire On Delete</b>	<p>Enable the events that fire the trigger.</p> <p>An INSERT trigger must be associated with an INSERT statement. For example, if a data load operation doesn't include an INSERT statement, the trigger won't be invoked.</p> <p>An UPDATE trigger can be associated with specific columns of the base table and will only be activated if those columns are updated.</p> <p>A <b>DELETE</b> trigger fires automatically after items in the table are deleted.</p>
<b>Trigger Type</b>	<p><b>STATEMENT:</b> (only fires once).</p> <p><b>ROW</b> (fires for each affected row): The trigger runs as many times as there are rows in the affected section. If the set of affected rows is empty, the trigger doesn't run.</p>
<b>Old Table Alias</b>	Type the name of a temporary table of rows as they exist before they're updated or deleted.
<b>New Table Alias</b>	Type a name for a temporary table of rows as they exist after they're inserted or updated.
<b>When Clause</b>	Type a WHEN clause or open a WHEN CLAUSE wizard to further qualify the trigger behavior.

### TRIGGERS (ORACLE) - COLUMN SELECTION

If you chose UPDATE as the **Trigger Event**, select the columns, select the check box beside each column that is to fire the trigger.

### TRIGGERS (ORACLE) - ACTION

Provide the body of trigger. If the SQL statement you have in mind for a trigger is longer than 60 lines of code, you would be better off creating a stored procedure.

INSTEAD OF statements can only be used by view triggers. BEFORE and AFTER options cannot be used for view triggers.

## OBJECT TYPE WIZARD (ORACLE)

Types define an abstract data type or object composed of a collection of similar types of data. For example, you can create an object type that defines a full address rather than the pieces of an address, such as city, state and postal code. An object type stores the pieces of an address in a single type, storing them in the same location, and allowing the full address to be accessed and manipulated as single unit rather than multiple units.

### To Open the Object Type Wizard

- 1 On the Explorer, find the datasource where you want to create a Type and expand the Schema node.
- 2 Right-click the Type node, and select New.

The single page wizard asks that you pick an owner for the type and name the type. The fun begins after you make your choices and click Finish. The Object Editor, where you can finalize the object type's creation, opens.

The Object Type Editor's tabs are:

**Header:** Here you flesh out the specification for the type you're creating including any methods (a subprogram), attributes, and any parameters or return types for a function.

**Body:** For every specification, the object type body defines the code for the method. If the object type header declares only attributes, not a method, the body is unnecessary.

You need to create the object type before you can see the following tabs. After you've added what you need to the Header and Body pages of the Editor, click Create on the toolbar.

**Information:** Read-only page that indicates the vital statistics for the object type.

**Dependencies:** A tree structure displays the objects that depend on the type you have created.

**Privileges:** Displays individual, group, and role permissions associated with the object type.

## UNIQUE KEY WIZARD (ORACLE)

A unique key constraint requires that every value in a column or set of columns be unique. Thus, no two rows of a table have duplicate values in the column or set of columns you identified. So, for example, you can use a unique key constraint to make sure you haven't duplicated a social security number in a list of employees.

### To create a new unique key using a wizard:

- 1 [Open an Object Wizard](#) for a unique key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Unique Keys \(Oracle\) - Properties](#).
  - **Columns** panel - for details, see [Unique Keys \(Oracle\) - Columns](#).
  - **Storage** panel - for details, see [Unique Keys \(Oracle\) - Storage](#).
  - **Partition** panel - for details, see [Unique Keys \(Oracle\) - Partition](#).
  - **DDL** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

### UNIQUE KEYS (ORACLE) - PROPERTIES

When creating or editing a unique key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Owner and Table Name</b>	Choose the owner and name of the table in which the unique key is being created.
<b>Name</b>	Provide a name for the unique key being created.
<b>No Sort</b>	Enable this feature if the rows in the table already stored in ascending order. This increases the speed of the index creation process. Oracle does not sort the rows.
<b>Logging</b>	Enabling logs this operation to the redo file.
<b>Reverse</b>	Enabling this feature stores the bytes of the index block in reverse order and excludes the ROWID. The ROWID is a globally unique identifier for a row in a database. It is created at the time the row is inserted into a table, and destroyed when it is removed from a table.
<b>Validate</b>	Enabling this option indicates that existing data is checked against the constraint when the primary key is enabled. Leaving it disabled indicates that only new data is to be checked against the constraint.
<b>Deferrable</b>	Dictates whether constraint checking can be deferred until the end of a transaction.
<b>Deferred</b>	This option is enabled only if you enabled the <b>Deferrable</b> option. Select IMMEDIATE to have the constraint checked at the end of every DDL statement. Select DEFERRED to have the constraint checked only at the end of a transaction.
<b>Enabled</b>	Enables or disables the primary key.

## UNIQUE KEYS (ORACLE) - COLUMNS

From the **Column** dropdown, select a column for the index and specify a **Sort** option. To add more columns, click the New button and then follow the steps in the last instruction. Use the Delete button to drop columns.

## UNIQUE KEYS (ORACLE) - STORAGE

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Data Block Storage</b> group	<p>Select the <b>DEFAULT Tablespace</b> only if you are creating a local partitioned index and want the partitions in the same tablespace as the partitions in the underlying table. (Each partition of a local index is associated with one partition of the table. Oracle can then keep the index partitions in synch with table partitions.)</p> <p>A transaction entry is needed for each INSERT, UPDATE, DELETE, etc. statement that accesses one or more rows in the block. Transaction entries in many operating systems require approx. 23 bytes.</p> <p><b>Percent Free</b> identifies how much space you want to allocate for new rows or updates to existing rows.</p> <p><b>Initial Transactions</b> ensures that a minimum number of concurrent transactions can update an index block, avoiding the overhead of allocating a transaction entry dynamically.</p> <p><b>Maximum Transactions</b> limits concurrency on an index block.</p>
<b>Extents</b> group	<p>An extent is the unit of space allocated to an object whenever the object needs more space.</p> <p><b>Initial Extent</b> - The initial space extent (in bytes) allocated to the object.</p> <p><b>Next Extent</b> - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required.</p> <p><b>Percentage Increase</b> - Lets you type the percentage.</p> <p>NOTE: You should be careful when setting Percent Increase because it magnifies how an object grows and, therefore, can materially affect available free space in a tablespace.</p> <p><b>Minimum Extents</b> - For a dictionary managed tablespace, this is the total number of extents to be allocated when the index is first created. For a locally managed tablespace, this is simply the initial amount of space allocated.</p> <p><b>Maximum Extents</b> - For a dictionary managed tablespace, this is the total number of extents that can ever be allocated to the index. In a locally managed tablespace, the database will automatically manage the extents.</p>
<b>Freelists</b> group	<p><b>Free lists</b> let you manage the allocation of data blocks when concurrent processes are issued against the index. You can potentially improve the performance of the index by identifying multiple free lists, which can reduce contention for free lists when concurrent inserts take place.</p> <p>The default and minimum value is 1. You should increase this number if multiple processes access the same data block.</p> <p><b>Free List Groups</b> is the number of groups of free lists.</p> <p>NOTE: This option is only applicable for the parallel server option.</p>
<b>Buffer Pool</b>	<p><b>DEFAULT</b> - Choose this if you want to use the default bufferpool.</p> <p><b>KEEP</b> - Use this to retain the object in memory to avoid I/O conflicts. This type of bufferpool stores frequently referenced data blocks in a separate cache.</p> <p><b>RECYCLE</b> - Select this option to save cache space by ridding data blocks from memory as soon as they are no longer in use.</p>

## UNIQUE KEYS (ORACLE) - PARTITION

Clicking Create Partition opens the [Add, Insert, or Modify Partition](#) wizard.

## USER WIZARD (ORACLE)

Whoever you add as a user will have access to the Oracle database. You also can set up the means by which the database recognizes the user.

## Important Notes

- To create a user, you need the CREATE USER system privilege.

## To create a new user using a wizard:

- 1 [Open an Object Wizard](#) for a user.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Users \(Oracle\) - Properties](#).
  - **Roles** panel - for details, see [Users \(Oracle\) - Roles](#).
  - **Quotas** panel - for details, see [Users \(Oracle\) - Quotas](#).
  - **DDL** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## USERS (ORACLE) - PROPERTIES

When creating or editing a user, this tab/panel lets you work with the following settings:

Setting	Description
<b>Name</b>	Provide a name for the user.
<b>Default Tablespace and Temporary Tablespace</b>	Select the default tablespace for objects the user creates and the tablespace to be used for the user's temporary segments.  If you do not specify a default tablespace when creating a user, the assigned tablespace will be that set using the Set Default function. For details, see .
<b>Profile</b>	The profile you choose affects the amount of database resources available to the user.
<b>Identified By</b>	REQUIRED_YES - A user identified by password is a local user.  REQUIRED_EXTERNAL - This is a user who is validated by an external service like an operating system or third-party service. The login authentication is handled by that external entity.  REQUIRED_GLOBAL - The user you're creating is a global user who is authenticated by the enterprise directory service.
<b>Password</b>	If you specified an <b>Identified By</b> value of REQUIRED_YES, provide a password for the user.



Setting	Description
<b>External Name</b>	If you specified an <b>Identified By</b> value of REQUIRED_EXTERNAL, provide an external name for the user.
<b>Account Locked</b>	When enabled, the account cannot be altered by anyone but the creator. It also means that after a specified number of failed attempts to access the database, the database will remain closed to the user for a period of time.
<b>Password Expired</b>	Marks the password as expired, forcing the user to change the password before being allowed to connect to the database.

## USERS (ORACLE) - ROLES

For each role to be assigned to the user, select the check box beside that role.

## USERS (ORACLE) - QUOTAS

To assign a tablespace quota for a user, select a tablespace, select the **Other** radio button, and provide the value in the **Quota** box.

## VIEW WIZARD (ORACLE)

Views are SQL queries stored in the system catalog that customize the display of data contained in one or more tables. Views behave like tables because you can query views and perform data manipulation operations on them. However, views do not actually store any data. Instead, they depend on data contained in their base tables. You can use views to help enforce corporate security policies by creating a view that limits information a user can see.

### Important Notes

- To create a view in your own schema, you need CREATE VIEW privileges. To create a view in someone else's schema, you need CREATE ANY VIEW privileges.

### To create a new view using a wizard:

- 1 [Open an Object Wizard](#) for a view.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Views \(Oracle\) - Properties](#).
  - **Definition** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## VIEWS (ORACLE) - PROPERTIES

When creating or editing a view, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the view. The owner of the view must have SELECT privileges for the tables in the CREATE view statement or DBADM authority on the database that contains the table.
<b>Name</b>	Provide a name for the view.

**Views (Oracle) - Definition**

This is where you build the query that will display the data you're interested in seeing. The template is CREATE VIEW <name> AS

SELECT: Identify the columns you want to show in the view

FROM: Identify the table(s) you want to draw the data from

WHERE: Write the equivalence you want to view.

You can use the Query Builder to help you write the appropriate SQL statement. For details, see [Build Query](#).

Finally, use the **Execute** button to [Create the New Object](#).

## SYBASE ASE OBJECT WIZARDS

- Rapid SQL lets you create Sybase objects using the following wizards:

- [Alias Wizard \(Sybase\)](#)
- [Database Wizard \(Sybase\)](#)
- [Default Wizard \(Sybase\)](#)
- [Extended Procedure Wizard \(Sybase\)](#)
- [Foreign Key Wizard \(Sybase\)](#)
- [Group Wizard \(Sybase\)](#)
- [Index Wizard \(Sybase\)](#)
- [Login Wizard \(Sybase\)](#)
- [Primary Key Wizard \(Sybase\)](#)
- [Procedure Wizard \(Sybase\)](#)
- [Rule Wizard \(Sybase\)](#)
- [Segment Wizard \(Sybase\)](#)
- [Table Wizard \(Sybase\)](#)

- [Trigger Wizard \(Sybase\)](#)
- [Unique Key Wizard \(Sybase\)](#)
- [User Datatype Wizard \(Sybase\)](#)
- [User Wizard \(Sybase\)](#)
- [View Wizard \(Sybase\)](#)

## ALIAS WIZARD (SYBASE)

The Alias Wizard lets you map a login to an existing user in the database. You can set up aliases so that multiple users log in to the same account and therefore have the same privileges. You can also set up an alias based on individual log ins and give those users access to the same alias with the advantage that you can track their activity within the database.

### To create a new alias using a wizard:

- 1 [Open an Object Wizard](#) for an alias.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Aliases \(Sybase\) - Properties](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## ALIASES (SYBASE) - PROPERTIES

When creating or editing an alias, this tab/panel lets you work with the following settings:

Setting	Description
<b>Name</b>	Provide a name for the alias being created.
<b>User</b>	Select the user to which the alias refers. The user has to have a valid account on SQL Server but cannot be a user in the current database.

## DATABASE WIZARD (SYBASE)

The Database Wizard lets you create a database (a structured collection of data that can be updated or queried) without knowing the underlying commands. Databases can be simple, that is one file with many records and the same fields, or much more complicated with multiple files with different fields.

**To create a new database using a wizard:**

- 1 [Open an Object Wizard](#) for a database.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Databases \(Sybase\) - Properties](#).
  - **Placement** panel - for details, see [Databases \(Sybase\) - Placement](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**DATABASES (SYBASE) - PROPERTIES**

When creating or editing a database, this tab/panel lets you work with the following settings:

Setting	Description
<b>Database Name</b>	Provide a name for the database.
<b>For Load</b>	This option speeds loading by eliminating the step for pre-initializing panels. The for load option is appropriate when creating a database for recovering from media failure or moving a database from one machine to another.
<b>With Override</b>	Enable this option to specify whether logs and data are to be kept on the same logical device
<b>Type</b>	Optionally, use the TEMP option to designate this as a temporary database.
<b>Options group</b>	Enable or disable the following Sybase database options: <b>abort tran on log full, allow nulls by default, auto identity, dbo use only, ddl in tran, identity in nonunique index, no chkpt on recovery, no free space acctg, read only, select into/bulkcopy/pllsort, single user, trunc log on chkpt</b> , and <b>unique auto_identity index</b> .

**DATABASES (SYBASE) - PLACEMENT**

In the **Fragment Properties** area provide or select values for the default fragment: **Device Name**, **Size** for the fragment, and a **Device Type** value of data only, log only, or data and log.

If necessary use the New button to add a new fragment and repeat the steps above to provide details for that fragment. Use the Delete button to drop a fragment.

**DEFAULT WIZARD (SYBASE)**

Here you create a default for table column or user-defined datatype in the event that no value is available when the data is inserted. The default value you specify will be inserted only in the current database. You can then bind the default to a specific column or user-datatype.

**To create a new default using a wizard:**

- 1 [Open an Object Wizard](#) for a default.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Defaults \(Sybase\) - Properties](#).
  - **Dependencies** panel - for details, see [Defaults \(Sybase\) - Dependencies](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**Defaults (Sybase) - Properties**

When creating or editing a default, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the schema that is to own the default.
<b>Name</b>	Provide a name for the default.
<b>Value</b>	Provide the value of the default.

**DEFAULTS (SYBASE) - DEPENDENCIES**

From the **Type** dropdown, choose Column or Datatype, and if you chose **Column**, choose a Table from the **Table** dropdown. The list on the left is populated with candidate columns or datatypes. To move a candidate from the list on the left to the dependencies column on the right, select the candidate and click **Add**. Remove columns or datatypes from the dependencies list on the right by selecting the column or datatype and clicking **Remove**.

**EXTENDED PROCEDURE WIZARD (SYBASE)**

Extended stored procedures provide a method for calling external procedural language functions from within the Adaptive Server. A procedural language is a language capable of calling a C language function or manipulating C language datatypes.

**To create a new extended procedure using a wizard:**

- 1 [Open an Object Wizard](#) for an procedure.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Extended Procedures \(Sybase\) - Properties](#).
  - **DDL View** panel - for details, see [Create the New Object](#).

- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

#### EXTENDED PROCEDURES (SYBASE) - PROPERTIES

When creating or editing an extended procedure, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the extended procedure.
<b>Name</b>	Provide the name of the procedure.
<b>Library Name</b>	Provide the name of the library containing the procedure.

#### FOREIGN KEY WIZARD (SYBASE)

Foreign keys are used to relate information from one table to another. Foreign keys are unique values that refer to specific columns of other tables. Thus, a foreign key links two tables together. Embarcadero Rapid SQL's Foreign Key Wizard makes it easy for you to create a relational link between two tables, thereby speeding queries and giving you faster access to data. The column in the initial table, the parent table, is called the primary key. The corresponding column in the (child) table that references the primary key, is the foreign key. Foreign keys can also refer to columns within the same table.

#### To create a new foreign key using a wizard:

- 1 [Open an Object Wizard](#) for a foreign key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Foreign Keys \(Sybase\) - Properties](#).
  - **Column Mapping** panel - for details, see [Foreign Keys \(Sybase\) - Column Mapping](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

#### FOREIGN KEYS (SYBASE) - PROPERTIES

When creating or editing a foreign key, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Owner</b>	Select the owner of the referring, or child, table
<b>Table Name</b>	Select the name of the referring, or child, table.
<b>Match Full</b>	Specify a referential integrity option.

## FOREIGN KEYS (SYBASE) - COLUMN MAPPING

Under **Referenced Table**, choose the **Owner** and then the **Name** of the referenced, or parent, table.

Under the **Main Table**, select checkboxes corresponding to the columns that are to reference columns in the referenced table. Then, under **Referenced Table**, select the corresponding column check boxes.

## GROUP WIZARD (SYBASE)

A group is a collection of privileges that the DBA assigns to designated users.

### To Open the Group Wizard

- 1 On the Explorer, find the datasource where you want to add a foreign key constraint and expand the Databases node.
- 2 Right-click Groups, and select New.

This single page wizard simply asks that you give the group a name. After you click Finish and Execute, the Group Editor opens where you assign users to the new group and give them particular system and object privileges.

## INDEX WIZARD (SYBASE)

Like an index in a book, a table index helps you get at the data you want without having to read through the whole table. Indexes can exist on single column or on multiple columns. Indexes appear in the form of B-trees. And, as for books, you can have multiple indexes for a single table.

### To create a new index using a wizard:

- 1 [Open an Object Wizard](#) for an index.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Indexes \(Sybase\) - Properties](#).
  - **Columns** panel - for details, see [Indexes \(Sybase\) - Columns](#).
  - **Partition** panel - for details, see [Indexes \(Sybase\) - Partition](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## INDEXES (SYBASE) - PROPERTIES

When creating or editing an index, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Owner, Table Name, and Name</b>	Select or provide identifying information.
<b>Attributes</b>	Provide or select values for the following Sybase options: <b>Index Type</b> , <b>Clustered</b> , <b>Ignore Duplicate Key</b> , <b>Ignore Duplicate Rows</b> , and <b>Maximum Rows Per Page</b> properties.
<b>Storage</b>	Provide or select values for the following Sybase options: <b>Reserve Page Gap</b> , <b>Segment Name</b> , <b>Fill Factor</b> , <b>Prefetch Strategy</b> , and <b>MRU Replacement Strategy</b> .

## INDEXES (SYBASE) - COLUMNS

Select a column from the **Columns** dropdown and specify a **Sort** option.

Use the New button to add more columns to the index. Use the Delete button to remove selected columns from the index.

## INDEXES (SYBASE) - PARTITION

Click **Add** to open the [Index Partition wizard \(Sybase\)](#).

## INDEX PARTITION WIZARD (SYBASE)

The Sybase Index Partition wizard can be opened from the following editors and wizards:

- [Index Wizard \(Sybase\)](#)
- [Indexes Editor \(Sybase\)](#)

Use the following topics as a guide to understanding and setting the options on this wizard:

Step	Settings and tasks	
<b>Properties</b>	Select a <b>Locality</b> .	
<b>Partition Definitions</b>	Click the Define a new partition button to open a dialog that lets you provide a name, select a segment, click the New button to add values, and then <b>Add</b> the partition definition.	

When finished, click the **Finish** button.

## LOGIN WIZARD (SYBASE)

Your login account controls access to the server and all of the databases within it. Only the System Administrator or System Security Officer can create logins. Once you can log into a server, you need additional privileges to access user databases. Specifically, each database owner adds the login as a user or alias to the database.



## To Open the Login Wizard

- 1 On the Explorer, open the database where you want to add a login.
- 2 Expand the Security node.
- 3 Right-click Logins and then click New.

The table that follows describes the fields you may encounter as you complete the wizard:

Required Information	Description
What is the ID for the login	Self-explanatory
What is the password for the login?	Self-explanatory
What is the full name of the login?	Optional (e.g., John Smith)
What should be the login's default database?	Self-explanatory. Be cautious about using Master to avoid the inadvertent creation/deletion of objects.
What is the login's default language?	If you do not select one, Rapid SQL uses the server's default language.
Select any roles that you wish to grant to the new login	Self-explanatory
Do you want to lock the new login after creating it?	Locking a login is an efficient way to disable a login on the master database without deleting it. Deleting a login can be complicated if the user owns any objects or has granted use of those objects to other users. Implementing roles is another reason to lock a login to prevent any single user from having all privileges and to be able to trace actions.
Select databases to create users	Select the check boxes that correspond to the databases or click the Check All button.  User Type - Lets you select the user type you want to add.  If you select Alias in the Alias box, click the User Name list, and then click the user name.  If you select dbo and you want to transfer aliases and permissions, click the Transfer Aliases and Permissions check box.  If you select User, in the User box, click the User Name list, and then click the user name, then click the Group list, and then click the user group.
Enter the password expiration, if any	Self-explanatory
Enter the minimum password length, if any	Self-explanatory
Enter the maximum number of login attempts, if any	Self-explanatory

## PRIMARY KEY WIZARD (SYBASE)

Primary key constraints make sure that no duplicate values or NULLS are entered in the columns you specify. You can use primary key constraints to enforce uniqueness and referential integrity. A table can only have a single primary key constraint.

### To create a new primary key using a wizard:

- 1 [Open an Object Wizard](#) for a primary key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Primary Keys \(Sybase\) - Properties](#).
  - **Columns** panel - for details, see [Primary Keys \(Sybase\) - Columns](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

### PRIMARY KEYS (SYBASE) - PROPERTIES

Provide or select the following:

- The **Table Owner** and **Table Name** of the table for which to create the index.
- A **Name** for the Index.
- Whether the index is **Clustered**.
- The **Maximum Rows Per Page**.
- A **Reserve Page Gap**.
- A **Segment Name**.
- A **Fill Factor** value.

### PRIMARY KEYS (SYBASE) - COLUMNS

Use the **Column** dropdown to choose a column name and choose a **Sort** option for that column. Use the New button to add an additional column to the index or use the Drop button to delete selected columns.

## PROCEDURE WIZARD (SYBASE)

Procedures are a reusable block of PL/SQL, stored in the database, that applications can call. Procedures streamline code development, debugging, and maintenance by being reusable. Procedures enhance database security by letting you write procedures granting users execution privileges to tables rather than letting them access tables directly.

**To create a new procedure using a wizard:**

- 1 [Open an Object Wizard](#) for a procedure.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Procedures \(Sybase\) - Properties](#).
  - **Definition** panel - for details, see [Procedures \(Sybase\) - Definition](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**PROCEDURES (SYBASE) - PROPERTIES**

When creating or editing a procedure, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the procedure.
<b>Name</b>	Provide the name of the procedure.
<b>Procedure Number</b>	Provide a procedure number if you want to group identically-named procedures.
<b>Recompile</b>	Optionally, enable this feature to prevent Sybase from saving a plan for the procedure.

**Procedures (Sybase) - Definition**

Complete the provided CREATE PROCEDURE statement by typing or pasting the body of the procedure.

**RULE WIZARD (SYBASE)**

Rules promote data integrity by allowing you to validate the values supplied to a table column. They are reusable objects that you can bind to table columns or user datatypes. Check constraints are similar to rules, and are in fact the preferred way of restricting data. A column or user-defined data type can have only one rule bound to it, but a column can have both a rule and one or more check constraints associated with it. Not that a rule cannot apply to data already existing in the database at the time you're creating the rule and can't be bound to a system-created data type. If you create a new rule when one already exists, the new rule will override the previous one.

**To create a new rule using a wizard:**

- 1 [Open an Object Wizard](#) for a rule.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Rules \(Sybase\) - Properties](#).
  - **Dependencies** panel - for details, see [Rules \(Sybase\) - Dependencies](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**RULES (SYBASE) - PROPERTIES**

When creating or editing a rule, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner and Name</b>	Select an <b>Owner</b> and provide a <b>Name</b> for the rule.
<b>Restriction</b>	Type the condition. The rule restriction is the condition that defines the rule and can be any expression valid in a WHERE clause and can include such elements as arithmetic operators, relational operators, and predicates (for example, IN, LIKE, BETWEEN).
<b>Type</b>	Choose STANDARD_RULE, AND_ACCESS_RULE, or OR_ACCESS_RULE.

**RULES (SYBASE) - DEPENDENCIES**

From the **Type** dropdown, choose Column or Datatype, and if you chose **Column**, choose a Table from the **Table** dropdown. The list on the left is populated with candidate columns or datatypes. To move a candidate from the list on the left to the dependencies column on the right, select the candidate and click **Add**. Remove columns or datatypes from the dependencies list on the right by selecting the column or datatype and clicking **Remove**.

**SEGMENT WIZARD (SYBASE)**

Segments allow you to control the placement of objects on database storage devices. A segment is a subset of the database device on which a specific database is stored. Each database can contain as many as 32 segments. Each database includes system, logsegment, and default segments. Others can be added using the Segment Wizard. By judiciously placing large tables and nonclustered indexes on segments on different devices or segments of specific sizes, you can improve I/O throughput or control space usage.

**To Open the Segment Wizard**

- 1 On the Explorer expand the database node where you want to add the Segment.
- 2 Right-click the Segments node and select New.

The table that follows describes the fields you may encounter as you complete the wizard.

Required Information	Description
What is the name of the segment?	Self-explanatory
On which devices do you wish to place the segment?	Self-explanatory

When you click **Execute**, Rapid SQL opens the **Segments Editor**. Here you can extend the segment and see the objects stored in the segment.

## TABLE WIZARD (SYBASE)

A table is a column-based arrangement of data in which the content of one column has a bearing on the other column(s). So, for example, a table might have a column for authors, another column for the books each author has written, and a third for the number of copies each title by a given author has sold. The data moves across the columns in rows.

**NOTE:** You must have CREATE TABLE permissions to generate a new table.

### To create a new table using a wizard:

- 1 [Open an Object Wizard](#) for a table.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Tables \(Sybase\) - Properties](#).
  - **Columns** panel - for details, see [Tables \(Sybase\) - Columns](#).
  - **Indexes** panel - for details, see [Tables \(Sybase\) - Indexes](#).
  - **Constraints** panel - for details, see [Tables \(Sybase\) - Constraints](#).
  - **Partitions** panel - for details, see [Tables \(Sybase\) - Partition](#).
  - **Permissions** panel - for details, see [Tables \(Sybase\) - Permissions](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

## TABLES (SYBASE) - PROPERTIES

When creating or editing a table, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the table.
<b>Name</b>	Provide a name for the table

Setting	Description
<b>Segment Name</b>	Specify the segment on which you want to place the table.
<b>Maximum Rows Per Page</b>	Specifying a number allows you to override the default. The default, 0, creates indexes with full pages and nonclustered indexes with full leaf pages. This number can be changed at any time.
<b>Reserve Page Gap</b>	
<b>Identity Gap</b>	Specify the size of the identity gap for the table.
<b>MRU Replacement Strategy</b>	When enabled, new pages are read into the least recent end of the page chain. When pages reach the most recent end of the chain, the pages are flushed.
<b>Prefetch Strategy</b>	Enabling this feature allows you to fetch as many as eight 2K data pages simultaneously instead of one at a time (the default).
<b>Lock Scheme</b>	Select a locking scheme of ALLPAGES, DATAPAGES, or DATAROWS.
<b>Expected Row Size</b>	If you specified a <b>Lock Scheme</b> of DATAROWS or DATAPAGES, provide an expected row size

### TABLES (SYBASE) - COLUMNS

For each column in the table, click the **Add Column** button to create a column, provide a **Name** for the column and provide or select the remaining column attributes.

Use the **Delete** button to drop a selected column.

Use the arrow buttons to reposition the columns.

**NOTE:** Because the smalldatetime datatype stores dates and time with less precision than the datetime datatype, before outputting you use the CAST or CONVERT functions to convert any boxes with the smalldatetime datatype to either VARCHAR or datetime datatypes.

### TABLES (SYBASE) - INDEXES

Click **Add** to open the [Index Wizard \(Sybase\)](#).

### TABLES (SYBASE) - CONSTRAINTS

Selecting a constraint type and clicking **Add** opens the object wizard for that object type. For details see:

- [Primary Key Wizard \(Sybase\)](#)
- [Unique Key Wizard \(Sybase\)](#)
- [Foreign Key Wizard \(Sybase\)](#)
- [Create Synonym](#)

### TABLES (SYBASE) - PARTITION

Click **Add** to open the [Table Partition wizard for Sybase ASE](#).

## TABLES (SYBASE) - PERMISSIONS

Set up the user permissions for this table.

## TABLE PARTITION WIZARD FOR SYBASE ASE

The Sybase Table Partition wizard can be opened from the following editors and wizards:

- [Table Wizard \(Sybase\)](#)
- [Tables Editor \(Sybase\)](#)

Use the following topics as a guide to understanding and setting the options on this wizard:

Step	Settings and tasks	
<b>Properties</b>	Select a <b>Partition Type</b> of ROUNDROBIN, RANGE, HASH, or LIST	
<b>Columns</b> (not available with a <b>Partition Type</b> of ROUNDROBIN)	For each column, click the Insert a new column button to create a column, provide a <b>Name</b> for the column and provide or select the remaining column attributes. Use the <b>Delete</b> button to drop a selected column. Use the arrow buttons to reposition the columns.	
<b>Partition Definition</b> (only available with a <b>Partition Type</b> of HASH or ROUNDROBIN)	Take one of the following actions: (1) Select the <b>Number Of Partitions</b> radio box, specify the <b>Number Of Partitions</b> , and for each partition, click the New button and choose a tablespace from the dropdown, or (2) select the <b>By Partition Name</b> radio box and for each partition, click the New button provide a name and then choose a tablespace from the dropdown.	
<b>Range Definitions</b> (only available with a <b>Partition Type</b> of RANGE)	Click the Define a new partition button to open a dialog that lets you provide a name, select a segment and <b>Add</b> the partition definition.	
<b>List Definitions</b> (only available with a <b>Partition Type</b> of LIST)	Click the Define a new partition button to open a dialog that lets you provide a name, select a segment, click the New button to add values, and then <b>Add</b> the partition definition.	

When finished, click the **Finish** button.

## TRIGGER WIZARD (SYBASE)

Triggers are a special type of procedure that automatically fire when defined data modification operations (insert, update, or delete) occur on a target table or view. Triggers fire after an insert, update or delete, but belong to the same transaction as the data modification operation. Triggers can be implemented to enforce business rules or referential data integrity.

## Important Notes

- For more information on the syntax for Trigger bodies, consult the Microsoft SQL Server Transact-SQL Documentation.

**To create a new trigger using a wizard:**

- 1 [Open an Object Wizard](#) for a trigger.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Triggers \(Sybase\) - Properties](#).
  - **Definition** panel - for details, see [Triggers \(Sybase\) - Definition](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**TRIGGERS (SYBASE) - PROPERTIES**

When creating or editing a trigger, this tab/panel lets you work with the following settings:

Setting	Description
<b>Table Owner, and Table Name</b>	Select the owner and specific table for which the trigger is being created.
<b>Name</b>	Provide a name for the trigger being created.
<b>Fire On Insert, Fire On Update, and Fire On Delete</b>	<p>Enable the events that fire the trigger.</p> <p>An INSERT trigger must be associated with an INSERT statement. For example, if a data load operation doesn't include an INSERT statement, the trigger won't be invoked.</p> <p>An UPDATE trigger can be associated with specific columns of the base table and will only be activated if those columns are updated.</p> <p>A <b>DELETE</b> trigger fires automatically after items in the table are deleted.</p>

**TRIGGERS (SYBASE) - DEFINITION**

Complete the CREATE TRIGGER statement by typing or pasting in content. Finally, use the **Execute** button to [Create the New Object](#).

**NOTE:** For more information on the syntax for Trigger bodies, consult the Sybase ASE Documentation.

**UNIQUE KEY WIZARD (SYBASE)**

Unique keys can enforce logical keys that are not chosen as the primary key. In other words, you can use a unique key to ensure no duplicate values are entered in specific columns that are not a part of the primary key. Although you can only attach one primary key to a table, you can attach multiple unique keys. Also, you can use unique keys on columns that allow null values.



**To create a new unique key using a wizard:**

- 1 [Open an Object Wizard](#) for a unique key.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Unique Keys \(Sybase\) - Properties](#).
  - **Columns** panel - for details, see [Unique Keys \(Sybase\) - Columns](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**Important Notes**

- If you are creating a non-clustered index constraint, you should place it on a separate segment from the target table.

## UNIQUE KEYS (SYBASE) - PROPERTIES

Provide or select the following:

- The **Table Owner** and **Table Name** of the table for which to create the index.
- A **Name** for the Index.
- Whether the index is **Clustered**.
- The **Maximum Rows Per Page**.
- A **Reserve Page Gap**.
- A **Segment Name**.
- A **Fill Factor** value.

## UNIQUE KEYS (SYBASE) - COLUMNS

Use the **Column** dropdown to choose a column name and choose a **Sort** option for that column. Use the New button to add an additional column to the index or use the Drop button to delete selected columns.

## USER DATATYPE WIZARD (SYBASE)

User datatypes promote domain consistency by streamlining the definition of commonly used table columns in a database. You can build a customized datatype from system datatypes and bind defaults and rules to it to enhance integrity. When you reference the user datatype in a column, the column assumes all of the properties of the user datatype.

**To create a new user datatype using a wizard:**

- 1 [Open an Object Wizard](#) for a user datatype.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [User Datatypes \(Sybase\) - Properties](#).
  - **DDL View** panel - for details, see [Create the New Object](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

**USER DATATYPES (SYBASE) - PROPERTIES**

When creating or editing a user datatype, this tab/panel lets you work with the following settings:

Setting	Description
<b>Datatype</b>	Provide a name for the datatype.
<b>Type</b>	Select the base datatype.
<b>Size</b>	Provide the size of the datatype.
<b>Allow Nulls</b>	Null has no explicitly assigned value. Null is not equivalent to zero or blank. A value of null is not considered to be greater than, less than, or equivalent to any other value, including another value of null.
<b>Identity</b>	This allows you to specify whether or not you want to limit access to the new datatype based on the user's privileges.
<b>Default Binding</b>	Defaults promote data integrity by supplying a default value to a column if the user does not explicitly provide one. They are reusable objects that you can bind to user datatypes.
<b>Rule Binding</b>	Rules promote data integrity by allowing you to validate the values supplied to a column. They are reusable objects that you can bind to user datatypes.

**USER WIZARD (SYBASE)**

The User Wizard lets you create a user who will then have access to the database where you are registering him or her. You can also identify the appropriate user group and the system privileges you want to assign to the new user.

**To create a new user using a wizard:**

- 1 [Open an Object Wizard](#) for a user
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Users \(Sybase\) - Properties](#).
  - **DDL View** panel - for details, see [Create the New Object](#).

- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

You can use the User Editor to assign System and Object privileges as well as create an alias.

#### USERS (SYBASE) - PROPERTIES

When creating or editing a user, this tab/panel lets you work with the following settings:

Setting	Description
<b>Login Name</b>	Choose a login ID from the drop-down list.
<b>Name</b>	Provide the user name.
<b>Group Name</b>	Specify a group for the user.

#### VIEW WIZARD (SYBASE)

Views are SQL queries stored in the system catalog that customize the display of data contained in one or more tables. Views behave like tables because you can query views and perform data manipulation operations on them. However, views do not actually store any data. Instead, they depend on data contained in their base tables. Views can only be

#### To create a new view using a wizard:

- 1 [Open an Object Wizard](#) for a view.
- 2 Use the following topics as a guide to setting properties and performing tasks as you pass through the wizard panels:
  - **Properties** panel - for details, see [Views \(Sybase\) - Properties](#).
  - **Definition** panel - for details, see [Views \(Sybase\) - Definition](#).
- 3 Finally, use the **Execute** button to create the object. For more information, see [Create the New Object](#).

You can use the User Editor to assign System and Object privileges as well as create an alias.

#### VIEWS (SYBASE) - PROPERTIES

When creating or editing a view, this tab/panel lets you work with the following settings:

Setting	Description
<b>Owner</b>	Select the owner of the view.
<b>Name</b>	Provide a name for the view.
<b>Check Type</b>	When enabled, when a row is modified through a view, this option makes sure the data remains visible through the view after the modification is committed.

### VIEWS (SYBASE) - DEFINITION

Complete the CREATE VIEW statement by typing or pasting in the relevant query. Finally, use the **Execute** button to [Create the New Object](#).

# FUNCTIONALITY

The table below lists functionality available in Rapid SQL.

Functionality	DB2 LUW	DB2 z/OS	InterBase/ Firebird	Oracle	SQL SVR	Sybase
<a href="#">Activate Logging</a>	X					
<a href="#">Add Database Fragment</a>	X	X		X	X	X
<a href="#">Add, Insert, or Modify Column</a>	X	X		X	X	X
<a href="#">Add, Insert, or Modify Parameter</a>	X	X		X	X	X
<a href="#">Add, Insert, or Modify Partition</a>	X	X		X	X	X
<a href="#">Add or Modify Cluster Column</a>				X		
<a href="#">Add or Modify Datafile</a>	X	X		X	X	X
<a href="#">Allocate Extent</a>				X	X	
<a href="#">Analyze</a>	X	X		X	X	
<a href="#">Bind Data Cache</a>						X
<a href="#">Bind Package</a>		X				
<a href="#">Bind Plan</a>		X				
<a href="#">Build Query</a>	X	X	X	X	X	X
<a href="#">Change Category</a>				X		
<a href="#">Change Container Size</a>	X	X		X	X	X
<a href="#">Change Database Comment</a>	X					
<a href="#">Change Password</a>				X	X	
<a href="#">Change Status (Tablespaces)</a>				X		
<a href="#">Change Status (Triggers)</a>						
<a href="#">Checkpoint</a>					X	X
<a href="#">Clone Partition</a>		X				
<a href="#">Clone Table</a>		X				
<a href="#">Coalesce</a>				X		
<a href="#">Compile</a>				X		
<a href="#">Connection Editor</a>	X	X		X	X	X
<a href="#">Convert to Partitioned</a>				X		
<a href="#">Copy Object Names</a>	X	X		X	X	X
<a href="#">Create Alias</a>	X	X				
<a href="#">Create Clone</a>		X				
<a href="#">Create Insert Statements</a>	X	X	X	X	X	X
<a href="#">Create or Edit Java Source</a>						
<a href="#">Create Like</a>	X	X	X	X	X	X

Functionality	DB2 LUW	DB2 z/OS	InterBase/ Firebird	Oracle	SQL SVR	Sybase
<a href="#">Create or Edit User Accounts</a>	X	X		X	X	X
<a href="#">Create or Modify Check Constraint</a>	X	X		X	X	X
<a href="#">Create Synonym</a>	X	X		X		
<a href="#">Create View</a>			X			
<a href="#">Creating objects</a>	X	X		X	X	X
<a href="#">DataLink Options</a>		X				
<a href="#">DBCC</a>					X	X
<a href="#">Deallocate Unused Space</a>				X		
<a href="#">Dependencies</a>	X	X	X	X	X	X
<a href="#">Detach/Attach</a>					X	
<a href="#">Disable/Enable Job Queues</a>				X		
<a href="#">Disable Index</a>			X		X	
<a href="#">Disable/Enable Triggers</a>				X	X	X
<a href="#">Disk Resize</a>						X
<a href="#">Drop</a>	X	X	X	X	X	X
<a href="#">Drop Clone</a>		X				
<a href="#">Edit Data</a>	X	X	X	X	X	X
<a href="#">Error</a>	X	X		X	X	X
<a href="#">Estimate Size</a>						X
<a href="#">Exchange Data With Clone</a>		X				
<a href="#">Execute</a>			X	X		X
<a href="#">Extract</a>	X	X	X	X	X	X
<a href="#">Flush Cache</a>	X					
<a href="#">Free Packages</a>		X				
<a href="#">Free Plans</a>		X				
<a href="#">Generate Anonymous Block</a>				X		
<a href="#">Generating Packages, Procedures, and Statements from Tables and Views</a>			X	X		
<a href="#">Impact Analysis</a>	X	X		X	X	X
<a href="#">Index Constraint</a>	X	X		X	X	X
<a href="#">Load Java</a>				X		
<a href="#">Lock</a>	X	X				
<a href="#">Move Log</a>					X	X
<a href="#">Open</a>	X	X		X	X	X
<a href="#">Partitioned Columns</a>				X		

Functionality	DB2 LUW	DB2 z/OS	InterBase/ Firebird	Oracle	SQL SVR	Sybase
<a href="#">Partition Upper Bound</a>				X		
<a href="#">Paste Columns for Check Constraints</a>	X	X		X	X	X
<a href="#">Place</a>					X	X
<a href="#">Preview: Create</a>	X	X		X	X	X
<a href="#">Preview</a>	X	X		X	X	X
<a href="#">Quiesce Tablespaces</a>	X	X				
<a href="#">Reassign by Category</a>				X		
<a href="#">Rebind Packages</a>	X	X				
<a href="#">Rebind Plans</a>	X					
<a href="#">Rebuild Index</a>			X	X	X	
<a href="#">Rebuild Outlines</a>				X		
<a href="#">Redistribute</a>	X					
<a href="#">Refresh Summary Table</a>	X					
<a href="#">Rename</a>	X	X	X	X	X	X
<a href="#">Reorganize</a>	X	X			X	X
<a href="#">Report</a>	X	X	X	X	X	X
<a href="#">Restart</a>	X					
<a href="#">Schema</a>	X	X	X	X	X	X
<a href="#">Select * From</a>	X	X	X	X	X	X
<a href="#">Set Default</a>				X		
<a href="#">Set Online/Offline</a>						X
<a href="#">Set Statistics</a>			X			
<a href="#">Set Tablespaces Quota</a>				X		
<a href="#">Set UNDO Tablespace</a>				X		
<a href="#">Shrink</a>				X	X	
<a href="#">Status</a>				X	X	X
<a href="#">Summary Definition</a>	X					
<a href="#">Switch Online</a>	X					
<a href="#">Transfer Ownership</a>						
<a href="#">Truncate</a>	X	X		X	X	X
<a href="#">Update Statistics</a>	X	X			X	X

## ACTIVATE LOGGING

**NOTE:** This functionality is available for IBM DB2 LUW for Linux, Unix, and Windows only.

The Activate Logging dialog box lets you store the most recent log records, when the table is not logging initially.

### Important Notes

None

For more information, see [Completing the Activate Logging Dialog Box](#).

## COMPLETING THE ACTIVATE LOGGING DIALOG BOX

To complete this dialog box, do the following:

## ADD/MODIFY DATABASE FILE

The table below describes the options and functionality on the Add/Modify Database File dialog box.

Option	Description
Database	Displays the database name.
Filegroup	Lets you select a filegroup to which the database belongs. Default - Sets filegroup as the default. Read Only - Sets the filegroup as read-only.
Database File	Name - Lets you name the database file. File Path - Lets you specify the file location. Type or browse and locate the file. The application automatically gives the database file a *.mdf extension. Size - Lets you specify a database size in KB or MB.
Allow Growth	Lets the database file grow. Growth Rate - Lets you specify the growth rate in percents, KB, or MB. Max Size - Specifies the maximum database file size. Select Unlimited to let the database file grow without restrictions. To specify a specific size, select Size and enter a size in KB or MB.

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Activate Logging.  
Rapid SQL opens the Activate Logging dialog box.
- 3 Do one of the following:



Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.

## ADD DATABASE FRAGMENT

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Add Database Fragment dialog box lets you add database fragments to a device.

### Important Notes

None

The table below describes the options and functionality on the Add Database Fragment dialog box.

Option	Description
Data Device	Lets you specify a data device. Select a device and set a device size in MBs.
Log Device	OPTIONAL: Lets you specify the placement of the transaction log.  NOTE: We strongly recommend that you place the transaction log on a separate device from all other database fragments.
Add Button	Lets you add more database fragments.
Edit Button	Lets you edit an existing database fragment.
Delete Button	Lets you delete an existing database fragment.

For more information, see:

[Completing the Add Database Fragment Dialog Box](#)

## COMPLETING THE ADD DATABASE FRAGMENT DIALOG BOX

To complete this dialog box, do the following:

- 1 In the *Data Device* box, click the *Device* list to specify the placement of the database.
- 2 In the *Size* box, type the *MB* value of the data device.

- 3 OPTIONAL: In the *Log Device* box, click the list to specify the placement of the transaction log.
- 4 If you specify a log device, in the *Size* box, type the *MB* value of the log device.
- 5 Click *Add*.

Rapid SQL adds the database fragment and closes the *Add Database File* dialog box.

- To add more database fragments, click the *Add* button.
- To edit an existing database fragment, click the fragment and click the *Edit* button.
- To delete an existing database fragment, click the fragment and click the *Delete* button.

For more information, see:

[Add Database Fragment](#)

## ADD, INSERT, OR MODIFY COLUMN

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Add, Insert, or Modify Columns dialog box lets you manage columns. You can open the dialog box in the following wizards and editors:

- [Indexes Editor \(IBM DB2 LUW\)](#)
- [Tables Editor \(IBM DB2 LUW\)](#)
- [Unique Keys Editor \(IBM DB2 LUW\)](#)
- [Views Editor \(IBM DB2 LUW\)](#)
- [Indexes Editor \(IBM DB2 Z/OS\)](#)
- [Tables Editor \(IBM DB2 Z/OS\)](#)
- [Unique Keys Editor \(IBM DB2 Z/OS\)](#)
- [Views Editor \(IBM DB2 Z/OS\)](#)
- [Tables Editor \(SQL Server\)](#)
- [Tables Editor \(Oracle\)](#)
- [Tables Editor \(Sybase\)](#)
- [Table Wizard \(DB2 LUW\)](#)
- [Table Wizard \(DB2 Z/OS\)](#)
- [Table Wizard \(SQL Server\)](#)
- [Table Wizard \(Oracle\)](#)

- [Table Wizard \(Sybase\)](#)

## Important Notes

None

The table below describes the options and functionality on the Add, Insert or Modify Column dialog box.

**NOTE:** The options differ by platform.

Option	Description
Column Name	Lets you enter the name of the column.
Datatype	Lets you select a datatype and set the width and scale.  NOTE: Oracle 9i supports Timestamp datatypes, including Timestamp, Timestamp with Time Zone, and Timestamp with Local Time Zone.
Allow Nulls?	Lets you allow columns with no text.
Default Value	Lets you accept the default value for the datatype.
Comment	OPTIONAL: Lets you type comments, which can be up to 254 characters long.
Lob Options	Lets you specify options for a large datatype.
Character Options	For a datatype that has character options. For Bit Data
Identity Column	Lets you specifies values for an identity column. The identity column will start at the value specified in Min Value, and will continue to progress until the Max value is reached.  Min Value - Lets you specify the minimum value. If you do not specify, Rapid SQL will use the max values for the datatype.  Max Value - Lets you specify the maximum value. If you do not specify, Rapid SQL will use the max values for the datatype.  Cycle - If you do not select, when the Max Value is reached the server will not allow anymore rows. If you do select, when the Max value is reached the server will start the identity column from the beginning.

**NOTE:** The options differ for [IBM DB2 LUW for OS/390 and z/OS](#), [Microsoft SQL Server 7 or later](#), [Oracle](#), and [Sybase ASE](#).

For more information, see:

[Completing the Add, Insert or Modify Column Dialog Box](#)

## COMPLETING THE ADD, INSERT OR MODIFY COLUMN DIALOG BOX

To complete this dialog box, do the following:

**NOTE:** The options differ for [IBM DB2 LUW for OS/390 and z/OS](#), [Microsoft SQL Server 7 or later](#), [Oracle](#), and [Sybase ASE](#).

- 1 In the *Column Name* box, type the target column name.
- 2 Click the *Datatype* list, and then click the target datatype.

**NOTE:** For SQL Server 2000, if you create a table with a Column datatype = text., you can set the storage and image values on the Storage Tab of the Tables Editor Storage Tab. When you have a text datatype, the Storage Tab displays a Text In Row box where you can specify the maximum size to be stored.

**NOTE:** Because the smalldatetime datatype stores dates and time with less precision than the datetime datatype, before outputting you use the CAST or CONVERT functions to convert any boxes with the smalldatetime datatype to either VARCHAR or datetime datatypes. For more information, see SQL Server Books Online, Transact-SQL Reference.

- 3 To allow nulls, click the *Yes* option button.
- 4 If you do not want to allow nulls, click the *No* option button.
- 5 To define a computed expression column, click the *Computed Expression* option button and then type the expression in the box.

**NOTE:** This option is only valid for Microsoft SQL Server 7.

- 6 To set row binding parameters, in the *Row Binding* grid.
  - To set a default value, click the *Default Value* list or type a value in the box.
  - To fill the rows with a default value, select the *Fill Rows with Default Value* check box.
  - To bind a default to the column (instead of declaring it), in the *Default Binding* list, click the default binding.
  - To bind a rule to the column, click the *Default Rule* list and click the default rule.
- 7 To set an identity column, select the *Identity Column* check box.
- 8 If you select the *Identity Column* check box, in the *Identity Seed* and *Identity Increment* boxes, type the appropriate values.
- 9 If the columns are not for replication, select the *Not for Replication* check box.

- 10 Click the *Add* button.

Rapid SQL adds the column.

- 11 When you finish adding columns, click the *Close* button.

Rapid SQL closes the Add Column dialog box.

- 12 To insert a column, click the Insert button.
- 13 To add, insert, edit or drop columns, click the corresponding button in *Add columns* that are belonging to this table.

For more information see:

[Add, Insert or Modify Column](#)

[Completing a Dialog Box](#)

## ADD, INSERT OR MODIFY COLUMN FOR IBM DB2 LUW FOR OS/390 AND Z/OS

The table below describes the options and functionality on the Add, Insert or Modify Column dialog box for IBM DB2 LUW for OS/390 and z/OS.

Options	Description
Column Name	Lets you enter the name of the column.
Datatype	Lets you select a datatype and set the width and scale.
Allow Nulls?	Lets you allow columns with no text.
Default Value	Lets you accept the default value for the datatype you selected.
Comment	OPTIONAL: Lets you type comments, which can be up to 254 characters long.
For Data	<p>Lets you specify the subtype If the datatype selected is a character string column such as CHAR, VARCHAR or LONG VARCHAR.</p> <p>SBCS - The column holds single-byte data.</p> <p>Mixed - The column holds mixed data. Do not set as MIXED if the installation box, MIXED DATA on the DSNTIPF panel is NO.</p> <p>Bit - The column holds BIT data. For CLOB datatypes, the default is SBCS when the installation box, MIXED DATA on the DSNTIPF panel is NO. The default for CLOB datatypes is Mixed when the installation box, MIXED DATA on the DSNTIPF panel is YES.</p>

For more information, see:

[Completing the Add, Insert or Modify Column Dialog Box](#)

## ADD, INSERT OR MODIFY COLUMN FOR MICROSOFT SQL SERVER

Microsoft SQL Server 7 introduced the concept of the computed expression column which lets you define a virtual column derived from an expression that computes a value based on other columns from the same table. For example, create a computed expression column that calculates the sales tax for an item based on the values stored in the price and tax rate columns of that table. Computed expression columns cannot be used as a key column.

Option for Microsoft SQL Server 7 or later	Description
Column Name	Lets you type the target column name.
Datatype	<p>Lets you select the target datatype, width and scale.</p> <p>NOTE: Because the smalldatetime datatype stores dates and time with less precision than the datetime datatype, before outputting you use the CAST or CONVERT functions to convert any boxes with the smalldatetime datatype to either VARCHAR or datetime datatypes. For more information, see SQL Server Books Online, Transact-SQL Reference.</p> <p>Allow Nulls - Lets you allow columns with no text.</p> <p>Collation - Lets you specify a collation sequence for the column. A column's collation sequence is used in any operation that compares values of the column to each other or to constant values. Click button to open the <a href="#">Collation Properties</a> when datatype char, varchar, text, nchar, nvarchar or ntext is selected.</p> <p>Computed Expression</p>
Row Binding	<p>Lets you bind a row to the column.</p> <p>Default Value - Lets you select a default value or type a value in the box.</p> <p>Fill Rows with Default Value - Select to fill rows with the value above.</p> <p>Default Binding - Lets you bind a default to the column, instead of declaring it.</p> <p>Rule Binding - Lets you bind a rule to the column.</p>
Identity Column	<p>Lets you set an identity column and set the identity seed and identity increment values.</p> <p>Not for Replication - Lets you indicate if the columns are not for replication.</p>

For more information, see:

[Completing the Add, Insert or Modify Column Dialog Box](#)

## COLLATION PROPERTIES

The table below describes the options and functionality on the Collation Properties dialog box:

Option	Description
SQL Collation	<p>SQL collations are a compatibility option to match the attributes of common combinations of code page number and sort orders that have been specified in earlier versions of SQL Server.</p> <p>Select to choose among the collation sequences defined by SQL Server. Then select a collation sequence from the list.</p>
Windows Collation	<p>Windows collations are collations defined for SQL Server to support Microsoft Windows® locales. By specifying a Windows collation for SQL Server, the instance of SQL Server uses the same code pages and sorting and comparison rules as an application running on a computer for which you have specified the associated Windows locale. For example, the French Windows collation for SQL Server matches the collation attributes of the French locale for Windows.</p> <p>Select to choose among the collation sequences defined by Windows. Then select a collation sequence from the list.</p>
Binary Sort	<p>Sorts and compares data in tables based on the bit patterns defined for each character. Binary sort order is case-sensitive, that is lowercase precedes uppercase, and accent-sensitive. This is the fastest sorting order.</p> <p>If this option is not selected, SQL Server follows sorting and comparison rules as defined in dictionaries for the associated language or alphabet.</p> <p>NOTE: Available only if you select Windows collation.</p>
Dictionary Sort	<p>Available only if you select Windows collation. Select this if you want comparison operations to use certain alphabetic comparison options. The alphabetic comparisons options are:</p> <p>Case Sensitive - Distinguishes between uppercase and lowercase letters. If not selected, SQL Server considers the uppercase and lowercase versions of letters to be equal. SQL Server does not define whether lowercase letters sort lower or higher in relation to uppercase letters when Case-sensitive is not selected.</p> <p>Accent Sensitive - Distinguishes between accented and unaccented characters. For example, 'a' is not equal to ' '. If not selected, SQL Server considers the accented and unaccented versions of letters to be equal.</p> <p>Kana Sensitive -Distinguishes between the two types of Japanese kana characters: Hiragana and Katakana. If not selected, SQL Server considers Hiragana and Katakana characters to be equal.</p> <p>Width Sensitive - Distinguishes between a single-byte character (half-width) and the same character when represented as a double-byte character (full-width). If not selected, SQL Server considers the single-byte and double-byte representation of the same character to be equal.</p>
Reset Defaults	<p>Applies to the column the default collation sequence for the database.</p>

## ADD, INSERT OR MODIFY COLUMN FOR ORACLE

The table below describes the options and functionality on the Add, Insert or Modify Column dialog box for Oracle.

Option	Description
Column Name	Lets you type the column name.
Datatype	Lets you select the datatype for the table. NOTE: For Oracle 9i, Rapid SQL supports the XMLType datatype. If you select CHAR, RAW or VARCHAR2, in the Width box, type the width value. If you select NUMBER, in the Width box, type the width value and in the Scale box, type the scale value.
Allow Nulls?	Lets you allow columns with no text.
Default Value	Lets you select or type a default on the column.
Comment	OPTIONAL: Lets you type a comment. A column comment can be up to 2000 characters long.
Add Button	Click to add the column.
Close Button	Click to return to Panel 2.

For more information, see:

[Completing the Add, Insert or Modify Column Dialog Box](#)

## ADD, INSERT OR MODIFY COLUMN FOR SYBASE ASE

The table below describes the options and functionality on the Add, Insert or Modify Column dialog box for Sybase ASE.

Option	Description
Column Name	Lets you type the column name.
Datatype	Lets you select the target datatype. If the Datatype requires width and/or scale, type the values in the corresponding boxes. NOTE: Because the smalldatetime datatype stores dates and time with less precision than the datetime datatype, before outputting you use the CAST or CONVERT functions to convert any boxes with the smalldatetime datatype to either VARCHAR or datetime datatypes.
Allow Nulls?	Lets you allow columns with no text.
Default Value	OPTIONAL: Lets you select or type a value.



Option	Description
Identity Column	Lets you set the identity column. The identity column will start at the value specified in Min Value, and will continue to progress until the Max value is reached.
Default Binding	Lets you bind a default to the column (instead of declaring it).
Rule Binding	Lets you bind a rule to the column.

For more information, see:

[Completing the Add, Insert or Modify Column Dialog Box](#)

## ADD, INSERT, OR MODIFY PARAMETER

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Add, Insert, or Modify Parameters dialog box lets you manage parameters. You can open the dialog box in the following wizards and editors:

- [Functions Editor \(IBM DB2 LUW\)](#)
- [Procedures Editor \(IBM DB2 LUW\)](#)
- [Functions Editor \(IBM DB2 Z/OS\)](#)
- [Procedures Editor \(IBM DB2 Z/OS\)](#)
- [Function Wizard \(IBM DB2 LUW\)](#)
- [Procedure Wizard \(DB2 LUW\)](#)
- [Function Wizard \(DB2 Z/OS\)](#)
- [Procedure Wizard \(DB2 Z/OS\)](#)

### Important Notes

None

The table below describes the options and functionality on the Add, Insert or Modify Parameter dialog box.

Option	Description
Parameter Name	Lets you enter the name of the parameter.
Datatype	Lets you select a datatype.
Width	Lets you set as appropriate to the datatype.

Option	Description
Scale	Lets you set as appropriate to the datatype.
Parameter Mode	Lets you select the parameter mode.
Add Button	Click to add the parameter.
Close Button	Click when you finish adding parameters.

## ADD, INSERT, OR MODIFY PARTITION

**NOTE:** This functionality is available for Oracle.

You can open the dialog box in the following wizards:

- [Index Wizard \(Oracle\)](#)
- [Materialized View Wizard \(Oracle\)](#)
- [Table Wizard \(Oracle\)](#)

The table below describes the options and functionality on the Add, Insert or Modify Partition dialog box.

**NOTE:** Options differ by platform and wizard.

Option	Description
Name	Lets you type the name of the new partition. Do not use single quotation marks (') in the name.
# of Partitions	Lets you specify the number of partitions.
Set Upper Bound Button	Lets you set the non-inclusive upper bound for the partitioning columns. Not available for the local index.  Opens the <a href="#">Partition Upper Bound dialog box</a> . The partitioning columns you chose in Step 1 appear in the column grid, with a default upper bound of MAXVALUE.
Tablespace	Lets you select the tablespace on which the partition should be placed.
No Logging	Select to make the partition log.
Physical	Lets you set the percent free, initial transactions and maximum transactions.
Space Attributes	Click View to view space attributes.
Storage	Lets you set any storage parameters by clicking the corresponding lists or typing the values in the corresponding boxes.

Option	Description
Subpartition definition	<p>None - Oracle creates one subpartition.</p> <p>Subpartition by name - For more information, see <a href="#">Rapid SQL and Oracle Partitioning</a>.</p> <p>Click Drop or to delete the subpartition.</p> <p>Use the tablespace(s) specified at the partition level - Select to use the tablespace(s) specified at the partition level.</p> <p>Create a list of two subpartitions - Displays a list of two subpartitions.</p> <p>Add - For more information, see <a href="#">Rapid SQL and Oracle Partitioning</a>.</p> <p>Remove - Click to remove a selected subpartition.</p>

For more information, see:

[Completing the Add, Insert or Modify Partition Dialog Box](#)

## COMPLETING THE ADD, INSERT OR MODIFY PARTITION DIALOG BOX

To complete this dialog box, do the following:

- 1 In the *Partition Definition* box, type the name of the new partition.

**NOTE:** Do not use single quotation marks (') in the name.

- If you do not want to set the non-inclusive upper bound for the partitioning columns, proceed to step 13.
- To set the non-inclusive upper bound for the partitioning columns, proceed to step 8.

- 2 Click the column in the grid and then click the *Set Upper Bound* button.

Rapid SQL opens the [Partition Upper Bound dialog box](#).

**NOTE:** The partitioning columns you chose in Step 1 appear in the column grid, with a default upper bound of MAXVALUE.

The following table describes when you should enable editing of the upper bound:

Partition Type	Index	Table	Notes
Hash	N/A	N/A	Does not have upper bound.
Range Local	Disable	Enable	
Range Global	Enable all but the last partition	Enable	The upper bound button will be disabled if there is only one partition.

Partition Type	Index	Table	Notes
Composite Global	Enable all but the last partition	Enable	The upper bound button will be disabled if there is only one partition.

- 3 Click the *Partitioning Column* list and then click the partitioning column.
- 4 In the *Upper Bound* list, click the default or type the value of the upper bound. All rows in the partition will have partitioning keys that compare to less than the partition bound and greater than or equal to the partition bound for the preceding partition.
- 5 Click *Set*.  
Rapid SQL establishes the upper bound.
- 6 Click *OK*.  
Rapid SQL closes the Rapid SQL message.
- 7 When you finish setting the upper bound for the partitioning columns, click *Close*.  
Rapid SQL returns to the Add Partition dialog box.
- 8 In the *Segment* box, click the *Tablespace* list and then click the tablespace on which the partition should be placed.
- 9 To make the partition log, click the *No Logging* check box to clear this option.
- 10 In the *Physical* box, set the percent free, initial transactions and maximum transactions by clicking the corresponding lists or typing the values in the corresponding boxes.
- 11 In the *Storage Parameters* box, set any storage parameters by clicking the corresponding lists or typing the values in the corresponding boxes.
- 12 Click *Add*.  
Rapid SQL adds the partition.
- 13 When you finish adding partitions, click the *Close* button.
  - To insert another partition into the list, click the *Insert* button, and then click *Insert* to insert the column.
  - To edit a partition, select the partition in the grid and then click the *Edit* button, and then click *Modify* to modify the column.
  - To drop a partition, click the *Drop* button and then click *Yes*.  
Rapid SQL confirms the drop, closes the Rapid SQL message, and omits the partition from the list.
- 14 When you finish specifying the partitioning attributes, click *Finish*.

For more information, see:

## [Add, Insert or Modify Partition](#)

### ADD OR MODIFY CLUSTER COLUMN

The Add or Modify Cluster Column dialog box lets you manage cluster columns. You can open the dialog box in the [Cluster Wizard \(Oracle\)](#).

#### Important Notes

None

The table below describes the options and functionality on the Add or Modify Cluster Column dialog box:

Option	Description
Column Name	Lets you type the column name.
Datatype	<p>Lets you select the datatype for the cluster.</p> <p>If you select CHAR, RAW or VARCHAR2, in the Width box, type the width value.</p> <p>If you select NUMBER, in the Width box, type the width value and in the Scale box, type the scale value.</p>

For more information, see:

#### [Completing the Add or Modify Cluster Column Dialog Box](#)

### COMPLETING THE ADD OR MODIFY CLUSTER COLUMN DIALOG BOX

To complete this dialog box, do the following:

- 1 In the *Add Cluster Column* dialog box, in the *Column Name* box, type the column name.
- 2 Click the *Datatype* list, click the datatype for the cluster.
  - If you clicked *CHAR*, *RAW* or *VARCHAR2*, in the *Width* box, type the width value.
  - If you clicked *NUMBER*, in the *Width* box, type the width value and in the *Scale* box, type the scale value.
- 3 Click the Add button.
- 4 To continue adding columns to the cluster, repeat steps 1-3.
- 5 When you finish adding columns, click Close.

Rapid SQL closes the Add Cluster Column dialog box.

For more information, see:

[Add or Modify Cluster Column](#)

## ADD OR MODIFY DATAFILE

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Add/Modify Datafile dialog box lets you add a datafile to a tablespace, or edit an existing one. You can open the dialog box in the [Tablespace Wizard \(Oracle\)](#).

### Important Notes

None

The table below describes the options and functionality on the Add/Modify Datafile dialog box:

Option	Description
Datafile Name	Type the name of the datafile in the box.
Datafile Size	Type the datafile size and then click the list to indicate MB or KB.
Reuse Existing File?	Lets you indicate whether you want to reuse an existing file or not. The default is No. To reuse the existing file, click the Yes option button.
Autoextend	Select this check box if you want to autoextend the datafile when more extents are required. This enables the options below.
Disk Space to Allocate to the Datafile When More Extents are Required	Type the disk space to allocate to the datafile when more extents are required by typing it and then click the list to indicate MB or KB.
Maximum Disk Space Allowed for Allocation to the Datafile	Specify the maximum disk space allowed for allocation to the datafile. Select Unlimited or Other. In the Other box, type the amount and then click the list to indicate MB or KB.

For more information, see:

[Completing the Add/Modify Datafile Dialog Box](#)

## COMPLETING THE ADD/MODIFY DATAFILE DIALOG BOX

To complete this dialog box, do the following:

- 1 In the *Datafile Name* box, type the datafile name.
- 2 In the *Datafile Size* box, type the datafile size and then click the list to indicate MB or KB.

### 3 Reuse Existing File?

- To reuse the existing file, click the *Yes* option button.
- If you do not want to reuse the existing file, click the *No* option button.

### 4 Select the *Autoextend* check box.

Rapid SQL auto extends the datafile.

- In the *Disk Space to Allocate to the Datafile When More Extents are Required* box, type the disk space to allocate to the datafile when more extents are required by typing it and then click the list to indicate MB or KB.
- Specify the maximum disk space allowed for allocation to the datafile by clicking the *Unlimited* option button or the *Other* option button.

**NOTE:** If you click *Other*, in the *Other* box, type the amount and then click the list to indicate MB or KB.

### 5 Click OK.

Rapid SQL adds the datafile to the tablespace definition.

### 6 Click Close.

Rapid SQL closes the Add Datafile dialog box.

For more information, see:

[Add or Modify Datafile](#)

## ALLOCATE EXTENT

This functionality is available for Oracle only.

The Allocate Extent dialog box lets you explicitly allocate extents for clusters, tables, and indexes in Oracle. Though Oracle dynamically allocates extents when additional space is required, explicit allocation of additional extents can be useful if you know that an object grows.

Explicit allocation of additional extents can be particularly helpful when using Oracle Parallel Server. When using Oracle Parallel Server and allocating additional extents, you can allocate an extent explicitly to a specific instance in order to minimize contention for free space among multiple instances.

### Important Notes

For composite-partitioned tables, you can allocate extents to subpartitions as well as partitions.

The table below describes the options and functionality on the Allocate Extent dialog box.

Option	Describe
Extent Size	Refers to a datafile in the tablespace of the table, index, or cluster. You can choose a specific datafile from which to take space for the added extent. If you choose (Default), Oracle takes the space from any accessible datafile in the tablespace containing the table, index, or cluster.
Datafile	Lets you select the new datafile.
Instance	<p>Lets you specify a freelist from which to draw the extent. If you are using Oracle Parallel Server, you can assign the new extent to a free list group associated with a specific instance. The number you enter in the Instance text box should be the number of the freelist group that you wish to use, rather than the number of the specific instance. If you are using Oracle Parallel Server and you omit this parameter, Oracle allocates the extent, but the extent is drawn from the master freelist by default. Only use this parameter for Oracle Parallel Server.</p> <p>NOTE: The number you enter in the Instance field should be the number of the free list group that you wish to use, rather than the number of the specific instance.</p>

For more information, see:

[Completing the Allocate Extent Dialog Box](#)

## COMPLETING THE ALLOCATE EXTENT DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Allocate Extent.  
Rapid SQL opens the Allocate Extent dialog box.
- 3 In the *Extent Size* box, type the new value of the extent, click the list to indicate the extent size: *MB* or *KB*.
- 4 Click the *Datafile* list, click the new datafile.
- 5 In the *Instance* box, specify the instance for Oracle to use when allocating the extent; this option only applies for Oracle with the Parallel Server in parallel query mode.
- 6 Do one of the following:



Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

For more information, see:

[Allocate Extent](#)

## ANALYZE

The Analyze dialog box lets you analyze the selected partitions. Analyzing a partition is helpful when you are trying to gather statistics about that partition, or if you are trying to validate the structure of that partition. If you are using cost-based optimization techniques, you can use the statistics to determine the most efficient use of SQL statements when accessing the partitions. In the case of hardware or system failure, analyze your partitions to make sure that their structure is not corrupt. Gathering information on chained rows lets you see if you have enough room for updates to rows in the partition.

### Important Notes

For composite-partitioned tables, you can analyze subpartitions as well as partitions.

The table below describes the options and functionality on the Analyze dialog box:

Option	Description
Partition Box	Lists all the partitions available for analysis. Click the target partition. To select all the partitions, use the Select All button.
Analysis Option	Lets you select the type of analysis. The <a href="#">table</a> below describes the analysis types and any additional information that you need to supply for each type.
Tables Only: Histogram Option	Performs the core calculations for a width-balanced histogram. For each row, Oracle returns the number of the histogram bucket appropriate for the data.  Click to open the <a href="#">Histogram Statistics dialog box</a> .

**Partition Analysis Type table**

Analysis Option	Description	Additional Information
Compute Statistics	Scans the partition in its entirety to gather the information and computes exact statistics. An analysis that computes statistics takes longer than an estimation of statistics.	None
Delete Statistics	Deletes statistics for a partition from the data dictionary. Deletes statistics not used when cost-based optimization is used for SQL statements accessing the partition.	None
For Indexes: Estimate Statistics	Scans the partition and gathers representative information from portions of the partition. Estimating statistics is not as accurate as computing statistics but the analysis is less time consuming.	<p>If you choose Estimate Statistics, you must also type in a percentage or a row count in Sample. Use the drop-down menu to choose Percent or Row Count.</p> <p>The higher the percentage or the row count, the better the estimation.</p>
For Tables: Estimate Statistics	Scans the partition and gathers representative information from portions of the partition. Estimating statistics is not as accurate as computing statistics but the analysis is less time consuming.	<p>If you choose Estimate Statistics, you must also type in a percentage or a row count in Sample. Use the drop-down menu to choose Percent or Row Count.</p> <p>The higher the percentage or the row count, the better the estimation.</p> <p>If you choose Estimate Statistics, you can also choose to generate a histogram in Histogram Option.</p> <p>If you choose to create a histogram for Selected Columns, you must also choose the columns in Selected Histogram Columns. If there are no available columns, click the Add button.</p>
For Indexes: Validate Structure	Verifies the integrity of the partition. If the partition is corrupt, Oracle returns an error message. If the partition is valid, no error is returned.	None

Analysis Option	Description	Additional Information
For Tables: Validate Structure	Verifies the integrity of the partition. If the partition is corrupt, Oracle returns an error message. If the partition is valid, no error is returned.  Oracle also verifies that the rows belongs to the correct partition. If rows do not collate correctly, they are considered invalid and Oracle lists them in a table.	You must specify a table in which to put any invalid rows.  If you do not have a table prepared to accept the invalid rows, you can check the Create Target Table if it does not exist box to have Rapid SQL create a table for you.  Clicking Cascade validates the structure of associated objects such as indexes.
For Tables: List Chained Rows	Lists the chained rows of the partition. The results are stored in the named table.	You must specify a table in which to list the chained rows.  If you do not have a table prepared to accept the results of the List Chained Rows command, you can check the Create Target Table if it does not exist box to have Rapid SQL create a table for you.

For more information, see:

[Completing the Analyze Dialog Box](#)

## COMPLETING THE ANALYZE DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Analyze.  
Rapid SQL opens the Analyze dialog box.
- 3 Select the target partition(s).  
**TIP:** To select all partitions in the index, click Select All.
- 4 In Analysis Option, select the type of analysis. The table below describes the analysis types and any additional information that you need to supply for each type.

Analysis Option	Description	Additional Information
Compute Statistics	Scans the partition in its entirety to gather the information and computes exact statistics. An analysis that computes statistics takes longer than an estimation of statistics.	None
Delete Statistics	Deletes statistics for a partition from the data dictionary. Deletes statistics not used when cost-based optimization is used for SQL statements accessing the partition.	None
Estimate Statistics	Scans the partition and gathers representative information from portions of the partition. Estimating statistics is not as accurate as computing statistics but the analysis is less time consuming.	If you choose Estimate Statistics, you must also type in a percentage or a row count in Sample. Use the drop-down menu to choose Percent or Row Count.  The higher the percentage or the row count, the better the estimation.
Validate Structure	Verifies the integrity of the partition. If the partition is corrupt, Oracle returns an error message. If the partition is valid, no error is returned.	None

5 Do one of the following:

Button	Description
	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

For more information, see:

[Analyze](#)

## HISTOGRAM STATISTICS DIALOG BOX

The Histograms Statistics dialog box lets you specify options for histograms.

## Important Notes

You can only generate a histogram for tables.

The table below describes the options and functionality on the Histogram Statistics dialog box.

Option	Description
Size	Number of bands, or sections, into which the column data is sorted. Each section refers to a range of values and Oracle fits each piece of data into one of those sections. For example, in the image above, the column data for the NUMBER_SEATS column is split into 10 sections.

For more information, see:

[Completing the Histogram Statistics Dialog Box](#)

## COMPLETING THE HISTOGRAM STATISTICS DIALOG BOX

To complete this dialog box, do the following:

- 1 Click the option button that corresponds to the level at which you want to capture histogram statistics.
- 2 Type any appropriate values in the corresponding boxes.
- 3 If you click *Selected Columns*, click the columns you want and then click the *Select* or *Unselect* button.
- 4 Click OK.

For more information, see:

[Histogram Statistics Dialog Box](#)

## BIND DATA CACHE

This functionality is available for Sybase ASE only.

The data cache holds the data, index, and log pages currently in use and pages used recently by Sybase ASE. The Bind Data Cache dialog box lets you bind the following items to the data cache:

- Entire Database
- Index(es)
- Table(s)

- Text and Image Columns from Tables

### Important Notes

- The data cache must be active before you can bind an object to it. To activate a new data cache, Sybase ASE requires that you restart the server. The bindings take effect immediately after the server restart.
- You can bind the syslogs table to a logonly cache.

## BIND PACKAGE

This functionality is available for IBM DB2 LUW for OS/390 and z/OS only.

The Bind Package Wizard lets you set package parameters, add environments, and set package properties.

### Important Notes

None

For more information, see [Bind Package Wizard - Panel 1](#).

### BIND PACKAGE WIZARD - PANEL 1

The table below describes the options and functionality on the first panel of the Bind Package wizard:

Option	Functionality
Location	Lets you select the name of the location to connect to.
Collection	Lets you select the location of the DBMS where the package binds and where the description of the package resides.
New package - PDS name	Lets you select a name.
Member	Lets you select what database request modules (DBRMs) to include in the package.
Copy from collection.	Lets you select a collection.
Package	Lets you select a package.
Version	Lets you select a version of the package.
Options	Lets you select an option.

**TIP:** To filter the Owner and Object lists, click the Filter button next to each list. Rapid SQL opens the [Filter](#) dialog box.

For more information, see [Bind Package](#).

## BIND PACKAGE WIZARD - PANEL 2

The table below describes the options and functionality on the second panel of the Bind Package wizard:

Option	Functionality
Owner	Lets you select a package owner.
Qualifier	Lets you select a qualifier, the package creator.
Action	Lets you select an action.
Version	Lets you select a version of the package.

For more information, see [Bind Package](#).

## BIND PACKAGE WIZARD - PANEL 3

The table below describes the options and functionality on the third panel of the Bind Package wizard:

Option	Functionality
Isolation	Determines how far to isolate an application from the effects of other running applications.
Keep Dynamic	Specifies that DB2 keeps dynamic SQL statements after commit points. The application does not need to prepare an SQL statement after every commit point. DB2 keeps the dynamic SQL statement until the application process ends, a rollback operation occurs or the application executes an explicit PREPARE statement with the same statement identifier.  If the prepared statement cache is active, DB2 keeps a copy of the prepared statement in the cache. If the prepared statement cache is not active, DB2 keeps only the SQL statement string past a commit point. DB2 then implicitly prepares the SQL statement if the application executes an OPEN, EXECUTE, or DESCRIBE operation for that statement.
Current Data	Determines whether to require data currency for read-only and ambiguous cursors when the isolation level of cursor stability is in effect. It also determines whether block fetching can be used for distributed, ambiguous cursors.
Degree	Determines whether to attempt to run a query using parallel processing to maximize performance. Lets you select an option.
DB Protocol	Specifies which protocol to use when connecting to a remote site that is identified by a three-part name statement.

Option	Functionality
Dynamic Rules	<p>Determines what values apply at run time for the following dynamic SQL attributes:</p> <p>The authorization ID that is used to check authorization</p> <p>The qualifier that is used for unqualified objects</p> <p>The source for application programming options that DB2 uses to parse and semantically verify dynamic SQL statements</p> <p>Whether dynamic SQL statements can include GRANT, REVOKE, ALTER, CREATE, DROP, and RENAME statements</p>
Release	<p>Determines when to release resources that a program uses, either at each commit point or when the program terminates.</p> <p>Commit - Releases resources at each commit point.</p> <p>Deallocate - Releases resources only when the program terminates.</p>
Validate	<p>Determines whether to recheck, at run time, errors found during bind. The option has no effect if all objects and needed privileges exist.</p> <p>Bind - If not all objects or needed privileges exist at bind time, the wizard displays an error messages, and does not bind the package.</p> <p>Run - If not all objects or privileges exist at bind time, the process issues warning messages, but the bind succeeds. DB2 checks existence and authorization again at run time for SQL statements that failed those checks during bind. The checks use the authorization ID of the package owner.</p>

For more information, see [Bind Package](#).

## BIND PACKAGE WIZARD - PANEL 4

The table below describes the options and functionality on the fourth panel of the Bind Package wizard:

Option	Functionality
Explain	Obtains information about how SQL statements in the package are to execute, and then inserts that information into the table owner.PLAN_TABLE, where owner is the authorization ID of the owner of the plan or package. This option does not obtain information for statements that access remote objects.
Reopt(VARS)	Re-determines the access path at run time.
Prepare	Prepares dynamic SQL statements that refer to remote objects.
ImmedWrite	Immediate writes will be done for updates made to group buffer pool dependent pagesets or partitions.
Opthint	Query optimization hints are used for static SQL.
Encoding	Lets you select type of language for the package.



Option	Functionality
Path	Lets you select a path that DB2 uses to resolve unqualified user-defined distinct types, functions, and stored procedure names (in CALL statements).
Flag	Lets you select a flags, messages to display: <ul style="list-style-type: none"> <li>- All informational, warning, error, and completion messages</li> <li>- Only warning, error, and completion messages</li> <li>- Only error and completion messages</li> <li>- Only completion messages.</li> </ul>

For more information, see [Bind Package](#).

## BIND PACKAGE WIZARD - PANEL 5

The table below describes the options and functionality on the fifth panel of the Bind Package wizard:

Option	Functionality
Enable or Disable	Lets you select system connection types that can use the package.
System	Lets you select a system.
Cname	Lets you select an option.

For more information, see [Bind Package](#).

## BIND PLAN

This functionality is available for IBM DB2 LUW for OS/390 and z/OS only.

The Bind Plan Wizard lets you set plan parameters, add packages, and set bind properties.

### Important Notes

None

For more information, see [Bind Plan Wizard - Panel 1](#).

## BIND PLAN WIZARD - PANEL 1

The table below describes the options and functionality on the first panel of the Bind Plan wizard:

Option	Functionality
Plan Name	Lets you select the plan name.
Qualifier	OPTIONAL: Lets you select a qualifier, the plan creator.
Action	OPTIONAL: Lets you select an action.
Sql Rules	OPTIONAL: Determines whether you can execute a type 2 CONNECT statement to an existing SQL connection, according to DB2 rules. Lets you select DB2 or STD.
Cache Size	OPTIONAL: Lets you select or type the cachesize in bytes, the authorization cache acquired in the EDM pool for the plan. At run time, the authorization cache stores user IDs authorized to run. Consulting the cache can avoid a catalog lookup for checking authorization to run the plan.
Plan Owner	OPTIONAL: Determines the authorization ID of the owner of the plan.
Current Server	OPTIONAL: Determines the location to connect to before running the plan.
Resource Acquire	OPTIONAL: Determines whether to acquire resources for DBRMs specified in the MEMBER list when the application first accesses them or when the plan is allocated. Local or remote packages associated with the plan acquire their resources when the application first accesses them.  Use - Acquires table space locks only when the application program bound to the plan first uses them.  Allocate - Acquires all table space locks when the plan is allocated. The value has no effect on dynamic SQL statements, which always use ACQUIRE(USE).
Disconnect	OPTIONAL: Determines which remote connections to destroy during commit operations. The option applies to any application process that uses the plan and has remote connections of any type. Regardless of the value of this option, a commit operation destroys all connections in the release pending state.  Explicit - Destroy only connections in the release pending state. This value allows you maximum flexibility for controlling remote connections.  Automatic - Destroy all remote connections.  Conditional - Destroy all remote connections unless an open cursor defined as WITH HOLD is associated with the connection.

**TIP:** To filter the Owner and Object lists, click the Filter button next to each list. Rapid SQL opens the [Filter](#) dialog box.

For more information, see [Bind Plan](#).

## BIND PLAN WIZARD - PANEL 2

The table below describes the options and functionality on the second panel of the Bind Plan wizard:

Option	Functionality
Member Name	Lets you select the member name.
PDS Name	Lets you select the partitioned data set.
Add	Click to enter each member and PDS name.

For more information, see [Bind Plan](#).

## BIND PLAN WIZARD - PANEL 3

The table below describes the options and functionality on the third panel of the Bind Plan wizard:

Option	Functionality
Location	Lets you select the name of the location to connect to.
Collection	Lets you select the location of the DBMS where the plan binds and where the description of the plan resides.
Package	Lets you select a package.

For more information, see [Bind Plan](#).

## BIND PLAN WIZARD - PANEL 4

The table below describes the options and functionality on the fourth panel of the Bind Plan wizard:

Option	Functionality
Isolation	Determines how far to isolate an application from the effects of other running applications.

Option	Functionality
Keep Dynamic	<p>Specifies that DB2 keeps dynamic SQL statements after commit points. The application does not need to prepare an SQL statement after every commit point. DB2 keeps the dynamic SQL statement until the application process ends, a rollback operation occurs or the application executes an explicit PREPARE statement with the same statement identifier.</p> <p>If the prepared statement cache is active, DB2 keeps a copy of the prepared statement in the cache. If the prepared statement cache is not active, DB2 keeps only the SQL statement string past a commit point. DB2 then implicitly prepares the SQL statement if the application executes an OPEN, EXECUTE, or DESCRIBE operation for that statement.</p>
Current Data	Determines whether to require data currency for read-only and ambiguous cursors when the isolation level of cursor stability is in effect. It also determines whether block fetching can be used for distributed, ambiguous cursors.
Degree	Determines whether to attempt to run a query using parallel processing to maximize performance. Lets you select an option.
Dynamic Rules	<p>Determines what values apply at run time for the following dynamic SQL attributes:</p> <ul style="list-style-type: none"> <li>The authorization ID that is used to check authorization</li> <li>The qualifier that is used for unqualified objects</li> <li>The source for application programming options that DB2 uses to parse and semantically verify dynamic SQL statements</li> <li>Whether dynamic SQL statements can include GRANT, REVOKE, ALTER, CREATE, DROP, and RENAME statements</li> </ul>
Release	<p>Determines when to release resources that a program uses, either at each commit point or when the program terminates.</p> <p>Commit - Releases resources at each commit point.</p> <p>Deallocate - Releases resources only when the program terminates.</p>
Validate	<p>Determines whether to recheck, at run time, errors found during bind. The option has no effect if all objects and needed privileges exist.</p> <p>Bind - If not all objects or needed privileges exist at bind time, the wizard displays an error messages, and does not bind the package.</p> <p>Run - If not all objects or privileges exist at bind time, the process issues warning messages, but the bind succeeds. DB2 checks existence and authorization again at run time for SQL statements that failed those checks during bind. The checks use the authorization ID of the plan owner.</p>

For more information, see [Bind Plan](#).

## BIND PLAN WIZARD - PANEL 5

The table below describes the options and functionality on the fifth panel of the Bind Plan wizard:

Option	Functionality
Explain	Obtains information about how SQL statements in the member list of the plan, are to execute, and then inserts that information into the table owner.PLAN_TABLE, where owner is the authorization ID of the owner of the plan or package. This option does not obtain information for statements that access remote objects.
Reopt(VARS)	Re-determines the access path at run time.
Prepare	Prepares dynamic SQL statements that refer to remote objects.
ImmedWrite	Immediate writes will be done for updates made to group buffer pool dependent pagesets or partitions.
Opthint	Query optimization hints are used for static SQL.
Encoding	Lets you select type of language for the package.
Path	Lets you select a path that DB2 uses to resolve unqualified user-defined distinct types, functions, and stored procedure names (in CALL statements).
Flag	Lets you select a flags, messages to display: <ul style="list-style-type: none"> <li>- All informational, warning, error, and completion messages</li> <li>- Only warning, error, and completion messages</li> <li>- Only error and completion messages</li> <li>- Only completion messages.</li> </ul>

For more information, see [Bind Plan](#).

## BIND PLAN WIZARD - PANEL 6

The table below describes the options and functionality on the sixth panel of the Bind Plan wizard:

Option	Functionality
Enable or Disable	Lets you select system connection types that can use the plan or package.
System	Lets you select a system.
Cname	Lets you select an option.

For more information, see [Bind Plan](#).

## BUILD QUERY

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

[Query Builder](#) is a database productivity tool that lets you construct, structure, and manipulate up to five different types of queries simultaneously. It includes a separate graphical interface that opens within your current workspace. You can run Query Builder against all Embarcadero Technologies supported database platforms.

## CHANGE CATEGORY

**NOTE:** This functionality is available for Oracle only.

The Change Category dialog box lets you change the category of a target Stored Outline and specify a new outline category to move the outline.

### Important Notes

None

The table below describes the options and functionality on the Change Category dialog box:

Option	Description
Owner	Displays the outline owner.
Outline Name	Display the outline name.
Current Category	Displays the current category of the outline.
New Category	Lets you select a new outline category to move the outline.

For more information, see [Completing the Change Category Dialog Box](#).

## COMPLETING THE CHANGE CATEGORY DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Outlines node.  
Rapid SQL displays the Outlines in the Database Explorer.
- 2 On the Database Explorer, right-click the Outline, and then select Change Category.  
Rapid SQL opens the Change Category dialog box.
- 3 In the New Category box, enter the target category.
- 4 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

For more information, see:

[Change Category](#)

## CHANGE CONTAINER SIZE

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Change Container Size lets you increase and resize a container, the allocation of space to a tablespace. Depending on the tablespace type, the container can be a directory, device, or file.

### Important Notes

None

The table below describes the options and functionality on the Change Container Size dialog box:

Option	Description
Increase Container Size by (Extent)	Select to increase the container size.
Change Container Size to (Resize)	Select to resize the container.
Size	Lets you select a container size.

## CHANGE DATABASE COMMENT

This functionality is available for IBM DB2 for Linux, Unix, and Windows.

The Change Database Comment dialog box lets you change the comment for a database.

**Important Notes**

None

The table below describes the options and functionality on the Change Database Comment dialog box:

Option	Description
Current Comment	Displays the current comment.
New Comment	Lets you type a new comment.

## CHANGE PASSWORD

This functionality is available for Microsoft SQL Server, Oracle, and Sybase ASE only.

The Change Password dialog box lets you change user passwords, which you should do on a regular basis to guard against security leaks.

**Important Notes**

None

The table below describes the options and functionality on the Change Password dialog box.

Option	Description
New Password	Lets you type the new password.
Confirm Password	Lets you retype the new password.

For more information, see [Completing the Change Password Dialog Box](#).

## COMPLETING THE CHANGE PASSWORD DIALOG BOX

- 1 On the *Database Explorer*, select the Users node.  
Rapid SQL displays the Users in the Database Explorer.
- 2 On the Database Explorer, right-click the User, and then select Change Password.  
Rapid SQL opens the Change Password dialog box.
- 3 OPTIONAL: In the Old Password box, type the old password.
- 4 In the *New Password* box, type the new password.
- 5 In the *Confirm Password* box, type the new password.



6 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

For more information, see:

[Change Password](#)

## CHANGE STATUS (TABLESPACES)

You can change the online, offline, or read only status of a tablespace to control access to its segments. In addition, when setting a tablespace offline, you can choose between NORMAL, TEMPORARY, or IMMEDIATE modes of taking the tablespace offline.

**NOTE:** This functionality is available for Oracle.

### To change the status of a tablespace

- 1 On the *Datasource Explorer*, expand the Tablespaces node.

Rapid SQL displays the Tablespaces in the right pane of the Explorer window.

- 2 Select one or more tablespaces, on the *Datasource Explorer* tool bar click Command, and then select Change Status.

Rapid SQL opens the Change Status wizard.

- 3 Use the following table as a guide to understanding and modifying settings in the wizard:

Step	Settings and tasks
<b>Action options</b>	From the Change Status dropdown, select ONLINE, OFFLINE, or READ ONLY. If you select OFFLINE, from the OfflineMode dropdown, select NORMAL, TEMPORARY, or IMMEDIATE.
<b>Dependencies</b>	Review the referring and referred objects that will be automatically resolved when you execute this operation.
<b>Preview</b>	Preview the DDL generated for the operation and when ready, click <b>Execute</b> .

## CHANGE STATUS (TRIGGERS)

This action lets you build an ALTER TRIIGGER statement with an ACTIVE or INACTIVE argument, and issue the statment against one or more triggers.

**NOTE:** This functionality is available for InterBase/Firebird.

### To change the status of a trigger

- 1 On the *Datasource Explorer*, expand the Triggers node.
- 2 Select one or more triggers, right-click the selected triggers, and then select Change Status from the context menu.

Rapid SQL opens the **Change Status** dialog.

- 3 Use the following table as a guide to understanding and modifying settings in the dialog:

Step	Settings and tasks
<b>Action options</b>	With <b>Enabled</b> selected, the ALTER TRIGGER statement is created with an ACTIVE argument. Deselected, the ALTER TRIGGER statement is created with an INACTIVE argument.
<b>Dependencies</b>	Lists referring and referenced objects potentially impacted by the change. For details, see <a href="#">Dependencies</a> .
<b>Preview</b>	Preview the DDL generated for the operation and when ready, click <b>Execute</b> .

## CHECKPOINT

This functionality is available for Microsoft SQL Server, and Sybase ASE only.

The Checkpoint Database(s) dialog box lets you force all dirty pages for the current database to be written to disk. A dirty page is any data or log page which, after being read into the buffer cache, is modified but not yet written to disk. The Checkpoint statement saves time in a subsequent recovery by creating a point at which all modifications to data and log pages are guaranteed to have been written to disk. If the current database is in log-truncate mode, CHECKPOINT also truncates the inactive portion of the log.

### Important Notes

The default permission for executing a checkpoint is the db\_owner fixed database role.

For more information, see [Completing the Checkpoint Database\(s\) Dialog Box](#).

## COMPLETING THE CHECKPOINT DATABASE(S) DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Databases node.

Rapid SQL displays the Databases in the Database Explorer.

- 2 On the Database Explorer, right-click the database, and then select Checkpoint.  
Rapid SQL opens the Checkpoint Database(s) dialog box.
- 3 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

## CLONE PARTITION

This functionality is available for IBM DB2 LUW for OS/390 and z/OS only.

The Clone Partition dialog box lets you clone a definition of a partition.

### Important Notes

None

The table below describes the options and functionality on the dialog box:

Option	Description
Select 1 or More Partitions	Lets you select partitions to clone.
Select All	Click select all partitions to clone.
Unselect All	Click to unselect all partitions to clone.

## CLONE TABLE

**NOTE:** This functionality is available for Oracle.

The Clone Table dialog box lets you clone a definition of a table.

### Important Notes

None

The table below describes the options and functionality on the dialog box:

Option	Description
Select 1 or More Tables	Lets you select tables to clone.
Select All	Click select all tables to clone.
Unselect All	Click to unselect all tables to clone.

## COALESCE

You can maximize the size of free space chunks in tablespaces to avoid the situation in which an object cannot acquire enough contiguous free space to accommodate its next extent size. Towards this goal, look for opportunities to coalesce adjacent blocks of free space into a single, larger block.

**NOTE:** The Coalesce dialog box is available for Oracle 7 or later.

Starting with Version 7, Oracle automatically coalesces adjacent free space chunks with a background process. However, it still supports the commands for coalescing free space manually. Depending on the size of the tablespace, coalescing its free space can take a long time. So determine when to perform this operation. If you coalesce immediately, Rapid SQL locks the tablespace.

### Important Notes

- You cannot coalesce on an UNDO tablespace.

### To coalesce a tablespace

- 1 On the *Datasource Explorer*, expand the Tablespaces node.

Rapid SQL displays the Tablespaces in the right pane of the Explorer window.

- 2 Select one or more tablespaces, on the *Datasource Explorer* tool bar click Command, and then select Coalesce.

Rapid SQL opens the Coalesce wizard.

- 3 Use the following table as a guide to understanding and modifying settings in the wizard:

Step	Settings and tasks
<b>Action options</b>	Verify display of the tablespaces to be coalesced.
<b>Dependencies</b>	Review the referring and referred objects that will be automatically resolved when you execute this operation.
<b>Preview</b>	Preview the DDL generated for the operation and when ready, click <b>Execute</b> .

## COMPILE

This functionality is available for IBM DB2 LUW for Linux, Unix, and Windows Oracle only.

The Compile dialog box lets you recompile specific objects by issuing the proper ALTER statement. The explicit recompilation of invalid objects eliminates the need for implicit run-time recompilation which, in turn, can cause run-time compilation errors and performance overhead. Recompile objects after you make changes to that object or dependent objects.

### Important Notes

None

The table below describes the options and functionality on the Confirm Compile dialog box for Oracle java classes:

Option	Description
Java class(es) to be compiled	Displays the java classes.

The table below describes the options and functionality on the Confirm Compile dialog box for Oracle:

Option	Description
Do not compile dependent objects	Compiles only the current object and to ignore the object's dependencies. This is the default setting.
Compile dependent objects	Compiles statements for all objects referenced by the object being compiled. For example, if you compile a function that references a specific procedure and you select to compile the dependent objects, an ALTER COMPILE statement is created for that referenced procedure.  Compile only invalid dependent objects - Creates ALTER COMPILE statements for only those objects that are currently invalid.  Compile dependent system objects - Compiles all of the referenced objects with the debug option.
Compile with debug option	Selecting the Compile with debug option check box instructs the Oracle PL/SQL compiler to generate and store the code for use in debugging sessions.

The table below describes the options and functionality on the Confirm Compile dialog box for IBM DB2 LUW for Linux, Unix, and Windows:

Option	Description
Compare with release option	Compiles the stored procedure into a state that will run optimally in a production environment.
Compile with debug option	Compiles the stored procedure with the debugging information and symbols required to service step, pause, and breakpoint requests.
Keep current specific name	Notifies DB2 that the stored procedure will keep its current system assigned name. With this option, DB2 is not required to reset internal references to the target object.

To recompile an object it must belong to your schema or you need ALTER ANY privileges on that object. You must also have appropriate privileges for any associated objects.

- [Functions](#)
- [Java Classes](#)
- [Java Sources](#)
- [Materialized Views](#)
- [Packages](#)
- [Package Bodies](#)
- [Procedures](#)
- [Type Bodies](#)
- [Types](#)
- [Users](#)
- [Views](#)

## COMPILING FUNCTIONS

Rapid SQL lets you recompile a function. Oracle first recompiles any invalid objects on which the function depends. In addition, it marks any objects that depend on the function as invalid.

To recompile a function that is part of a package, compile the package itself. Rapid SQL uses the ALTER FUNCTION statement to compile a stand-alone function. However, you should not use the ALTER FUNCTION statement to individually recompile a function that is part of a package.

## COMPILING JAVA CLASSES

Rapid SQL lets you compile a Java class. Oracle resolves references to other Java classes.

## COMPILING JAVA SOURCES

Oracle lets you compile a Java source. Oracle resolves references to other Java sources.

## COMPILING MATERIALIZED VIEWS

Rapid SQL lets you compile materialized views. If a materialized view fails to revalidate after you recompile, that materialized view cannot be fast refreshed ON DEMAND or used for query rewrite.

For more information, see:

[Compile](#)

## COMPILING PACKAGES AND PACKAGE BODIES

Rapid SQL lets you recompile a package, and recompiles all package objects together.

Recompiling a package in Rapid SQL compiles both the package specification and the package body by issuing two consecutive ALTER statements. However, Rapid SQL shows only the ALTER statement for the specification in the Preview: Confirm Compile dialog box. You can recompile only the package body by explicitly compiling the package body itself.

When recompiling the entire package, Oracle recompiles the package even if it is invalid. However, if there are compilation errors, the package remains invalid and Oracle invalidates all dependent objects.

Recompiling only a package body does not invalidate objects that depend upon the package specification regardless of whether or not the package body has compilation errors.

For more information, see:

[Compile](#)

## COMPILING PROCEDURES

Rapid SQL lets you compile a procedure that is part of a package, by compiling the package itself. Rapid SQL uses the ALTER PROCEDURE statement to compile a stand-alone procedure. However, you should not use the ALTER PROCEDURE statement to individually recompile a procedure that is part of a package.

The table below describes the options and functionality on the Confirm Compile dialog box for IBM DB2 LUW for Linux, Unix, and Windows:

Option	Description
Compare with release option	Compiles the stored procedure into a state that will run optimally in a production environment.

Option	Description
Compile with debug option	Compiles the stored procedure with the debugging information and symbols required to service step, pause, and breakpoint requests.
Keep current specific name	Notifies DB2 that the stored procedure will keep its current system assigned name. With this option, DB2 is not required to reset internal references to the target object.

The table below describes the options and functionality on the Confirm Compile dialog box for Oracle:

Option	Description
Do not compile dependent objects	Compiles only the current object and to ignore the object's dependencies. This is the default setting.
Compile dependent objects	Compiles statements for all objects referenced by the object being compiled. For example, if you compile a function that references a specific procedure and you select to compile the dependent objects, an ALTER COMPILE statement is created for that referenced procedure.  Compile only invalid dependent objects - Creates ALTER COMPILE statements for only those objects that are currently invalid.  Compile dependent system objects - Compiles all of the referenced objects with the debug option.
Compile with debug option	Selecting the Compile with debug option check box instructs the Oracle PL/SQL compiler to generate and store the code for use in debugging sessions.

For more information, see:

[Compile](#)

## COMPILING TYPES AND TYPE BODIES

Rapid SQL lets you recompile a type. Rapid SQL recompiles both the type specification and the type body.

For more information, see:

[Compile](#)

## CONNECTION EDITOR

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Connection Editor lets you modify plan and package connections.



The table below describes the options and functionality of the Connection Editor:

Option	Description
Connections	Displays the connections for the plan or package.
Add	Click to add the connection.

## CONVERT TO PARTITIONED

**NOTE:** This functionality is available for Oracle only.

The Convert to Partitioned Wizard lets you select partitioning methods, columns, and add partitions. Index-organized tables take up less storage space and quickly access table rows. Index-organized tables stores rows in primary key order reducing the amount of storage space needed.

### Completing the Convert to Partitioned Wizard

To complete this wizard, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select **Convert to Partitioned**.  
Rapid SQL opens the **Convert to Partitioned** wizard.
- 3 Complete the wizard panels. See the following topics for detailed information:
  - [Convert to Partitioned Wizard - Panel 1](#)
  - [Convert to Partitioned Wizard - Panel 2](#)
  - [Convert to Partitioned Wizard - Panel 3](#)
- 4 When ready, click **Finish** to open the **Preview:Create** dialog box.

### CONVERT TO PARTITIONED WIZARD - PANEL 1

The table below describes the options and functionality on this panel of the Convert to Partitioned Wizard:

Option	Description
Owner	Displays the table owner.
Table	Displays the table name.
Tablespace	Lets you select the tablespace.
Partitioning Method	Lets you select a partitioning method.
Do you want to enable row movement?	Lets you enable row movement.

For more information, see [Convert to Partitioned Wizard](#).

## CONVERT TO PARTITIONED WIZARD - PANEL 2

The table below describes the options and functionality on this panel of the Convert to Partitioned Wizard:

Option	Description
Select the partitioning columns	Available Columns - Lets you select at least one available column, and then click the right arrow to move the column to the Selected Columns.
Create a list of ordered partitions	Add - Opens the <a href="#">Add Partition dialog box</a> . Insert - Opens the <a href="#">Insert Partition dialog box</a> . Edit - Opens the <a href="#">Modify Partition dialog box</a> . Drop - Lets you drop a partition.

For more information, see [Convert to Partitioned Wizard](#).

## CONVERT TO PARTITIONED WIZARD - PANEL 3

The table below describes the options and functionality on this panel of the Convert to Partitioned Wizard:

Option	Description
Enter a table description	Lets you type an optional table description.

For more information, see [Convert to Partitioned Wizard](#).

## COPY OBJECT NAMES

This functionality is available for Microsoft SQL Server and Oracle only.

The Copy Object Names functionality lets you copy and then paste object name(s) into other applications.

### Important Notes

None

## CREATE ALIAS

This functionality is available for IBM DB2 LUW for Linux, Unix, and Windows, IBM DB2 LUW for OS/390 and z/OS, and Microsoft SQL Server only.

The Create Alias dialog box lets you create an alternate name for an object without requiring you to know any of the underlying SQL commands. The dialog box constructs the necessary CREATE ALIAS statement based on the information you supply.

**TIP:** If an object name is lengthy, create an alias, and then use the alias during a query.

### Important Notes

None

For more information, see:

[Creating an Alias for One Object](#)

[Creating an Alias for Multiple Objects](#)

[Completing the Create Alias Dialog Box](#)

## CREATING AN ALIAS FOR ONE OBJECT

The Create Alias dialog box lets you create an alternate name for an object without requiring you to know any of the underlying SQL commands. The dialog box constructs the necessary CREATE ALIAS statement based on the information you supply.

The table below describes the options and functionality on the Create Alias dialog box:

Option	Description
Alias	Lets you type the alias name.
Owner	Lets you select the target alias owner.

For more information, see:

## [Create Alias](#)

### CREATING AN ALIAS FOR MULTIPLE OBJECTS

The Create Aliases dialog box lets you assign an Alias Owner to multiple objects. The Create Alias dialog box lets you create an alternate name for an object without requiring you to know any of the underlying SQL commands. The dialog box constructs the necessary CREATE ALIAS statement based on the information you supply.

The table below describes the options and functionality on the Create Alias dialog box:

Option	Description
Alias Owner	Lets you select the target alias owner.

For more information, see:

## [Create Alias](#)

### COMPLETING THE CREATE ALIAS DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Create Alias.  
Rapid SQL opens the Create Alias dialog box.
- 3 In Alias, type the alias name.
- 4 For one object, in Owner, select the target alias owner.
- 5 Do one of the following:

Button	Description
	Opens the <a href="#">Preview dialog box</a> .
	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
	Executes the task.

Button	Description
	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

For more information, see [Create Alias](#).

## CREATE CLONE

Rapid SQL provides support for the ADD CLONE option functionality, letting you create a clone on an existing base table. This creates a table with a cloned definition but with no data.

**NOTE:** This functionality is available as of IBM DB2 for z/OS version 9. The base table must reside on a universal tablespace.

### To create a clone for one or more tables:

- 1 On the *Datasource Explorer*, expand the Tables node.  
Rapid SQL displays the tables in the right pane of the Explorer window.
- 2 In the right pane of the application, select one or more tables.
- 3 From the Command menu select Clone Actions > Create Clone.  
Rapid SQL opens the **Create Clone** dialog box.
- 4 Use the following table as a guide to understanding and modifying settings in this wizard:

Step	Settings and tasks
<b>Action Options</b>	Lets you provide a <b>Clone Name</b> .
<b>Dependencies</b>	Lists referring and referenced objects potentially impacted by the change. For details, see <a href="#">Dependencies</a> .
<b>Preview</b>	Displays the DDL that will execute the object action. For details, see <a href="#">Preview</a> .

- 5 Click **Execute**. For information on the other options, see [Preview](#) and [Scheduling](#).

The clone for a table is not displayed in the Datasource Explorer listing of tables for a data source. To determine whether a table has a clone, see the **Dependencies** tab/panel on a Table editor or on the wizard for object actions such as Drop or Rename.

After creating a clone for a table, you can exchange data between the base table and the clone. For related information, see the following topics:

- [Exchange Data With Clone](#)
- [Drop Clone](#)

## CREATE INSERT STATEMENTS

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Create Insert Statements dialog box lets you create Insert Statements for selected columns and data.

### Important Notes

None

The table below describes the options and functionality on the Create Insert Statements dialog box:

Option	Description
Columns	Lets you select columns for a Select statement.
Where	OPTIONAL: Lets you enter a WHERE statement.
Include owner information in insert statements	Select to include the owner information.
Set row count	Lets you specify a row count, the number of rows in a table that were affected by the Insert statement executed against the table, or a view based on the table.

For more information, see:

[Completing the Create Insert Statements Dialog Box](#)

## COMPLETING THE CREATE INSERT STATEMENTS DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Tables node.  
Rapid SQL displays the tables in the Database Explorer.
- 2 On the Database Explorer, right-click the table, and then select Create Insert Statements.  
Rapid SQL opens the Create Insert Statements dialog box.
- 3 In Columns, select columns for a SELECT statement.
- 4 OPTIONAL: In Where, lets you enter a WHERE statement.
- 5 Click OK.

## CREATE OR EDIT JAVA SOURCE

**NOTE:** This functionality is available for Oracle only.

The Java Editor lets you enter Java code. The table below describes the Java Editor toolbar options:

Option	Description
Lock/Unlock Connection	Click to lock or unlock connection.
Create	Click to open the Create Options dialog box, which lets you select the options for creating the java source.
Errors	Click to split the workspace in half, displaying the error messages in the lower half of the workspace.

### Important Notes

None

## CREATE LIKE

This functionality is available for InterBase/Firebird and Oracle.

The Create Like Editor lets you create a new object based on an existing object. The Create Like Editor is available for [logins](#), [tables](#), and [users](#).

### Important Notes

None

## CREATE LIKE EDITOR FOR LOGINS

The Create Like Editor lets you create a new login based on an existing login.

The following tabs are available in the Create Like Editor for Logins:

- [Definition](#)
- [Users](#)
- [Roles](#)

### CREATE LIKE EDITOR FOR LOGINS - DEFINITION TAB

The table below describes the options and functionality of the Definition Tab:

Option	Description
New Login Name	Lets you type the new login name.
SQL Server Login	OPTIONAL: Lets you type the password.
NT User or Group	Lets you select the domain.
Default Database	Lets you select the default database.
Default Language	Lets you select the default language.

**NOTE:** Click Create to open the [Preview:Create dialog box](#).

For more information, see [Create Like Editor for Logins](#).

#### CREATE LIKE EDITOR FOR LOGINS - USERS TAB

The table below describes the options and functionality of the Users Tab:

Option	Description
Databases where the login does NOT have a user account.	Lets you select the target database.
Databases where the login HAS a user account.	Lets you select the target database.
Add User	Click to open the <a href="#">Create User Account dialog box</a> .
Drop User	Click to move the database to the Databases where the login does NOT have a user account window.
Edit User	Click to open the <a href="#">Edit User Accounts dialog box</a> .

**NOTE:** Click Create to open the [Preview:Create dialog box](#).

For more information, see [Create Like Editor for Logins](#).

#### CREATE LIKE EDITOR FOR LOGINS - ROLES TAB

The table below describes the options and functionality of the Roles Tab:



Option	Description
Grant	Click to open the <a href="#">Grant Role dialog box</a> .
Revoke	Click to open the <a href="#">Revoke Roles dialog box</a> .

**NOTE:** Click Create to open the [Preview>Create dialog box](#).

For more information, see [Create Like Editor for Logins](#).

## CREATE LIKE EDITOR FOR TABLES

The Create Like Editor lets you create a new table based on an existing table. The Create Like Editor lets you manage:

- Table columns
- Table constraints
- Table storage
- Table partitions

The following tabs are available in the Create Like Editor:

- [Columns](#)
- [Constraints](#)
- [Storage](#)
- [Performance](#)
- [Partitions](#)
- [Comment](#)

## CREATE LIKE EDITOR FOR TABLES - COLUMNS TAB

The Columns Tab of the Create Like Editor lets you manage columns for the new table. The table below describes the options and functionality of the Columns Tab:

Option	Description
New Table Name	Lets you type the new table name.
Add	Click to open the <a href="#">Add Column dialog box</a> .
Insert	Click to open the <a href="#">Insert Column dialog box</a> .
Edit	Click to open the <a href="#">Modify Column dialog box</a> .
Drop	Lets you drop the selected column.
Up	Lets you move the selected column up.
Down	Lets you move the selected column down.
LOB Storage	Click to open the <a href="#">Lob Storage Definition dialog box</a> .
Convert to LOB	Click to open the <a href="#">Lob Storage Definition dialog box</a> .

**NOTE:** Click Create to open the [Preview:Create dialog box](#).

For more information, see [Create Like Editor for Tables](#).

## CREATE LIKE EDITOR FOR TABLES - CONSTRAINTS TAB

The Constraints Tab of the [Create Like Editor](#) lets you manage constraints for the new table. Rapid SQL arranges the constraints in a tree structure. The tree includes folders which contain all constraints associated with the target table. The objects are organized in folders based on the type of constraint:

- Primary Key
- Unique Key
- Check Constraint
- Foreign Key

**TIP:** Double-click a constraint to open a dialog box with detailed information on the target constraint.

The table below describes the options and functionality on the Constraints Tab:

Option	Description
Add	Click to open the <a href="#">Index Constraint dialog box</a> .
Edit	Click to open the <a href="#">Index Constraint dialog box</a> .
Drop	Drops the selected constraint.

**NOTE:** Click Create to open the [Preview:Create dialog box](#).

For more information, see [Create Like Editor for Tables](#).

#### CREATE LIKE EDITOR FOR TABLES - STORAGE TAB

The Storage Tab of the [Create Like Editor](#) lets you manage storage for the new table.

The table below describes the options and functionality on the Storage Tab:

Option	Description
Data Block Storage	The storage parameter lets you tune performance by minimizing the occurrence of row migration and chaining caused by update operations that extend the length of rows stored on the data block.  Percent Used - Lets you type a value in the corresponding box.  Maximum - The maximum parameter limits concurrency on a data block.
Extents	Initial - The initial parameter ensures that a minimum number of concurrent transactions can update a data block, avoiding the overhead of allocating a transaction entry dynamically.  Percent Increase - Magnifies how an object grows and, can materially affect available free space in a tablespace. Select a value in the corresponding box.
Filegroup	Lets you select the filegroup within the database the table is stored. This is for Microsoft SQL Server 7.0 or later.
Text Image Filegroup	Lets you select the text image filegroup within the database on which to place any text, image, and/or next columns. This is for Microsoft SQL Server 7.0 or later.

**NOTE:** Click Create to open the [Preview>Create dialog box](#).

For more information, see [Create Like Editor for Tables](#).

#### CREATE LIKE EDITOR FOR TABLES - PERFORMANCE TAB

The table below describes the options and functionality on the Performance Tab:

Option	Description
Parallel Query Option	Lets you process queries using many query server processes running against multiple CPUs, which provides substantial performance gains such as reduction of the query completion time.
Logging	Select to have all changes stored.
Cache	Keeps the blocks in memory by placing it at the most recently used end. This option is useful for small lookup tables.

**NOTE:** Click Create to open the [Preview>Create dialog box](#).

For more information, see [Create Like Editor for Tables](#).

#### CREATE LIKE EDITOR FOR TABLES - PARTITIONS TAB

The table below describes the options and functionality on the Partitions Tab:

Option	Description
Is Table Partitioned	Lets you select if the table is partitioned.
Partitioned Method	Lets you select the partitioned method.
Add	Click to open the <a href="#">Partitioned Columns dialog box</a> .

**NOTE:** Click Create to open the [Preview:Create dialog box](#).

For more information, see [Create Like Editor for Tables](#).

#### CREATE LIKE EDITOR FOR TABLES - COMMENT TAB

The Comment Tab lets you type a comment up to 2000 characters long.

**NOTE:** Click Create to open the [Preview:Create dialog box](#).

For more information, see [Create Like Editor for Tables](#).

### CREATE LIKE EDITOR FOR USERS

The Create Like Editor lets you create a new user based on an existing user. The following tabs are available in the Create Like Editor:

- [Definition](#)
- Object Privileges
- System Privileges

#### CREATE LIKE EDITOR FOR USERS - DEFINITION TAB

The table below describes the options and functionality of the Definition Tab:

Option	Description
Login Name	Lets you select login name.
Full Name	Lets you type the full name.
Add	Click to open the <a href="#">Add Alias</a> dialog box.

Option	Description
Drop	Click to open the <a href="#">Drop Aliased Login</a> dialog box.
Add/Change	Lets you add or change a group.
Drop	Lets you drop a group.

For more information, see [Create Like Editor for Users](#).

#### ADD ALIAS

The Add Alias dialog box lets you create an aliased login.

For more information, see [Create Like Editor for Users](#).

#### DROP ALIASED LOGIN

The Drop Aliased Login dialog box lets you drop an aliased login.

For more information, see [Create Like Editor for Users](#).

#### CREATE LIKE EDITOR FOR USERS - OBJECT PRIVILEGES

The table below describes the options and functionality of the Object Privileges Tab:

Option	Description
Object Type	Lets you select an object type.
Owner	Lets you select an owner.
Grant	Opens the <a href="#">Grant Privilege(s) To</a> dialog box.
Revoke	Opens the <a href="#">Revoke Privilege(s) From</a> dialog box.

For more information, see [Create Like Editor for Users](#).

#### CREATE LIKE EDITOR FOR USERS - SYSTEM PRIVILEGES TAB

The table below describes the options and functionality of the System Privileges Tab:

Option	Description
Grant	Opens the <a href="#">Grant Privilege(s) To</a> dialog box.
Revoke	Opens the <a href="#">Revoke Privilege(s) From</a> dialog box

For more information, see [Create Like Editor for Users](#).

## CREATE OR EDIT USER ACCOUNTS

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Create or Edit User Accounts dialog box lets you manage user accounts.

### Important Notes

None

The table below describes the options and functionality on the Create or Edit User Accounts dialog box:

Option	Description
User Name	Lets you type the user name.
Database roles	Lets you select the roles for the user account.

## CREATE OR MODIFY CONSTRAINT

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Interbase/Firebird, Oracle, SQL Server, and Sybase.

The following table describes the settings included in the Check Constraints editors and wizards for each supported platform.

Step	Settings and tasks	
<b>Definition</b>	<b>Table Owner</b> or <b>Table Schema</b> (DB2 LUW, DB2 z/OS, SQL Server, Oracle, Sybase)	Select the owner of the table for which the check constraint is to be created.
	<b>Table</b> or <b>Table Name</b> (DB2 LUW, DB2 z/OS, Interbase/Firebird, SQL Server, Oracle, Sybase)	Lets you select where you want to place the check constraint.
	<b>Name</b> (DB2 LUW, DB2 z/OS, Interbase/Firebird, SQL Server, Oracle, Sybase)	Lets you type the name of the constraint, which can be up to 30 characters long and must be unique across the entire database.
	<b>Enabled</b> or <b>Enforced</b> (DB2 LUW, DB2 z/OS, SQL Server, Oracle)	Enables or disables the check constraint.
	<b>Not for Replication</b> (SQL Server)	Enabling this feature ensures that the constraint is not enforced if the table is replicated for another database.
	<b>Check Condition</b> (DB2 LUW, DB2 z/OS, Interbase/Firebird, SQL Server, Oracle, Sybase)	Type the check constraint condition. As a convenience, use the <b>Table Columns</b> button to paste in column names as part of the condition.

Step	Settings and tasks
<b>Comment</b> (DB2 LUW)	Lets you type a comment.
<b>DDL View</b>	Preview the DDL generated by your choices. Finally, use the <b>Execute</b> button to create the check constraint.

## CREATE SYNONYM

**NOTE:** The Create Synonym functionality is available for IBM DB2 LUW for Linux, Unix, and Windows and Oracle.

The Create Synonym dialog box lets you create synonyms for objects.

A synonym is an alternate name for an object in the database. You can define synonyms on tables, views, sequences, procedures, functions, packages, package bodies, materialized views, and Java class objects.

There are two types of synonyms:

- Public
- Private

Public synonyms can be used by everyone in the database. Private synonyms belong to the synonym owner and the owner can control the synonym's availability to others.

Synonyms simplify object names and provide added security. Synonyms let you:

- Simplify SQL statements.
- Reference an object without needing to specify its owner.
- Reference a remote object without needing to specify its database.
- Alias an object so its purpose becomes more understandable.
- Mask the name, owner, and location of an object.

**NOTE:** The schema object for which you define a synonym cannot be contained in a package.

Rapid SQL lets you specify options for creating synonyms with the Create Synonym and Create Synonyms dialog boxes. The dialog box that Rapid SQL returns depends on whether you are creating a synonym for a single object or are creating synonyms for multiple objects at the same time.

### Important Notes

None

The table below describes the options and functionality on the Create Synonym dialog box.

Option	Description
Synonym	<p>A synonym can have the same name as the base object if that object is contained in another schema. For example, you can define the synonym EMPLOYEE on the table COMPANY.EMPLOYEE as long you are not working in the COMPANY schema.</p> <p>If you choose to create synonyms for multiple objects in a single operation, Rapid SQL uses the object name as the synonym name. After creating synonyms in such a manner, users no longer need to specify the owner name.</p>
Synonym Scope	Lets you choose whether you want the synonym to be private or public by clicking the Public or Private option buttons. If you click the Private option button, choose a schema in which to create the synonym.

## COMPLETING THE CREATE SYNONYM DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Create Synonym.  
Rapid SQL opens the Create Synonym dialog box.
- 3 In the *Synonym Scope* box, click the *Public* or *Private* option button.
- 4 If you clicked the *Private* option button, click the list and then click the owner.
- 5 Do one of the following:

Button	Description
	Opens the <a href="#">Preview dialog box</a> .
	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
	Executes the task.
	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.



For more information, see:

[Create Synonym](#)

## CREATE VIEW

Right-clicking an object in the right pane of the explorer and selecting Create View opens a wizard that lets you create a new view.

## DATALINK OPTIONS

**NOTE:** This functionality is available for IBM DB2 UBB for OS/390 and z/OS only.

The DataLink Options dialog box lets you manage datalinks, datatypes that enable logical references from the database to a file that is stored outside the database.

### Important Notes

None

The table below describes the options and functionality on the DataLink Options dialog box:

Option	Description
Datalink Options	Lets you select datalink options: No Link Control File Link Control File Link Options
File-Link Options	Lets you select file-link options: Read On Unlink Write Permission Recovery

## DBCC

This functionality is available for [Microsoft SQL Server](#) and [Sybase ASE](#) only.

The DBCC (Database Consistency Check) dialog box lets you:

- Specify single or multiple tables or indexes for validation.
- Perform database-wide validations.

- Perform object-level validations of databases.

### Important Notes

None

For more information, see [Completing the DBCC Dialog Box](#).

## COMPLETING THE DBCC DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select DBCC.  
Rapid SQL opens the DBCC dialog box.
- 3 Click the *DBCC* list and click the *DBCC* operation you want to perform.  
**NOTE:** DBCC options differ for [Microsoft SQL Server](#) and [Sybase ASE](#) objects.
- 4 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

For more information, see:

[DBCC](#)

## DBCC FOR MICROSOFT SQL SERVER

The DBCC (Database Consistency Check) dialog box lets you specify single or multiple tables or indexes for validation in Microsoft SQL Server. Use this dialog box to perform table-level or index-level validations of databases which are too large to undergo database-level DBCC operations in a time-efficient manner.

The DBCC dialog box includes the following elements:

- A window displaying the target database objects
- A drop-down list of DBCC Operations
- Buttons for previewing the operation's SQL code, scheduling the operation, and executing the operation

## Important Notes

None

For more information, see:

[DBCC for Databases](#)

[DBCC for Tables](#)

[DBCC for Indexes](#)

## DBCC FOR MICROSOFT SQL SERVER DATABASES

The DBCC dialog box for databases lets you perform database-wide validations. You should validate your databases as part of regular database maintenance to guard against corruption and failure. Microsoft SQL Server offers a set of DBCC commands to validate the integrity of your databases. Generally, you should perform these DBCC commands prior to dumping your databases to ensure that you are capturing clean backups of your databases.

The fundamental difference between the DBCC dialog box for databases, tables and indexes is the content of the DBCC Operation drop-down list.

The table below describes the options and functionality on the DBCC dialog box.

DBCC Operation	Description
Check Allocation	Executes a DBCC CHECKALLOC command. Makes sure that all data and index panels are correctly allocated and used. It reports on the amount of space allocated and used in the database. When checking allocation, you have the option to skip non-clustered indexes by selecting the Skip non-clustered indexes check box.
Check Catalog	Executes a DBCC CHECKCATALOG command. Checks for consistency in and between system tables.
Check Database	Executes a DBCC CHECKDB command. Verifies that all tables and indexes are properly linked, that indexes are in proper sorted order, that all pointers are consistent, that the data on each panel is reasonable, and that panel offsets are reasonable. When checking a database, you have the option to skip non-clustered indexes by selecting the Skip non-clustered indexes check box.

DBCC Operation	Description
Check FileGroup	Executes a DBCC CHECKFILEGROUP command. Verifies that all tables and indexes for the specified filegroup are properly linked, that indexes are in proper sorted order, that all pointers are consistent, that the data on each page is reasonable, and that page offsets are reasonable. When checking filegroups, you have the option to skip non-clustered indexes by selecting the Skip non-clustered indexes check box.
Show Oldest Transaction	Executes a DBCC OPENTRAN command. Displays information on the oldest active transaction and the oldest distributed and non distributed replicated transactions, if any, within the specified database.
Update Usage	Executes a DBCC UPDATEUSAGE command. Reports and corrects the rows, used, reserved, and data pages columns of the sysindexes table for any clustered indexes on objects of the type U (user-defined table) or S (system table).

For more information, see:

[Completing the DBCC Dialog Box](#)

[DBCC](#)

#### DBCC FOR MICROSOFT SQL SERVER TABLES

The DBCC dialog box for tables lets you perform table-level validations of databases. The fundamental difference between the DBCC dialog box for tables and indexes is the content of the DBCC Operation drop-down list.

The table below describes the options and functionality on the DBCC dialog box.

Option	Description
Check Current Identity Value	Checks the current identity value for the target objects, correcting values if needed depending on parameter specifications. Identity columns created with a NOT FOR REPLICATION clause in either the CREATE TABLE or ALTER TABLE statement are not corrected by this operation.

Option	Description
Check Fragmentation	<p>Displays the target table's data and index fragmentation information, determining whether the table is heavily fragmented.</p> <p>When a table is heavily fragmented, you can reduce fragmentation and improve read performance by dropping and recreating a clustered index (without using the SORTED_DATA option). Doing so reorganizes the data, resulting in full data pages. To adjust the level of fullness, use the Rebuild Index operation's FILLFACTOR option. When INSERT, UPDATE, and DELETE statements fragment tables, they usually do so with unequal distribution across the entire database so that each page varies in fullness over time, forcing additional page reads for queries that scan part or all of a table.</p>
Check Table	<p>Checks the linkages and sizes of text, ntext and image pages for selected tables. For the data, index, text, ntext, and image pages of the target tables, this operation also checks that index and data pages are correctly linked, indexes are in their proper sorted order, pointers are consistent, the data on each page is reasonable, and the page offsets are reasonable. DBCC CHECKTABLE requires a shared lock on all tables and indexes in the database for the duration of the operation. However, DBCC CHECKTABLE does not check the allocations of pages in the specified table (for this, use DBCC CHECKALLOC). To perform DBCC CHECKTABLE on every table in the database, use DBCC CHECKDB.</p>
Check Text/Image Allocation	<p>NOTE: This option is not available Microsoft SQL Server version 7 or later.</p> <p>Checks the allocation of text, ntext, or image columns for a table. In later versions of Microsoft SQL, use DBCC CHECKTABLE to check the integrity of the data, index, text, ntext, and image pages for the target table.</p>
Pin Table	<p>Pins target tables in memory so that they are not flushed when Microsoft SQL Server needs space to read in new pages. DBCC PINTABLE is best used for keeping small, frequently referenced tables in memory.</p> <p>Pinning a large table can consume a large portion of the buffer cache, leaving inadequate memory to service other tables in the system. A pinned table that is larger than the buffer cache itself can fill the entire cache, necessitating a shut down of the system by a sysadmin user, who must then restart Microsoft SQL Server and unpin the table. Pinning too many small tables can result in a similar problem.</p>
Rebuild Index	<p>Dynamically rebuilds one, multiple, or all indexes for a table in the target database, allowing indexes which enforce either primary key or unique constraints to be rebuilt without need for dropping and recreating. This operation is not supported for use on system tables.</p>
Unpin Table	<p>Marks target tables as unpinned, rendering their pages flushable from the buffer cache if space is needed to read in a new page from disk.</p>

Option	Description
Update Usage	Reports and corrects inaccuracies in the sysindexes table (which can result in incorrect space usage reports by the sp_spaceused system stored procedure) and corrects the rows, used, reserved, and dpages columns of the sysindexes table for tables and clustered indexes. If there are no inaccuracies in sysindexes, DBCC UPDATEUSAGE returns no data. Use this operation to synchronize space-usage counters. Executing this operation on large tables or databases can require some time, so it should typically be used only when you suspect incorrect values returned by sp_spaceused.

For more information, see:

[Completing the DBCC Dialog Box](#)

[DBCC](#)

#### DBCC FOR MICROSOFT SQL SERVER INDEXES

The DBCC dialog box for indexes lets you perform index-level validations of databases. The fundamental difference between the DBCC dialog box for tables and indexes is the content of the DBCC Operation drop-down list.

The table below describes the options and functionality on the DBCC dialog box.

Option	Description
Check Fragmentation	Displays the target index's table data and fragmentation information, determining whether the table is heavily fragmented. For more information, see tip below.  When a table is heavily fragmented, you can reduce fragmentation and improve read performance by dropping and recreating a clustered index (without using the SORTED_DATA option). Doing so reorganizes the data, resulting in full data pages. To adjust the level of fullness, use the Rebuild Index operation's FILLFACTOR option. When INSERT, UPDATE, and DELETE statements fragment tables, they usually do so with unequal distribution across the entire database so that each page varies in fullness over time, forcing additional page reads for queries that scan part or all of a table.
Check Index	Checks the linkages and sizes of text and image pages for selected indexes. DBCC CHECKTABLE requires a shared lock on all tables and indexes in the database for the duration of the operation.
Rebuild Index	Dynamically rebuilds all target indexes, allowing those which enforce either primary key or unique constraints to be rebuilt without need for dropping and recreating. This operation is not supported for use on system tables.

Option	Description
Show Statistics	Displays the current distribution statistics for the target indexes. The results returned indicate the selectivity of each target index (a lower density equals a higher selectivity) and provide the basis for determining the usefulness of target indexes to the optimizer.
Update Usage	Reports and corrects inaccuracies in the sysindexes table (which can result in incorrect space usage reports by the sp_spaceused system stored procedure) and corrects the rows, used, reserved, and dpages columns of the sysindexes table for tables and clustered indexes. If there are no inaccuracies in sysindexes, DBCC UPDATEUSAGE returns no data. Use this operation to synchronize space-usage counters. Executing this operation on large tables or databases can require some time, so it should typically be used only when you suspect incorrect values returned by sp_spaceused. Additional Options: Update Index Option, and DBCC General Option.

For more information, see:

[Completing the DBCC Dialog Box](#)

[DBCC](#)

## DBCC FOR SYBASE ASE

The DBCC (Database Consistency Check) dialog box lets you specify single or multiple databases, tables or indexes for validation in Sybase ASE. Use this dialog box to perform table-level or index-level validations of databases which are too large to undergo database-level DBCC operations in a time-efficient manner.

The DBCC dialog box includes the following elements:

- A window displaying the target database objects
- A drop-down list of DBCC Operations
- Buttons for previewing the operation's SQL code, scheduling the operation, and executing the operation

For more information, see:

[DBCC for Tables](#)

[DBCC for Indexes](#)

[DBCC for Databases](#)

## DBCC FOR SYBASE ASE DATABASES

The DBCC dialog box for databases lets you perform database-wide validations. The fundamental difference between the DBCC dialog box for databases, tables and indexes is the content of the DBCC Operation drop-down list.

The table below describes the options and functionality on the DBCC dialog box.

Option	Description
Check Allocation	Checks the allocation and use of all pages in the target database.
Check Catalog	Checks for consistency in and between system tables in the target database.
Check Database	Checks the allocation and structural integrity of all the objects in the target database.
Check Storage	Checks the target database for allocation, OAM page entries, page consistency, text valued columns, allocation of text valued columns, and text column chains. The results of this operation are stored in the dbccdb database.
Database Repair	Drops a damaged database.

For more information, see:

[Completing the DBCC Dialog Box](#)

[DBCC](#)

#### DBCC FOR SYBASE ASE TABLES

The DBCC dialog box for tables lets you perform table-level validations of databases. The fundamental difference between the DBCC dialog box for tables and indexes is the content of the DBCC Operation drop-down list.

The table below describes the options and functionality on the DBCC dialog box.

Option	Description
Check Allocation	Checks the database to see that every page is correctly allocated, and that no allocated page is unused. Use TABLEALLOC frequently (daily) to check page linkages in the Adaptive Server before performing a database dump to ensure the integrity of the dumped data.



Option	Description
Check Table	Checks the linkages and sizes of text, ntext and image pages for selected tables. For the data, index, text, ntext, and image pages of the target tables, this operation also checks that index and data pages are correctly linked, indexes are in their proper sorted order, pointers are consistent, the data on each page is reasonable, and the page offsets are reasonable. DBCC CHECKTABLE requires a shared lock on all tables and indexes in the database for the duration of the operation. However, DBCC CHECKTABLE does not check the allocations of pages in the specified table (for this, use DBCC CHECKALLOC). To perform DBCC CHECKTABLE on every table in the database, use DBCC CHECKDB.
Check Text	Upgrades text values after you have changed an Adaptive Server's character set to a multibyte character set.
Rebuild Index	Dynamically rebuilds one, multiple, or all indexes for a table in the target database, allowing indexes which enforce either primary key or unique constraints to be rebuilt without need for dropping and recreating. This operation is not supported for use on system tables.

### DBCC Operation Options

Rapid SQL offers additional options for selected operations which you can specify to further customize a database consistency check. The table below describes each option:

Option	Description
Error Option	Click Fix Error to instruct Rapid SQL to fix any allocation errors it finds. You must put your database in single-user mode to fix errors, so specify this option during times of low usage.
Job Scope	Select Optimize to produce a report based on the allocation pages listed in the object allocation map (OAM) pages for the table. It does not report and cannot fix unreferenced extents on allocation pages that are not listed in the OAM pages. The optimized option is the default. Select Full to perform the equivalent of a table-level CHECKALLOC, reporting all types of allocation errors. Select Fast to produce an exception report of pages that are referenced but not allocated in the extent. Fast does not produce an allocation report.
Update Index Option	Click this check box to skip non-clustered indexes when updating index options.

For more information, see:

[Completing the DBCC Dialog Box](#)

[DBCC](#)

## DBCC FOR SYBASE ASE INDEXES

The DBCC dialog box for indexes lets you perform index-level validations of databases. Unlike the DBCC dialog box for tables, this DBCC dialog box offers only one option on the DBCC Operation drop-down list: Check Allocation. This option checks the specified database to see that all pages are correctly allocated and that no allocated page is unused.

The table below describes the options and functionality on the DBCC dialog box.

Option	Description
DBCC Option	Checks the specified database to see that all pages are correctly allocated and that no page that is allocated is not used.
Error Option	Rapid SQL to fixes any allocation errors it finds. You must put your database in single-user mode to fix errors, so specify this option during times of low usage
Job Scope	Produces a report based on the allocation pages listed in the object allocation map (OAM) pages for the table. It does not report and cannot fix unreferenced extents on allocation pages that are not listed in the OAM pages. The optimized option is the default. A full job is the equivalent to a table-level CHECKALLOC, reporting all types of allocation errors. A fast job does not produce an allocation report, but produces an exception report of pages that are referenced but not allocated in the extent.

For more information, see:

[Completing the DBCC Dialog Box](#)

[DBCC](#)

## DEALLOCATE UNUSED SPACE

This functionality is available for Oracle only.

The Deallocate Unused Space dialog box lets you deallocate space from clusters, indexes, and tables. You can also deallocate unused space from table partitions and subpartitions. When you find that allocated space is not being used, you can free that space for use by other objects by explicitly deallocating space. Oracle releases the freed space to the user quota for the tablespace in which the deallocation occurs.

Oracle deallocates unused space from the end of the object toward the high water mark. In other words, Oracle frees space starting with the space that would have been used last. If an extent is completely contained in the space to be deallocated, then the whole extent is freed. If an extent is only partially contained in the space to be deallocated, then Oracle shrinks that extent to the size of the used space up to the high water mark, and frees the unused space in the extent.

If you are deallocating unused space from an index and the index is range-partitioned or hash-partitioned, Oracle deallocates unused space from each partition in the index. If an index is a local index on a composite-partitioned table, Oracle deallocates unused space from each of the subpartitions in the index.

### Important Notes

None

The table below describes the options and functionality on the Deallocate Unused Space dialog box:

Option	Description
Specify the number of bytes above the high-water mark that the objects will have after deallocation. If no value is specified, all unused space will be freed.	<p>If you do not specify an amount of unused space and the high water mark is above the size of INITIAL and MINEXTENTS, then all of the unused space is freed. If the high water mark is less than the size of INITIAL or MINEXTENTS, then all unused space above MINEXTENTS is freed.</p> <p>If you specify an amount of unused space and the remaining number of extents is smaller than MINEXTENTS, then the MINEXTENTS value changes to reflect the new number. If the initial extent becomes smaller as a result of the deallocation, the INITIAL value changes to reflect the new size of the initial extent.</p>

**TIP:** You can verify that the deallocated space is freed by going to the Space Tab in the appropriate Rapid SQL editor.

For more information, see [Completing the Deallocate Unused Space Dialog Box](#).

## COMPLETING THE DEALLOCATE UNUSED SPACE DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Tables node.

Rapid SQL displays the tables in the Database Explorer.

- 2 On the Database Explorer, right-click the table, and then select Deallocate Unused Space.  
Rapid SQL opens the Deallocate Unused Space dialog box.
- 3 Specify the number of bytes above the high-water mark that the objects will have after deallocation. If no value is specified, all unused space will be freed.
- 4 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

## DEPENDENCIES

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, InterBase/Firebird, Oracle, SQL Server, and Sybase.

The Dependencies window displays and lets you modify the referencing objects for an object.

### Important Notes

None

For more information, see [Viewing Dependencies](#).

## VIEWING DEPENDENCIES

To open the Dependencies Window, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Dependencies.  
Rapid SQL opens the Dependencies Window.

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

## DETACH/ATTACH

This functionality is available for Microsoft SQL Server 2000 or later only.

The Detach/Attach dialog box lets you detach and reattach data and transaction log files. Detaching a database removes the database from the server but leaves the database intact within the data and transaction log files that compose the database. These data and transaction log files can then be used to attach the database to any instance of Microsoft SQL Server, including the server from which the database was detached. This makes the database available in exactly the same state it was in when it was detached. The Rapid SQL commands are:

- [Detach](#)
- [Attach](#)

### Important Notes

None

## DETACHING A DATABASE

The table below describes the options and functionality on the Detach dialog box.

Option	Description
Skip Checks	Select to skip the UPDATE STATISTICS operation when detaching the database. This option is useful for databases that are to be moved to read-only media.

For more information, see [Completing the Detach Dialog Box](#).

### COMPLETING THE DETACH DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Databases node.  
Rapid SQL displays the Databases in the Database Explorer.
- 2 On the Database Explorer, right-click the database, and then select Detach.  
Rapid SQL opens the Detach dialog box.
- 3 To skip the UPDATE STATISTICS operation when detaching the database, select the *Skip Checks* check box for the target database(s).

**TIP:** This option is useful for databases that are to be moved to read-only media.

- 4 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

For more information, see:

[Detach](#)

## ATTACHING A DATABASE

The table below describes the options and functionality on the Attach dialog box.

Option	Description
Database Name to be Attached	Lets you type the name of the target database to attach the file(s) to.
Add Button	Click to add database file(s), and then type the name of the MDF (master data file) of the database to attach.

For more information, see:

[Completing the Attach Dialog Box](#)

### COMPLETING THE ATTACH DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Databases node.  
Rapid SQL displays the Databases in the Database Explorer.
- 2 On the Database Explorer, right-click the Database, and then select Attach.  
Rapid SQL opens the Attach dialog box.
- 3 In the Database Name to be Attached box, type the database name.

4 In the grid, do one of the following:

#### Microsoft SQL Server 7.0

- Select the target database file(s).
- To add the \*.mdf file(s), click Add, enter the name of the MDF (master data file) of the database to attach, and then click OK.
- To add the \*.ldf file(s), click Add, enter the name of the LDF (log data file) of the database to attach, and then click OK.
- To drop database file(s), click Drop and then select the target file(s).

#### Microsoft SQL Server 2000

- Select the target database file(s).
- To add database file(s), click Add and then enter the name of the MDF (master data file) of the database to attach.

Rapid SQL automatically adds the appropriate \*.ldf file.

5 To drop database file(s), click Drop and then select the target file(s).

6 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

For more information, see:

[Attach](#)

## DISABLE INDEX

The **Disable Indexes** dialog box lets you mark an index as disabled or inactive and unavailable for use.

**NOTE:** This functionality is available for InterBase/Firebird and Microsoft SQL Server 2005 and above.

**NOTE:** On SQL Server this action can also be applied to primary keys and unique keys.

**To disable an index**

- 1 On the *Datasource Explorer*, expand nodes until index nodes are displayed.
- 2 Right-click an index and select **Disable Index** from the context menu. The **Disable Index** wizard opens.
- 3 Use the following table as a guide to understanding and modifying settings in the wizard:

Step	Settings and tasks
<b>Action options</b>	Displays the name of the index, primary key, or unique key being disabled.
<b>Dependencies</b>	Review the referring and referred objects that will be automatically resolved when you execute this operation.
<b>Preview</b>	Preview the DDL generated for the operation and when ready, use the <b>Schedule</b> or <b>Execute</b> button to perform this action.

- 4 Click **Execute**. For information on the other options, see [Preview](#) and [Scheduling](#).

**DISABLE/ENABLE JOB QUEUES**

This functionality is available for Oracle only.

The Disable/Enable dialog box lets you enable or disable any job queue. Job Queues are built-in mechanisms that let you schedule a variety of SQL-based or command-line driven tasks.

**Important Notes**

None.

For more information, see [Completing the Enable/Disable Dialog Box](#).

**COMPLETING THE ENABLE/DISABLE DIALOG BOX FOR JOB QUEUES**

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Job Queues node.  
Rapid SQL displays the job queues in the Database Explorer.
- 2 On the Database Explorer, right-click the target job queue, and then select Enable or Disable.  
Rapid SQL opens the Enable or Disable dialog box.
- 3 Do one of the following:



Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

## DISABLE/ENABLE TRIGGERS

This functionality is available for Microsoft SQL Server, Oracle, and Sybase ASE only.

The Disable/Enable dialog box lets you enable or disable any triggers.

Triggers are a special type of procedure that automatically fire when defined data modification operations (insert, update or delete) occur on a target table. Triggers fire after an insert, update or delete, but belong to the same transaction as the data modification operation.

Loading a database from a previous dump causes any triggers defined in the database to fire. To speed the time required to load a database you should disable triggers.

### Important Notes

Disabling triggers can lead to problems with maintaining referential integrity and business rules.

For more information, see [Completing the Enable/Disable Dialog Box](#).

## COMPLETING THE ENABLE/DISABLE DIALOG BOX FOR TRIGGERS

To complete this dialog box, do the following:

- 4 On the *Datasource Explorer* tool bar, click Command, and then select Enable or Disable.

OR

In the right pane of the application, right-click the target object, and then select Enable or Disable. On the *Database Explorer*, select the Triggers node.

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.

Rapid SQL displays the triggers in the Database Explorer.

- 5 On the Database Explorer, right-click the target trigger, and then select Enable or Disable.

Rapid SQL opens the Enable or Disable dialog box.

- 6 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

## DISK RESIZE

This functionality is available for Sybase ASE 12.5.01 only.

The Disk Resize dialog box lets you dynamically increase the size of database devices without having to add and configure additional database devices.

### Important Notes

None

The table below describes the options and functionality on the Disk Resize dialog box.

Option	Description
Increase current size by:	Lets you type the amount in server pages to increase the device(s). 512 pages equals 1MB.  NOTE: Whether the Sybase database is configured for 2K, 4K, or 8K page sizes does not impact the sizing.

For more information, see [Completing the Disk Resize Dialog Box](#).

## COMPLETING THE DISK RESIZE DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Database Devices node.  
Rapid SQL displays the Database Devices in the Database Explorer.
- 2 On the Database Explorer, right-click the database device, and then select Disk Resize.  
Rapid SQL opens the Disk Resize dialog box.
- 3 In Increase current size by: type the amount in server pages to increase the device(s). 512 pages equals 1MB.
- 4 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

## DROP

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Interbase/Firebird, Oracle, SQL Server, and Sybase.

The Confirm Drop dialog box lets you drop one or more database objects and remove their definition from the system catalog. The Confirm Drop dialog box lets you:

- Drop the object(s) immediately

- Schedule dropping the object(s)
- Preview the SQL to drop the object(s)
- Save the SQL to drop the object(s)
- Print the SQL to drop the object(s)
- Send an e-mail attachment with the SQL to drop the object(s)

### Completing the Confirm Drop Dialog Box

To complete this dialog box, do the following:

- 1 On the *Datasource Explorer*, select the target object node.  
Rapid SQL displays the target objects in the right pane of the Explorer window.
- 2 On the *Datasource Explorer* tool bar, click Command, and then select Drop.  
OR  
In the right pane of the application, right-click the target object, and then select Drop.  
RAPID SQL opens the **Drop/Confirm Drop** dialog box on the **Action Options** pane.
- 3 On the Action Options panel, select any options specific to the object type or configuration of the object you are dropping.
- 4 Navigate to the **Dependencies** pane to view database objects that are dependent on the object you are dropping.  
**NOTE:** For Microsoft SQL Server 2005, when dropping Logins, you also have the option to delete corresponding user objects.  
  
**NOTE:** For Microsoft SQL Server, when dropping Databases, you also have the option to delete the backup and restore history for the database.
- 5 Navigate to the **Preview** pane to verify that the correct DDL was created, and if necessary navigate backward and modify your choices.
- 6 Click **Execute** to drop the object.

### DROP CLONE

If you have created a clone from a base table, you can drop that clone. To determine whether a table has a clone, see the **Dependencies** tab/panel on a Table editor or on the wizard for object actions such as Drop or Rename.

**NOTE:** This functionality is available as of IBM DB2 for z/OS version 9.

### To drop the clones of one or more tables:

- 1 On the *Datasource Explorer*, expand the Tables node.  
Rapid SQL displays the tables in the right pane of the Explorer window.
- 2 In the right pane of the application, select one or more tables.
- 3 From the Command menu select Clone Actions > Drop Clone.  
Rapid SQL opens the **Confirm Drop Clone** dialog box.
- 4 Use the following table as a guide to understanding and modifying settings in this wizard:

Step	Settings and tasks
<b>Action Options</b>	Lets you review the action you initiated.
<b>Dependencies</b>	Lists referring and referenced objects potentially impacted by the change. For details, see <a href="#">Dependencies</a> .
<b>Preview</b>	Displays the DDL that will execute the object action. For details, see <a href="#">Preview</a> .

- 5 Click **Execute**. For information on the other options, see [Preview](#) and [Scheduling](#).

For related information, see the following topics:

- [Create Clone](#)
- [Exchange Data With Clone](#)

## EDIT DATA

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, InterBase/Firebird, Oracle, SQL Server, and Sybase.

The Edit Data function opens the [Data Editor](#). You can use the Data Editor to edit your tables in real-time. The Data Editor supports all editable datatypes and is an alternative way to add, edit, or delete data from your tables.

**NOTE:** You can use Data Editor within [Query Builder](#) to edit data in tables while you create SELECT statements. You can open multiple Data Editor sessions so that you can continue to change your data until you find the best match query.

The Data Editor includes a [Data Editor Filter](#) that lets you select the columns in your table that you want to edit. You must select at least one column to use the Data Editor. The Data Editor Filter is not available for the Query Builder.

## ERROR

The Error message displays a warning. After reading the warning, click Close to continue.

## ESTIMATE SIZE

This functionality is available for Oracle and Sybase ASE only.

Estimate Size dialog box for [tables](#) and [indexes](#) lets you estimate how large a table or index will become given a row growth projection. The results let you proactively plan your space-related object needs.

### Important Notes

None

## ESTIMATE TABLE SIZES

The table below describes the options and functionality on the Estimate Size dialog box for tables:

Option	Description
Database	Lets you select the table database.
Owner	Lets you select the table owner.
Table Name	Lets you select the table by name.
Number of Rows	Lets you type number of rows to estimate size. Click Calculator to open a calculator.
Estimate Table Size (KB)	Lets you type table size to estimate size.
Add Table	Click to add a table to the grid.
Remove Table	Click to remove a table from the grid.
Estimate Size	Click to start the estimate size functionality. Rapid SQL then calculates the information, and displays it in a browser-ready HTML report that you can print or save.
Update Statistics	Opens the <a href="#">Update Statistics</a> dialog box.
Save As	Lets you save as a result grid (.rsl) file.

**TIP:** To generate an HTML report, click the Report button.

## ESTIMATE INDEX SIZES

The table below describes the options and functionality on the Estimate Size dialog box for indexes:

Option	Description
Database	Lets you select the table database.
Owner	Lets you select the table owner.
Table Name	Lets you select the table.
Index Name	Lets you select the index.
Fill Factor	Lets you specify a percentage of how full each index page can become. Click Calculator to open a calculator.
Number of Rows	Lets you type number of rows to estimate size. Click Calculator to open a calculator.
Estimate Index Size (KB)	Lets you type index size to estimate size.
Add Index	Click to add a index to the grid.
Remove Index	Click to remove a index from the grid.
Estimate Size	Click to start the estimate size functionality. Rapid SQL then calculates the information, and displays it in a browser-ready HTML report that you can print or save.
Update Statistics	Opens the <a href="#">Update Statistics</a> dialog box.
Save As	Lets you save as a result grid (.rsl) file.

**TIP:** To generate an HTML report, click the Report button.

## EXCHANGE DATA WITH CLONE

This action lets you issue an EXCHANGE statement with the DATA BETWEEN TABLE *table1* AND *table2* syntax, letting you swap data between a base table and its clone. To determine whether a table has a clone, see the **Dependencies** tab/panel on a Table editor or on the wizard for object actions such as Drop or Rename.

**NOTE:** This functionality is available as of IBM DB2 for z/OS version 9.

### To exchange data between one or more tables and their clones:

- 1 On the *Datasource Explorer*, expand the Tables node.  
Rapid SQL displays the tables in the right pane of the Explorer window.
- 2 In the right pane of the application, select one or more tables.
- 3 From the Command menu select Clone Actions > Exchange Data With Clone.  
Rapid SQL opens the **Exchange Data With Clone** dialog box.

4 Use the following table as a guide to understanding and modifying settings in this wizard:

Step	Settings and tasks
<b>Action Options</b>	Lets you review the action you initiated.
<b>Dependencies</b>	Lists referring and referenced objects potentially impacted by the change. For details, see <a href="#">Dependencies</a> .
<b>Preview</b>	Displays the DDL that will execute the object action. For details, see <a href="#">Preview/</a> .

5 Click **Execute**. For information on the other options, see [Preview](#) and [Scheduling](#).

For related information, see the following topics:

- [Create Clone](#)
- [Drop Clone](#)

## EXECUTE

This functionality is available for InterBase/Firebird, Oracle and Sybase ASE only.

The Execution dialog box lets you execute [functions](#) and [procedures](#). The Execution dialog box displays the necessary parameters. Depending on the parameter datatype, you can use the built-in calculator and calendar to select the entered data.

Rapid SQL lets you enter arguments and set execution options in the Execution dialog box.

### Important Notes

None

For more information, see:

[Executing Functions](#)

[Executing Procedures](#)

[Completing the Execution Dialog Box](#)

## COMPLETING THE EXECUTION DIALOG BOX

To complete this dialog box, do the following:

**TIP:** The [Code Analyst](#) is a tool to identify time-consuming lines of code. Code Analyst lets you perform detailed response time analysis on the execution of Procedures and Functions.

1 On the *Database Explorer*, select the target object node.

Rapid SQL displays the target objects in the Database Explorer.



- 2 On the Database Explorer, right-click the target object, and then select Execute.

Rapid SQL opens the Execute dialog box.

- 3 To use the calendar, click the Down arrow.
- 4 To use the calculator, click the Down arrow.

**NOTE:** These built-in functions are context-sensitive to the type. of parameter required

- 5 To execute the function with show plan, click Query Plan.

**TIP:** [Code Analyst](#) is a tool to identify time-consuming lines of code. Code Analyst lets you perform detailed response time analysis on the execution of Procedures and Functions.

- 6 To modify your results options, click Results.
- 7 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

For more information, see:

[Executing Functions](#)

Executing Procedures

[Execute](#)

## EXECUTING FUNCTIONS

The table below describes the options and functionality on the Execution dialog box:

Option	Description
Query Plan	<p>Click the Query Plan button to activate the Oracle EXPLAIN PLAN functionality. By default, Rapid SQL stores the EXPLAIN PLAN output in a table named EMBARCADERO_EXPLAIN_PLAN.</p> <p>Rapid SQL's Query Plan window now displays data for the estimated costs, number of rows, and bytes returned by each plan step.</p> <p>The <a href="#">Code Analyst</a> is a tool to identify time-consuming lines of code. Code Analyst lets you perform detailed response time analysis on the execution of Procedures and Functions.</p>
Result Options	Click the Result Options button lets you set options for the current execution. For more information on Result Options, see the Options Editor.
Arguments	Enter arguments in the blank text box. The datatype to the left of the blank field lets you know what type of argument is expected. For numeric arguments, you can use the calculator to enter information. For date arguments, you can use the built-in calendar or you can override the calendar by selecting the check box and entering your argument in field that follows. NULL is the default argument in that field.
Default	Select to gather default information from the data dictionary. ORACLE ONLY: You can not specify non-default parameters after specifying a default parameter.

For more information, see:

[Completing the Execution Dialog Box](#)

## EXECUTING PROCEDURES

The Procedure Execution dialog box lets you:

- Save input parameters as \*.prm files to preserve specific input parameter configurations.
- Open \*.prm files to save the effort of reentering specific input parameters.
- Reset parameters to their default setting.

The table below describes the options and functionality of the Procedure Execution dialog box:

Option	Description
Owner	Displays the current procedure's owner.
Procedure	Displays the name of the current procedure.

Option	Description
Parameter	Specify the required input parameters in this window. If input parameters are not required for the execution of the target procedure, a message displays in this window, stating that the procedure "has no input parameters. Press execute to run it."
Open	Click to open an Open dialog box, from which you can open an existing *.prm file. The saved parameters immediately populate the dialog box upon opening.
Save	Click to save the values of your input parameters as a *.prm file. You can reopen a saved *.prm file from this dialog box at any time.
Reset	Click to reset the parameters in the Parameter window to their default values.
Default	Select to gather default information from the data dictionary. ORACLE ONLY: You can not specify non-default parameters after specifying a default parameter.
Continue	Click to execute the procedure once you have entered values for all required parameters in the Parameter window.

For more information, see:

[Completing the Execution Dialog Box](#)

## EXTRACT

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Interbase/Firebird, Oracle, SQL Server, and Sybase.

Rapid SQL lets you extract data from one database to another database and extract the statements required to create objects into an Interactive SQL window. You can extract an object from the Datasource Explorer or the Objects Editor.

## Important Notes

None

For more information, see:

[Extracting Data](#)

## EXTRACTING DATA

To extract data, do the following:

- 1 On the *Database Explorer*, click the *target object* node, and then click the target object.
- 2 In the Explorer window, right-click the target object, and then click *Extract*.

Rapid SQL opens the DDL Editor.

## EXTRACT DATA AS XML

**NOTE:** This functionality is available for Oracle 9i and SQL Server 8.0.

- 1 On the Datasource Explorer, expand the target datasource.
- 2 Expand the Tables node.
- 3 In the right pane of the Explorer window, right-click any table listed, and then select Extract Data as XML.
- 4 Select the columns to include in the Insert statement.
- 5 You can also filter what rows are included by adding your own Select statement.
- 6 Click OK.

The resulting XML document is created and presented in an active XML Editor. At this point the document can be saved in XML format.

## FLUSH CACHE

This functionality is available for IBM DB2 LUW for Linux, Unix, and Windows 8.1 only.

The Flush Cache dialog box lets you clear all dynamic SQL in the cache and forces IBM DB2 LUW for Linux, Unix, and Windows to recompile the SQL the next time it is called.

### Important Notes

None

For more information, see [Completing the Flush Cache Dialog Box](#).

## COMPLETING THE FLUSH CACHE DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.

Rapid SQL displays the target objects in the Database Explorer.

- 2 On the Database Explorer, right-click the target object, and then select Flush Cache.

Rapid SQL opens the Flush Cache dialog box.

- 3 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.

## FREE PACKAGES

This functionality is available for IBM DB2 LUW for OS/390 and z/OS only.

The Free Packages dialog box lets you delete a specific version of a package, all versions of a package, or whole collections of packages.

**CAUTION:** Confirmation is required because this is a destructive action.

To complete the Free Plans dialog box, do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Execute	Executes the task.

## FREE PLANS

This functionality is available for IBM DB2 LUW for OS/390 and z/OS only.

The Free Plans dialog box lets you delete application plans. A Plan is an executable application created in the bind process. It can include one or more packages or debris.

### Important Notes

Confirmation is required because this is a destructive action.

For more information, see [Completing the Free Plans Dialog Box](#).

## COMPLETING THE FREE PLANS DIALOG BOX

To complete the Free Plans dialog box, do one of the following:

- 1 On the *Database Explorer*, select the Plans node.  
Rapid SQL displays the plans in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Free.  
Rapid SQL opens the Free Plans dialog box.
- 3 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Execute	Executes the task.

## GENERATE ANONYMOUS BLOCK

This functionality is available for Oracle only.

An anonymous block is an alternative to a stored procedure or function. An anonymous block has no explicit name and, unlike a procedure or function, is not stored by in the Oracle server database as a schema object. This means that Oracle server compiles and executes only the current instance of the anonymous block. Anonymous blocks are normally sent to the Oracle server from another application and require compilation before execution. Since functions and procedures are stored by Oracle server, they exist in a precompiled state and require no explicit compilation before execution.

Rapid SQL lets you generate anonymous blocks from existing functions and procedures. Rapid SQL produces the code appropriate for an anonymous block in the ISQL Window and provides you with the normal ISQL Window options to check, manipulate, and execute the SQL.

### Important Note

Rapid SQL also lets you generate anonymous blocks for functions and procedures that are part of a package.

## GENERATING PACKAGES, PROCEDURES, AND STATEMENTS FROM TABLES AND VIEWS

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, InterBase/Firebird, Oracle, SQL Server, and Sybase.

Rapid SQL lets you generate simple packages, procedures, and statements for selected tables and lets you generate a simple select statement for views. Because the packages and procedures generated by Rapid SQL are rudimentary, they are intended merely as a starting point and should be modified to reflect your specific needs. Rapid SQL opens the generated statements in the ISQL Window.

When Rapid SQL creates a package from a table, it generates a series of procedures designed to emulate the typical variety of procedures in a package. Rapid SQL lets you choose the IN and OUT columns and generates the procedures based on your selections. Rapid SQL creates procedures and statements from tables in the same way.

### Important Notes

None

For more information, see:

[Generate Packages, Procedures, and Statements](#)

[Generate Select Statement](#)

## GENERATE PACKAGES, PROCEDURES, AND STATEMENTS

Rapid SQL opens this dialog box when you want to generate code for an Insert, Update, or Delete statement. This dialog box lets you specify the columns you want to include in the generation of an Insert, Update, or Delete statement or procedure.

## GENERATE SELECT STATEMENT

Rapid SQL opens this dialog box when you want to generate code for a Select statement. This dialog box includes two panes: one for specifying Input Columns, one for specifying Output Columns.

**NOTE:** This dialog box mirrors the functionality of Rapid SQL's Embarcadero Code Generator dialog box, which is accessible from the application's Tools menu and Tools toolbar.

The Select 1 or More Columns dialog box lets you specify the IN and OUT columns for a package or procedure, or the columns to select for a Select statement and any associated WHERE clause. Rapid SQL uses the Input Columns to generate the WHERE clause.

## IMPACT ANALYSIS

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Impact Analysis dialog box displays the objects dependent on the object you are changing. You can use the Impact Analysis dialog box to review the objects that could be impacted should you proceed with the action in question.

**Important Notes**

None

**INDEX CONSTRAINT**

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Index Constraint dialog box lets you create constraints.

**Important Notes**

None

The table below describes the options and functionality on the Index Constraint dialog box:

**NOTE:** Options differ by platform.

Option	Description
Constraint Name	Lets you type the name of the constraint, which can be up to 30 characters long and must be unique across the entire database.
Properties	Lets you select options:  Clustered - A cluster comprises of a group of tables that share the same data blocks, and which are grouped together because they share common columns and are often used together.  Filegroup - Lets you select the filegroup within the database the constraint is stored.  Fill Factor - Lets you specify a percentage of how large each constraint can become.
Specify Columns in Constraint	Lets you click the key columns, and then click the right arrow button to move them to the Constraint Columns grid.
Columns Button	Click to open the <a href="#">Paste Columns for Check Constraint dialog box</a> .

**LOAD JAVA**

**NOTE:** This functionality is available for Oracle 8i or later.

Before you can call Java stored procedures, you must load them into the Oracle database and publish them to SQL.

The Java Load Wizard lets you:

- Select the owner of the Java object and the files to load.



- Select options for the loading of the files.
- Select Resolver options.

### Important Notes

None

For more information, see:

[Java Load Wizard - Panel 1](#)

### JAVA LOAD WIZARD - PANEL 1

The table below describes the options of the first panel of the Java Load Wizard.

Option	Description
Who owns the Java Object?	Lets you select the owner of the Java object.
Select files to be loaded	Select a file, and then click Add.

For more information, see:

[Java Load Wizard](#)

### JAVA LOAD WIZARD - PANEL 2

The table below describes the options of the second panel of the Java Load Wizard.

Option	Description
When do you want the Java files to be resolved?	Lets you specify when the source file is loaded as a source schema object, the source file is compiled, class schema objects are created for each class defined in the compiled .java file, and the compiled code is stored in the class schema objects.
Select the Encoding Options	Lets you specify the encoding of the .java file.
Grant Access to the following users	Lets you select one or more users.

For more information, see:

[Java Load Wizard](#)

## JAVA LOAD WIZARD - PANEL 3

The table below describes the options of the third panel of the Java Load Wizard.

Option	Description
Other Load Options	OPTIONAL: Lets you select options.
Add Resolver Options	<p>Lets you specify the objects to search within the schemas defined.</p> <p>Add - Click to open the Select a Resolver Option dialog box to add a new resolver option in the list.</p> <p>Edit - Click to open the Resolver Edit dialog box to modify a resolver option.</p> <p>Remove - Select one or more resolver option and click to delete.</p>

For more information, see:

[Java Load Wizard](#)

## LOCK

This functionality is available for IBM DB2 LUW for Linux, Unix, and Windows only.

The Lock Table dialog box lets you lock tables to prevent other users from reading or updating the table data. Locking a table saves the time of locking each row of the table that needs to be updated. Rapid SQL releases locks at the end of a transaction.

### Important Notes

None

The table below describes the options and functionality on the Lock Table dialog box.

Option	Description
Lock Mode	<p>Share - Lets other users view but not modify the table data.</p> <p>Exclusive - Prevents other users from viewing or modifying the table data.</p>

For more information, see [Completing the Lock Table Dialog Box](#).

## COMPLETING THE LOCK TABLE DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Tables node.  
Rapid SQL displays the Tables in the Database Explorer.
- 2 On the Database Explorer, right-click the table, and then select Lock.  
Rapid SQL opens the Lock dialog box.
- 3 Select a Lock Mode option:
  - Share - Lets other users view but not modify the table data.
  - Exclusive - Prevents other users from viewing or modifying the table data.
- 4 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

## MODIFY CONSTRAINT

The table below describes the options and functionality on the Modify Constraint dialog box:

Option	Description
Owner	Displays owner.
Table	Displays table.
Specify Columns in Constraint	Lets you select the columns and click the right arrow to include in constraint.

## MOVE LOG

This functionality is available for Microsoft SQL Server and Sybase ASE only.

The Move Transaction Log dialog box lets you move a transaction log from one device to another.

## Important Notes

None

The table below describes the options and functionality on the Move Transaction Log dialog box.

Option	Description
New Device	Let you select the device to which you want to move the transaction log.

For more information, see [Completing the Move Log Dialog Box](#).

## COMPLETING THE MOVE LOG DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Databases node.  
Rapid SQL displays the Databases in the Database Explorer.
- 2 On the Database Explorer, right-click the database, and then select Move Log.  
Rapid SQL opens the Move Log dialog box.
- 3 In New Device, select the device to which you want to move the transaction log.
- 4 Do one of the following:

## OPEN

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

Rapid SQL stores information about server and object types in object editors. An editor is a tabbed dialog box that groups related information together.

Each editor contains a context-sensitive Commands menu with pertinent functionality for the object. Many of the object editors contain a DDL Tab. This tab displays the underlying target object's SQL.

**TIP:** The Object Editor tool bar has a refresh button, that lets you refresh Object Editor contents, clear the Object Editors, and log SQL if SQL Logging is turned on.

If an objects has dependent objects, such as tables, triggers, procedures and views, you can view and access their dependent objects in the editor.

For more information, see [Modifying objects using editors](#)

## PACKAGE EDITOR

This functionality is available for IBM DB2 LUW for OS/390 and z/OS only.

The Package Editor lets you modify a package.

### Important Notes

None

The table below describes the options and functionality on the Package Editor:

Option	Description
Plan Name	Displays the plan name.
Location	Lets you select the name of the location to connect to.
Collection	Lets you select the location of the DBMS where the plan binds and where the description of the plan resides.
Package	Lets you select a package.
Add	Click to add the plan.
Close	Click to close the editor and return to the wizard.

## PARTITIONED COLUMNS

This functionality is available for Oracle only.

The Partitioned Columns dialog box lets you partition columns. Partitioning lets you break large columns into smaller pieces, which are called partitions. Partitions make the data in your table easier to manage and analyze. Your SQL statements can access the partitions rather than the entire table. Partitions are most useful in data warehouse applications, which store large amounts of data.

**Important Notes**

None

The table below describes the options and functionality on the Partitioned Columns dialog box.

Option	Description
Available Columns	Select column and click the right arrow button to move to Partitioning Columns.
Partitioning Columns	Select column and click the left arrow to move to Available Columns.
Apply	Click to apply changes.

**PARTITION EDITOR**

The Partition Editor lets you specify partition management and partition options.

**Important Notes**

None

The table below describes the options and functionality on the Partition Editor.

Option	Description
Partition number	Identifies a partition of the object.
DB2 manages the data sets on a volume of the specified storage group	<p>A DB2-managed data set that resides on a volume of the identified storage group.</p> <p>Stogroup Name - Lets you specify the stogroup.</p> <p>Primary Space Allocation - Lets you specify the minimum primary space allocation for a DB2-managed data set of the partition.</p> <p>Secondary Space Allocation - Lets you specify the minimum secondary space allocation for a DB2-managed data set of the partition.</p>
User manages the data sets on a specified VCAT catalog-name	A user-managed data set with a name that starts with catalog-name.
Free space portion of each page	Lets you specify what percentage of each page to leave as free space when the object is loaded or reorganized.
Free page frequency	Lets you specify how often to leave a page of free space when the object is loaded or reorganized. a number between 0 and 255.

Option	Description
Pages in group buffer pool	In a data sharing environment, lets you specify what pages of the partition are written to the group buffer pool in a data sharing environment.

## PARTITION UPPER BOUND

This functionality is available for Oracle only.

The Partition Upper Bound dialog box lets you specify the non-inclusive upper bound for the partitioning columns.

### Important Notes

None

The table below describes the options and functionality on the Partition Upper Bound dialog box.

Option	Description
Partitioning Column	Lets you select the partitioning column.
Upper Bound	Lets you select the default or type the value of the upper bound. All rows in the partition have partitioning keys that compare to less than the partition bound and greater than or equal to the partition bound for the preceding partition.  NOTE: The set button is not available if your selection is higher than the upper bound for the column in the original partition. If the upper bound is MAXVALUE, you need to type a value in the list.

The table below describes when you should enable editing of the upper bound:

Partition Type	Index	Table	Notes
Hash	N/A	N/A	Does not have upper bound.
Range Local	Disable	Enable	
Range Global	Enable all but the last partition	Enable	The upper bound button is disabled if there is only one partition.
Composite Global	Enable all but the last partition	Enable	The upper bound button is disabled if there is only one partition.

For more information, see [Completing the Partition Upper Bound Dialog Box](#).

## COMPLETING THE PARTITION UPPER BOUND DIALOG BOX

To complete this dialog box, do the following:

- 1 Click the *Partitioning Column* list and then click the partitioning column.
- 2 In the *Upper Bound* list, click the default or type the value of the upper bound. All rows in the partition will have partitioning keys that compare to less than the partition bound and greater than or equal to the partition bound for the preceding partition.
- 3 Click *Set*.
- 4 Click *OK*.

For more information, see:

[Partition Upper Bound](#)

## PASTE COLUMNS FOR CHECK CONSTRAINT

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Paste Columns for Check Constraint dialog box lets you paste selected column(s) in the check constraint definition.

### Important Notes

None

The table below describes the options and functionality on the Paste Columns for Check Constraint dialog box.

Option	Description
Columns	Lets you select column(s) to include in the check constraint definition.
Paste Button	Click to paste selected column(s) in the check constraint definition.

For more information, see:

[Completing the Paste Columns for Check Constraint Dialog Box](#)



## COMPLETING THE PASTE COLUMNS FOR CHECK CONSTRAINT DIALOG BOX

To complete this dialog box, do the following:

- 1 In the *Columns* grid, click the columns you want to paste or click the Select All button.
- 2 When you finish selecting columns, click *Paste*.

For more information, see:

[Paste Columns for Check Constraint](#)

## PLACE

**NOTE:** Place functionality is available for Sybase ASE.

The Placement dialog box lets you place tables and indexes on different segments. From a performance standpoint it is not recommended to have a table and its supporting indexes on the same device or disk segment. It is also good to have more frequently accessed indexes and tables grouped together on higher speed devices, if possible.

### Important Notes

None

The table below describes the options and functionality on the Placement dialog box.

Option	Description
New Segment Location	Lets you specify the segment on which you can place objects, the default, logsegment or system.

For more information, see [Completing the Placement Dialog Box](#).

## COMPLETING THE PLACEMENT DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Place.  
Rapid SQL opens the Placement dialog box.
- 3 In New Segment Location, select the segments to place the object.
- 4 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

## PREVIEW: CREATE

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Preview:Create dialog box lets you preview any SQL related command you want to perform on an object.

### Important Notes

None

The table below describes the options and functionality of the Preview:Create dialog box:

Button	Description
Save	Saves SQL.
Print	Prints SQL.
Email	Sends SQL via e-mail.
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.

## PREVIEW

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Preview dialog box lets you preview any SQL related command you want to perform on an object.

## Important Notes

None

The table below describes the options and functionality on the Preview dialog box:

Button	Description
Save	Saves SQL.
Print	Prints SQL.
Email	Sends SQL via e-mail.
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.

## JOB CARDS FOR BATCH EXECUTION

The Job Cards for Batch Execution dialog box lets you specify the job cards that should be used to execute the utility in batch.

## QUIESCE TABLESPACES

Option	Description
For user	Lets you select and type user name to restrict.
For group	Lets you select and type group name to restrict.
Connections	Lets you select a connection option to restrict.

### Quiesce Tablespaces

This functionality is available for IBM DB2 LUW for Linux, Unix, and Windows and IBM DB2 LUW for OS/390 and z/OS only.

The Quiesce Tablespaces dialog box lets you temporarily prevent other users from updating or reading selected tables on a tablespace. For example, you may not want the table to be updated during a backup procedure to capture a complete backup of the table as it exists.

## Important Notes

None

The table below describes the options and functionality on the Quiesce Tablespaces dialog box.

Option	Description
Quiesce Mode	<p>Exclusive - Prevents other users from viewing or modifying the table data.</p> <p>Intent to Update - Lets other users view but not update the table data.</p> <p>Reset - Lets you reset a table's quiesce mode.</p> <p>Share - Lets all users view (including you) but not modify the table data.</p>

For more information, see [Completing the Quiesce Tablespaces Dialog Box](#).

## COMPLETING THE QUIESCE TABLESPACES DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Tables node.  
Rapid SQL displays the Tables in the Database Explorer.
- 2 On the Database Explorer, right-click the table, and then select Quiesce Tablespaces.  
Rapid SQL opens the Quiesce Tablespaces dialog box.
- 3 Select a Quiesce Mode:
  - Exclusive - Prevents other users from viewing or modifying the table data.
  - Intent to Update - Lets other users view but not update the table data.
  - Reset - Lets you reset a table's quiesce mode.
  - Share - Lets all users view (including you) but not modify the table data.
- 4 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

## REASSIGN BY CATEGORY

This functionality is available for Oracle only.

The Reassign by Category dialog box lets you reassign the category of stored outlines in Oracle.

Outlines are a set of results for the execution plan generation of a particular SQL statement. When you create an outline, plan stability examines the optimization results using the same data used to generate the execution plan. That is, Oracle uses the input to the execution plan to generate an outline, and not the execution plan itself.

### Important Notes

None

The table below describes the options and functionality on the Reassign by Category dialog box:

Option	Description
Source Category	Lets you move all outlines from one category to a new category. Select the original category.
Destination Category	Reassigns the original category to the new category. Select a new destination category.

For more information, see [Completing the Reassign by Category Dialog Box](#).

## COMPLETING THE REASSIGN BY CATEGORY DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Outlines node.  
Rapid SQL displays the Outlines in the Database Explorer.
- 2 On the Database Explorer, right-click the outline, and then select Reassign by Category.  
Rapid SQL opens the Reassign by Category dialog box.
- 3 In the Source Category box, enter the target source outline category.
- 4 In the Destination Category box, enter the new category for the target outlines.
- 5 You can click the list to display existing categories, or you can type a new category in the box.
- 6 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .

Button	Description
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

For more information, see:

[Reassign by Category](#)

## REBIND PACKAGES

This functionality is available for IBM DB2 LUW for Linux, Unix, and Windows only.

The Rebind Package dialog box lets you update the best access path for SQL statements when the contents of a package changes.

### Important Notes

None

**TIP:** If the physical storage of a package is changed or dropped, rebinding updates the path of the SQL statements.

For more information, see [Completing the Rebind Packages Dialog Box](#).

## COMPLETING THE REBIND PACKAGES DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Packages node.  
Rapid SQL displays the Packages in the Database Explorer.
- 2 On the Database Explorer, right-click the package, and then select Rebind.  
Rapid SQL opens the Rebind dialog box.
- 3 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Execute	Executes the task.

Button	Description
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

## REBIND PLANS

This functionality is available for IBM DB2 LUW for Linux, Unix, and Windows only.

The Rebind Plans dialog box lets you update the best access path for SQL statements when the contents of a plan change.

**TIP:** If the physical storage of a plan is changed or dropped, rebinding updates the path of the SQL statements.

**TIP:** To filter the Owner and Object lists, click the Filter button next to each list. Rapid SQL opens the [Filter](#) dialog box.

### Important Notes

None

For more information, see [Completing the Rebind Plans Dialog Box](#).

## COMPLETING THE REBIND PLANS DIALOG BOX

To complete the Rebind Plans dialog box, do one of the following:

- 1 On the *Database Explorer*, select the Plans node.  
Rapid SQL displays the plans in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Rebind.  
Rapid SQL opens the Rebind Plans dialog box.
- 3 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Execute	Executes the task.

## REBUILD INDEX

See the following topics for DBMS-specific instructions on rebuilding indexes:

- [Rebuild Index \(Oracle\)](#)
- [Rebuild Index \(SQL Server\)](#)
- [Rebuild Index \(InterBase/Firebird\)](#)

## REBUILD INDEX (ORACLE)

**NOTE:** This functionality is available for Oracle only.

The Rebuild Indexes dialog box lets you rebuild an index that has become fragmented. Rebuilding an index is a good alternative to coalescing an index because you can move the index to a different tablespace and change both tablespace and storage parameters while eliminating fragmentation. However, rebuilding an index has a higher cost than coalescing an index. These same qualities also make rebuilding an index a viable alternative to dropping an index then re-creating it.

As a rule of thumb, check indexes for rebuilds when their level (or tree depth) reaches four or greater, or many deleted leaf rows are found. The Rebuild Indexes dialog box can also be used to easily move an index from one tablespace to another.

### Important Notes

- If you are rebuilding a function-based index, the index is enabled when the rebuild is finished.
- You cannot rebuild a partitioned index. You must rebuild each partition or subpartition individually.

The table below describes the options and functionality on the Rebuild dialog box.

Option	Description
New Tablespace	Defaults to the tablespace which currently includes the index. To change the tablespace containing the index, choose a new tablespace from the list.
Logging	Recoverable - The creation of the index logs in the redo log file. Non-Recoverable - The creation of the index is not logged in the redo log file.
Use Parallel Processes	Performs processes for the sequential execution of a SQL statement in parallel using multiple parallel processes. One process, known as the parallel execution coordinator, dispatches the execution of a statement to several parallel execution servers and coordinates the results from all of the server processes to send the results back to the user.  NOTE: Only available for Oracle with the Parallel Server option.  NOPARALLEL execution - Select this if you are concerned that the cost of synchronizing parallel processes will impede the throughput of data.
Order	Reverse - Instructs Oracle to store the bytes of the index block in reverse order and to exclude the ROWID when rebuilding the index.  No Reverse - Instructs Oracle to store the bytes of the index block in normal order when rebuilding the index.



## Completing the Rebuild Indexes Dialog Box

To complete this dialog box, do the following:

- 1 On the *Datasource Explorer*, select the Indexes node.  
Rapid SQL displays the Indexes in the right pane of the Explorer window.
- 2 On the *Datasource Explorer* tool bar, click Command, and then select Rebuild.  
OR  
In the right pane of the application, right-click the index, and then select Rebuild.  
Rapid SQL opens the Rebuild Indexes dialog box.
- 3 To move the index to a new tablespace, click the *New Tablespace* list and then click the new tablespace.
- 4 In the *Logging* box, click:
  - The *Recoverable* option button to make the operation log in the redo file.
  - The *Non-Recoverable* option button if you do not want the operation logged in the redo file.
- 5 If you are using Parallel Server, select the *Parallel Server* check box and:
  - Type a value indicating the number of query server processes that should be used in the operation in the *Degree* box.
  - Type a value indicating how you want the parallel query partitioned between the Parallel Servers in the *Instances* box.
- 6 In the *Order* box:
  - Click the *Reverse* option button to rebuild the index to store the bytes of the index block in reverse order.
  - Click the *No Reverse* option button to rebuild the index to store the bytes of the index block in order.

**NOTE:** This option is only available for Oracle8.
- 7 Click **Execute**. For information on the other options, see [Preview](#), [Scheduling](#), and [Impact Analysis](#).

## REBUILD INDEX (SQL SERVER)

The **Rebuild Indexes** dialog box lets you rebuild an entire index, primary key, or unique key or a single partition of those objects. Depending on your choice, a number of REBUILD WITH clause options are available.

**NOTE:** This functionality is available for Microsoft SQL Server 2005 and above.

**To rebuild an index**

- 1 On the *Datasource Explorer*, expand nodes until Microsoft SQL Server index, primary key, or unique key nodes are displayed.
- 2 Right-click an index, primary key, or unique key and select **Rebuild** from the context menu. The **Rebuild Index** wizard opens.
- 3 Use the following table as a guide to understanding and modifying settings in the wizard:

Step	Settings and tasks
<b>Action options</b>	To rebuild a single partition, specify a <b>Partition number</b> and optionally, provide <b>Sort in tempdb</b> and <b>MaxDOP</b> property values used to create the REBUILD WITH clause.  To rebuild the entire index, primary key, or unique key, DO NOT provide a <b>Partition Number</b> , and optionally, provide <b>Pad Index</b> , <b>Sort in tempdb</b> , <b>Ignore Duplicate Key</b> , <b>Statistics no recompute</b> , <b>Online</b> , <b>Allow Row Locks</b> , <b>Allow Page Locks</b> , <b>MaxDOP</b> , and <b>Fill Factor</b> property values used to create the REBUILD WITH clause.  For more information on these properties, see <a href="#">Index Wizard (SQL Server)</a> .
<b>Dependencies</b>	Lets you review the objects potentially impacted by this action.
<b>Preview</b>	Preview the DDL generated for the operation and when ready, use the <b>Schedule</b> or <b>Execute</b> button to perform this action.

- 4 Click **Execute**. For information on the other options, see [Preview](#) and [Scheduling](#).

**REBUILD INDEX (INTERBASE/FIREBIRD)**

This action rebuilds indexes by issuing ALTER INDEX/INACTIVE followed by ALTER INDEX/ACTIVE statements.

**To rebuild an index**

- 1 On the *Datasource Explorer*, expand nodes until the **Indexes** node is displayed.
- 2 Expand the indexes node and select one or more indexes.
- 3 Right-click the selected indexes and select **Rebuild Indexes** from the context menu. The **Rebuild Index** dialog opens.
- 4 Use the following table as a guide to understanding and modifying settings in the wizard:

Step	Settings and tasks
<b>Action options</b>	Displays the indexes you chose to rebuild.
<b>Dependencies</b>	Lists referring and referenced objects potentially impacted by the change. For details, see <a href="#">Dependencies</a> .
<b>Preview</b>	Preview the DDL generated for the operation and when ready, use the <b>Schedule</b> or <b>Execute</b> button to perform this action.

- 5 Click **Execute**. For information on the other options, see [Preview](#) and [Scheduling](#).

## REBUILD OUTLINES

This functionality is available for Oracle only.

The Rebuild Outlines dialog box lets you rebuild an outline.

Outlines are a set of results for the execution plan generation of a particular SQL statement. When you create an outline, plan stability examines the optimization results using the same data used to generate the execution plan. That is, Oracle uses the input to the execution plan to generate an outline, and not the execution plan itself.

### Important Notes

None

The table below describes the options and functionality on the Rebuild Outlines dialog box:

Option	Description
Outline(s) to be rebuilt box	Displays the outline(s) available to rebuild.

## RECOMPILE

This functionality is available for Microsoft SQL Server and Sybase ASE only.

The Recompile dialog box lets you specify single or multiple objects for recompilation. Recompilation causes each procedure and trigger that uses the target table to be recompiled the next time it runs.

The queries used by procedures and triggers are optimized only once, when they are compiled. As you add indexes or make other changes to your database that affect its statistics, your compiled procedures and triggers may lose efficiency. By recompiling the procedures and triggers that act on a table, you can optimize the queries for maximum efficiency.

## REDISTRIBUTE

This functionality is available for IBM DB2 LUW for Linux, Unix, and Windows only.

The Redistribute dialog box lets you redistribute the data in a nodegroup.

### Important Notes

None

The table below describes the options and functionality on the Redistribute dialog box:

Option	Description
Redistribute Options	Lets you click the appropriate option button: Uniform Using Dist. File Using Target Map Continue Rollback
Using Dist. File	Lets you select and type the distinct file name in the corresponding box.
Using target File	Lets you select and type the target file name in the corresponding box

The table below describes the options and functionality on the Refit dialog box:

Option	Description
Redistribute Options	Lets you click the appropriate option button: Uniform Using Dist. File Using Target Map Continue Rollback
Using Dist. File	Lets you select and type the distinct file name in the corresponding box.
Using target File	Lets you select and type the target file name in the corresponding box

### Important Notes

None

## REFRESH SUMMARY TABLE

This functionality is available for IBM DB2 LUW for Linux, Unix, and Windows 8.1 only.

The Refresh Summary Table Dialog box lets you reload summary tables that have been defined with refresh options.

### Important Notes

None

The table below describes the options and functionality on the Refresh Summary Table dialog box:

Option	Description
Refresh	Lets you select a refresh option.  NOTE: For IBM DB2 LUW for Linux, Unix, and Windows 7, summary tables refresh when the system determines the refresh patterns.

For more information, see [Completing the Refresh Summary Table Dialog Box](#).

## COMPLETING THE REFRESH SUMMARY TABLE DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click a summary table, and then select Refresh Summary Table.  
Rapid SQL opens the Refresh Summary Table dialog box.
- 3 Select a refresh option.
- 4 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.

## Important Notes

## RENAME

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Interbase/Firebird, Oracle, SQL Server, and Sybase.

The Rename dialog box lets you rename an object.

### Microsoft SQL Server

Microsoft SQL Server lets you rename a database if you own it. Before renaming a database, set it to single-user mode.

Microsoft SQL Server does not rename the table if it is referenced within the body of other objects that call it, such as tables, triggers or views. As a result, renaming a table can result in broken dependencies with other objects. Also, Microsoft SQL Server does not let you rename System Tables.

### IBM DB2 LUW for OS/390 and z/OS

Rapid SQL lets you rename a primary key if the underlying table has only one owner.

The rename operation does not rename the table if it is referenced within the body of other objects, such as tables, triggers or views, that call it. As a result, renaming a table can result in broken dependencies with other objects.

### Sybase ASE

Before renaming a database, set it to single-user mode.

System indexes can not be renamed.

The rename operation does not rename the stored procedure if it is referenced within the body of other objects, such as another stored procedure, that call it. As a result, renaming a stored procedure can result in broken dependencies with other objects.

The rename operation does not rename the table if it is referenced within the body of other objects, such as tables, triggers or views, that call it. As a result, renaming a table can result in broken dependencies with other objects.

The rename operation does not rename the view if it is referenced within the body of other objects, such as stored procedures, triggers or other views, that call it. As a result, renaming a view can result in broken dependencies with other objects.

The table below describes the options and functionality on the Rename dialog box.

**NOTE:** Options differ by platform.

Option	Description
New Name	Lets you type the new name of the object.

For more information, see [Completing the Rename Dialog Box](#).

## COMPLETING THE RENAME DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Rename.  
Rapid SQL opens the Rename dialog box.
- 3 In New Name, type the new name.
- 4 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

## REORGANIZE

The Reorganize dialog box lets you alter a table's physical storage to ensure that it is stored efficiently and to eliminate fragmentation. For details, see the following topics:

- [Reorganizing IBM DB2 LUW for Linux, Unix, and Windows Objects](#)
- [Reorganizing Oracle Objects](#)
- [Reorganizing Sybase ASE Objects](#)
- [Reorganizing SQL Server indexes, primary keys, and unique keys](#)

### Important Notes

None

For more information, see [Completing the Reorganize Dialog Box](#).

## COMPLETING THE REORGANIZE DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Reorganize.  
Rapid SQL opens the Reorganize dialog box.
- 3 If you want to move the table(s) to a new tablespace, click the New Tablespace list and select the target tablespace.
- 4 In the *Data Block Storage* box, indicate the data block storage parameters:
  - In the *Percent Free* box, type the appropriate percent free value for the table.
  - In the *Percent Used* box, type the appropriate percent used value for the table.
  - In the *Initial Transactions* box, type the appropriate initial transactions value for the table.
  - In the *Max Transactions* box, type the appropriate maximum transactions value for the table.
- 5 In the *Extents* box, indicate the extents parameters:
  - In the *Initial Extent* box, type the appropriate initial extent KB value for the table.
  - In the *Next Extent* box, type the appropriate next extent KB value for the table.
  - In the *Percent Increase* box, type the appropriate percent increase value for the table.
  - In the *Minimum Extents* box, type the appropriate minimum extents value for the table.
  - In the *Maximum Extents* box, type the appropriate maximum extents value for the table.
- 6 To modify freelist parameters, in the *Freelists* box:
  - In the *Freelists* box, type the new value.
  - In the *Freelist Group* box, type the new value.
- 7 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.

For more information, see:

[Reorganize](#)



## REORGANIZING IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS OBJECTS

### Reorganize Dialog Box (One Table)

The table below describes the options and functionality on the Reorganize dialog box.

Option	Description
Index	Associates a temporary index with the table's index. You can select another index from the list.
Temporary Tablespace	Associates a temporary tablespace with the table's tablespace. You can select another tablespace from the list.

### Reorganize Dialog Box (Multiple Tables)

The table below describes the options and functionality on the Reorganize dialog box.

Option	Description
Temporary Tablespace	Associates a temporary tablespace with the table's tablespace. You can select another tablespace from the list.

For more information, see [Completing the Reorganize Dialog Box](#).

## REORGANIZING ORACLE OBJECTS

The Reorganize dialog box lets you reduce query processing time against tables. This functionality is available for both tables.

For more information, see [Completing the Reorganize Dialog Box](#).

## REORGANIZING SYBASE ASE OBJECTS

The Reorganize dialog box lets you reduce query processing time against tables. This functionality is available for both [tables](#) and [indexes](#).

For more information, see [Completing the Reorganize Dialog Box](#).

## REORGANIZING SYBASE ASE TABLES

The Reorganize Table dialog box lets you reduce the query processing time against a table by reorganizing the table to ensure that space is properly allocated to it. For lengthy reorganization processes, this dialog box also lets you [execute a process in increments](#), lets you resume an incomplete process, and lets you specify the duration of each increment.

**TIP:** Frequent update activity on a table can cause data rows to migrate and to chain over multiple data pages. Chained or forwarded rows can degrade performance because more physical reads are required to access a row of data. Consequently, you should monitor chained rows regularly to spot performance bottlenecks before they become severe. In addition, altering physical storage parameters can lead to fragmentation of space on your data pages, which also results in reduced performance levels.

You should consider reorganizing a table if you are experiencing slow performance due to:

- A large number of chained or forwarded rows on your data pages
- A large amount of fragmentation in your data pages

The table below describes the options and functionality on the Reorganize Table dialog box:

Option	Description
Compact	Lets you reclaim space and undo row forwarding.  Minimizes interference with other activities by using multiple small transactions of brief duration. Each transaction is limited to eight pages of reorg processing. These three commands also provide <a href="#">resume</a> and time options that allow you to set a time limit on how long a reorg runs and to resume a reorg from the point at which the previous reorg stopped, making it possible to use a series of partial reorganizations at off-peak times to reorg a large table.
Reclaim Space	Lets you reclaim unused space resulting from deletions and row-shortening updates on a page.  Minimizes interference with other activities by using multiple small transactions of brief duration. Each transaction is limited to eight pages of reorg processing. These three commands also provide <a href="#">resume</a> and time options that allow you to set a time limit on how long a reorg runs and to resume a reorg from the point at which the previous reorg stopped, making it possible to use a series of partial reorganizations at off-peak times to reorg a large table.

Option	Description
Rebuild	<p>Lets you undo row forwarding and reclaim unused page space. It also rewrites all rows to comply with the target table's clustered index, writes rows to data pages to comply with space management setting changes (via <code>sp_chgattribute</code>), and drops and re-creates all the target table's (or tables') indexes.</p> <p>Reorg rebuild holds an exclusive table lock for its entire duration. On a large table this can be a significant amount of time. However, reorg rebuild accomplishes everything that dropping and re-creating a clustered index does and takes less time. In addition, reorg rebuild rebuilds the table using all of the table's current space management settings. Dropping and re-creating an index does not use the space management setting for <code>reservepagegap</code>. In most cases, reorg rebuild requires additional disk space equal to the size of the table it is rebuilding and its indexes.</p>
Undo Row Forwarding	Lets you undo row forwarding, a process that occurs when an update increases a row's length in a data-only-locked table such that the row is too large to fit on its original page.
Options	<p>Start at the point where a previous reorg left off - Select to resume a previously initiated but incomplete partial reorganization. Then specify the duration for which you want the resumed reorganization to continue before stopping again. This box is disabled for the rebuild command.</p>

For more information, see [Completing the Reorganize Dialog Box](#).

#### INCREMENTAL REORGANIZATIONS

If target tables are too long to reorganize in one session, Rapid SQL lets you reorganize them in increments over multiple sessions by specifying a maximum duration for each session. After Rapid SQL reorganizes tables for the specified duration, the operation stops until you resume it again from the Options box of the ReOrganize Table dialog box. The Options box lets you specify to resume a previously initiated but incomplete partial reorganization. It also lets you specify the duration for which you want a resumed reorganization to continue before stopping again. The Option box is disabled for the rebuild command.

**NOTE:** The duration you specify refers to elapsed time, not CPU time

In the option box, if you select the check box without specifying a duration, Rapid SQL executes the reorg at the point where the previous reorg stopped and continues to the end of the target tables. If you clear the check box and specify a duration, the reorg starts at the beginning of the target tables and continues for the specified number of minutes. If you select the check box and specify a duration, Rapid SQL runs the reorg from the point where it last left off, and continues for the specified number of minutes.

**NOTE:** If you reorganize a table using one command (Compact, Reclaim Space, or Undo Forwarding) for a specified duration, you cannot resume the process from its resume point using a different command. For example, you cannot compact a table for an hour, and then reclaim space on the remainder of the table. A resumed reorganization process must utilize the same command from start to finish. Selecting a different command begins a new reorganization process.

**CAUTION:** While this option lets you reorganize a large table in multiple manageable pieces, any updates to the table between reorganization runs can cause pages to be skipped or processed more than once.

## REORGANIZE SYBASE ASE INDEXES

The Reorganize Index dialog box lets you reduce the query processing time against a table by running a reorg rebuild command on the target index.

This operation:

- Undoes row forwarding and reclaim unused page space
- Rewrites all rows in the table to comply with the table's clustered index
- Writes rows to data pages to comply with space management setting changes (via `sp_chgattribute`)
- Drops and re-creates the table's indexes

Reorg rebuild holds an exclusive table lock for its entire duration. On a large table this can be a significant amount of time. However, reorg rebuild accomplishes everything that dropping and re-creating a clustered index does and takes less time. In addition, reorg rebuild rebuilds the table using all of the table's current space management settings. Dropping and re-creating an index does not use the space management setting for `reservepagegap`. In most cases, reorg rebuild requires additional disk space equal to the size of the table it is rebuilding and its indexes.

For more information, see [Completing the Reorganize Dialog Box](#).

## REORGANIZING SQL SERVER INDEXES, PRIMARY KEYS, AND UNIQUE KEYS

The **Reorganize Indexes** dialog box lets you reorganize an entire index, primary key, or unique key or a single partition of that object. It also lets you specify a `LOB_COMPACT` option.

**NOTE:** This functionality is available for Microsoft SQL Server 2005 and above.

**NOTE:** You cannot reorganize an index (primary key, or unique key) that has an **Allow Page Locks** property set to `FALSE`. For information on setting index properties, see [Indexes Editor \(SQL Server\)](#).

### To reorganize an index

- 1 On the *Datasource Explorer*, expand nodes until Microsoft SQL Server index nodes, primary keys, or unique keys are displayed.
- 2 Right-click an index, primary key, or unique key and select **Reorganize** from the context menu. The **Reorganize Index** wizard opens.

3 Use the following table as a guide to understanding and modifying settings in the wizard:

Step	Settings and tasks
<b>Action options</b>	To reorganize a single partition, specify a <b>Partition number</b> and optionally, specify <b>Lob Compaction</b> used to create the REORGANIZE WITH clause.  To reorganize the entire index, primary key, or unique key, DO NOT provide a <b>Partition Number</b> , and optionally, specify <b>Lob Compaction</b> used to create the REORGANIZE WITH clause.
<b>Dependencies</b>	Review the referring and referred objects that will be automatically resolved when you execute this operation.
<b>Preview</b>	Preview the DDL generated for the operation and when ready, use the <b>Schedule</b> or <b>Execute</b> button to perform this action.

4 Click **Execute**. For information on the other options, see [Preview](#) and [Scheduling](#).

## REPORT

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Interbase/Firebird, MySQL, Oracle, SQL Server, and Sybase.

The Generate Report dialog box lets you generate detailed reports about all the objects in your database. The reports give you the complete information about the properties of your objects. They derive their contents from the information displayed in the object editors. To facilitate the online publication of these reports, Rapid SQL generates these reports in HTML. You can view these reports directly in Rapid SQL's built-in HTML browser.

### Important Notes

None

For more information, see [Completing the Generate Report Dialog Box](#).

## COMPLETING THE GENERATE REPORT DIALOG BOX

To complete the Generate Report dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Report.  
Rapid SQL opens the Generate Report dialog box.

### Generating a Summary Report

- 1 Select the Summary Report (Listing all items) option.

Rapid SQL displays the report. You can find data in the report, save, send, and print the report.

- 2 To exit the report, click Close.

### Generating a Detail Report

- 1 Select the Detail Report (Report on each supported item) option.

Rapid SQL opens the Report dialog box.

- 2 In Report Home Page File Name, type the report name or click Browse to locate the report.

- 3 In Report Title, type the report title.

Rapid SQL displays Object Name and Object Type.

- 4 Click Execute.

Rapid SQL displays the report. You can find data in the report, save, send, and print the report.

- 5 To exit the report, click Close.

For more information, see:

[Report](#)

## RESTART

This functionality is available for IBM DB2 LUW for Linux, Unix, and Windows only.

The Restart Sequence dialog box lets you restart a sequence.

A sequence is a programmable database objects that generates a definable sequence of values. A sequence can be made available to many users.

### Important Notes

None

## SCHEMA

This functionality is available for all platforms.

The Schema dialog box lets you view everything that makes up a selected table.

**Important Notes**

None

For more information, see [Viewing Schema](#).

**VIEWING SCHEMA**

To view schema, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Schema.  
Rapid SQL opens the Schema window.

**SELECT \* FROM**

**NOTE:** This functionality is available for DB2 LUW, DB2 z/OS, Interbase/Firebird, Oracle, SQL Server, and Sybase.

The Select \* From dialog box lets you retrieve all data from a selected table or view.

**Important Notes**

None

For more information, see [Completing the Select \\* From Dialog Box](#).

**COMPLETING THE SELECT \* FROM DIALOG BOX**

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Select \* From.  
Rapid SQL opens the Select \* From dialog box.

## SET DEFAULT

This function lets you set a tablespace as the default tablespace. Users created without a specified default tablespace will be assigned this default tablespace. If no default tablespace is set, users created without a specified tablespace will have their default tablespace set to SYSTEM.

**NOTE:** This functionality is available for Oracle 10g tablespaces.

### To set a tablespace as the default

- 1 On the *Datasource Explorer*, expand nodes until Oracle tablespace nodes are displayed.
- 2 Right-click an Oracle 10g tablespace and select **Set Default** from the context menu. The **Set Default** wizard opens.
- 3 Use the following table as a guide to understanding and modifying settings in the wizard:

Step	Settings and tasks
<b>Action options</b>	Verify that the panel displays the tablespace to be set as default.
<b>Dependencies</b>	Review the referring and referred objects that will be automatically resolved when you execute this operation.
<b>Preview</b>	Preview the DDL generated for the operation and when ready, use the <b>Schedule</b> or <b>Execute</b> button to perform this action.

## SET ONLINE/OFFLINE

This functionality is available for Microsoft SQL Server, Oracle, and Sybase ASE only.

The Set Database(s) Online/Offline dialog box lets you disable your databases to prevent access, and enable your databases to grant access through the Datasource menu.

### Important Notes

For Sybase, Rapid SQL only lets you set databases online.

For more information, see [Completing the Set Database\(s\) Online/Offline Dialog Box](#).

## COMPLETING THE SET DATABASE(S) ONLINE/OFFLINE DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Databases node.

Rapid SQL displays the Databases in the Database Explorer.



- 2 On the Database Explorer, right-click the database, and then select Set Online/Offline.  
Rapid SQL opens the Set Database(s) Online/Offline dialog box.
- 3 Select an option:
  - Set Online
  - Set Offline
- 4 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

## SET STATISTICS

This action creates and submits a SET STATISTICS statement, letting you recompute the selectivity of an index.

**NOTE:** The action is available for Interbase/Firebird.

To recompute the selectivity of an index:

- 1 On the *Datasource Explorer*, expand nodes until the **Indexes** node is visible, and then expand the **Indexes** node.
- 2 Select one or more indexes.
- 3 Right click the selected indexes and select Set Statistics from the context menu.  
DBArtisan opens the **Set Statistics** dialog box.
- 4 Use the following table as a guide to understanding and modifying settings in this wizard:

Step	Settings and tasks
<b>Action Options</b>	Displays the indexes you selected.
<b>Dependencies</b>	Lists referring and referenced objects potentially impacted by the change. For details, see <a href="#">Dependencies</a> .
<b>Preview</b>	Displays the DDL that will execute the object action. For details, see <a href="#">Preview</a> .

- 5 Click **Execute**. For information on the other options, see [Preview](#) and [Scheduling](#).

## SET TABLESPACES QUOTA

This functionality is available for Oracle only.

The Set Tablespaces Quota dialog box lets you set a quota for a tablespace. When you assign a quota:

- Users with privileges to create certain types of objects can create those objects in the specified tablespace.
- Oracle limits the amount of space that can be allocated for storage of a user's objects within the specified tablespace to the amount of the quota.

### Important Notes

None

The table below describes the options and functionality on the Set Tablespaces Quota dialog box.

Option	Description
Quota	Lets you set a quota on the selected tablespace. You can select an unlimited, or specified size.  Unlimited - Lets you place an unlimited quota on the tablespace.  Other - Lets you place a specified quota in KB or MB on the tablespace.

For more information, see [Completing the Set Tablespaces Quota Dialog Box](#).

## COMPLETING THE SET TABLESPACES QUOTA DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Tablespaces node.  
Rapid SQL displays the Tablespaces in the Database Explorer.
- 2 On the Database Explorer, right-click the tablespace, and then select Quota.  
Rapid SQL opens the Set Tablespaces Quota dialog box.
- 3 Click the Unlimited or Other option button.
- 4 If you click Other, type the value of the quota and then click list to indicate KB or MB.
- 5 Click OK.
- 6 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

For more information, see:

[Set Tablespaces Quota](#)

## SET UNDO TABLESPACE

**NOTE:** This functionality is available for Oracle 9 or later.

Set UNDO Tablespace dialog box lets you dynamically set an UNDO tablespace if the tablespace is running in AUTO UNDO mode.

### To dynamically set an UNDO tablespace

- 1 On the *Datasource Explorer*, select the Tablespaces node.

Rapid SQL displays the Tablespaces in the right pane of the Explorer window.

- 2 On the *Datasource Explorer* tool bar, click Command, and then select Set UNDO Tablespace.

Rapid SQL opens the Set UNDO Tablespace dialog box.

Use the following table as a guide to understanding and modifying the settings on this wizard:

Step	Functionality
<b>Action Options</b>	Displays the tablespace you selected to be set as an Undo tablespace.
<b>Dependencies</b>	View the dependencies on the tablespace.
<b>Preview</b>	Preview the DDL generated from your choices.

- 3 Click **Execute**.

## COMPLETING THE SET UNDO TABLESPACE DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the Tablespaces node.  
Rapid SQL displays the Tablespaces in the Database Explorer.
- 2 On the Database Explorer, right-click the tablespace, and then select Set Undo Tablespace.  
Rapid SQL opens the Set Undo Tablespace dialog box.
- 3 In New Undo Tablespace select from the possible UNDO tablespace choices. Select NULL to un-assign all UNDO tablespaces
- 4 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.

For more information, see:

[Set Tablespaces Quota](#)

## SHRINK

**NOTE:** The Shrink functionality is available for Microsoft SQL Server and Oracle 7.3 or later.

### Microsoft SQL Server

The Shrink Database dialog box lets you reclaim space from a database that is too large.

### Important Notes

None

The table below describes the options and functionality on the Shrink Database dialog box.

Option	Description
Move data pages to beginning of file before shrink	Select to move data pages to the beginning of the file before the shrink.

Option	Description
Release All Unused Space	Deselect to set the target free space to retain, and then in the Target Free Space to Retain (percent) box, type the new value of free space to retain.  The new size for the database must be at least as large as the Minimum Allowable Size displayed in the Current File Size box.
Target free space percent after shrink	Lets you specify the target free space percent after the shrink.

## Oracle

The Shrink Rollback Segments dialog box lets you shrink the size of rollback segments. The proper sizing of rollback segments is critical to their overall performance. Performance degrades whenever a rollback segment must extend, wrap or shrink in response to transaction loads. Ideally, you want to make the extents of rollback segments as small as possible while still ensuring that each transaction can fit into a single extent.

After an abnormally large transaction load, you might consider shrinking a rollback segment to eliminate unnecessary space. Oracle 7.3 or later lets you shrink a rollback segment manually by a specific amount or back to its Optimal Size.

## Important Notes

For Oracle 9i or later, Shrink is not available if auto-UNDO management is enabled.

The table below describes the options and functionality on the Shrink Rollback Segments dialog box:

Option	Functionality
Specify the size...	If you do not provide a specific number the Rollback Segment uses the OPTIMAL value specified in the Storage clause. If an OPTIMAL value is not specified, the size defaults to the MINEXTENTS value of the Storage clause.

## STATUS

**NOTE:** The Status functionality is available for [Microsoft SQL Server](#), [Oracle](#), and Sybase ASE.

The Set Constraint(s) Status dialog box lets you change the status of check constraints, foreign key constraints, primary key constraints, and unique key constraints.

## Important Notes

None

For more information, see [Completing the Set Constraint\(s\) Status Dialog Box](#).

## COMPLETING THE SET CONSTRAINT(S) STATUS DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Status.  
Rapid SQL opens the Set Constraint(s) Status dialog box.
- 3 Select dialog box options.
- 4 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.
Analysis	Opens the <a href="#">Impact Analysis dialog box</a> . This option is not available for every dialog box.

For more information, see:

[Setting Constraint Status for Microsoft SQL Server Objects](#)

[Setting Constraint Status for Oracle](#)

## SETTING CONSTRAINT STATUS FOR MICROSOFT SQL SERVER OBJECTS

The Set Constraint Status dialog box lets you specify the ability of a group of constraints to be replicated, and (for Microsoft SQL Server version 7 or later) enable or disable check constraints, foreign key constraints, primary key constraints, and unique key constraints.

The table below describes the options and functionality on the Set Constraint Status dialog box.

Option	Description
Enabled	Deselect to temporarily override listed check constraints. Useful when you need to execute special processes that would ordinarily incur constraint-related errors.
Not for Replication	When you duplicate the table schema and data of a source database containing constraints marked "Not for Replication", these objects are not carried over to the duplicate of the schema.

For more information, see:

[Status](#)

## SETTING CONSTRAINT STATUS FOR ORACLE

The Set Constraint(s) Status dialog box lets you change the status of check constraints, foreign key constraints, primary key constraints, and unique key constraints. Rapid SQL lets you enable or disable selected constraints and, in the case of primary key and unique key constraints, lets you enable with or without validation and disable with or without the changes cascading.

When enabled, the rule defined by the constraint is enforced on the data values in the columns on which the constraint is placed. When disabled, the constraint rule is not enforced but the constraint continues to be stored in the data dictionary.

Temporarily disabling constraints can improve performance when you are loading large amounts of data or when you are making massive changes to a table. Disabling constraints also can be useful if you are importing or exporting one table at a time.

**NOTE:** Primary keys for index-organized tables cannot be disabled.

**NOTE:** You cannot drop a unique or primary key constraint that is part of a referential integrity constraint without also dropping the foreign key. To drop the referenced key and the foreign key together, select the Cascade check box in the Set Constraint(s) Status dialog box.

The table below describes the options and functionality on the Set Constraint(s) Status dialog box.

**NOTE:** The options differ by object.

Option	Description
Enable	Enabling the constraint and not selecting the Validate check box automatically uses Oracle ENABLE NOVALIDATE clause which enables a constraint so that it does not validate the existing data. A table using constraints in enable novalidate mode can contain invalid data but you cannot add new invalid data to that table. The enable novalidate mode is useful as an intermediate state or when you do not want the constraint to check for possible exceptions (e.g., after a data warehouse load).
Validate	Enabling the constraint and selecting the Validate check box causes Oracle to validate all existing data in the key columns of the table with the constraint. If an exception exists, Oracle returns an error and the constraint remains disabled.

Option	Description
Cascade	Selecting the Cascade check box when disabling a primary key or foreign key constraint instructs Oracle to simultaneously disable any constraints that depend on the primary or unique key. Selecting the Delete Cascade check box instructs Oracle to delete data in the child table (on which the foreign key is defined) if the referenced data in the parent table is deleted.

For more information, see:

[Status](#)

## SUMMARY DEFINITION

This functionality is available for IBM DB2 LUW for Linux, Unix, and Windows only.

The Summary Definition dialog box lets you enter the query for your table. If you are creating a table as a result of a query, you need to include the query in the AS FullSelect box.

**TIP:** You can open the query in the ISQL Editor and use the copy/paste function to add the query to the AS FullSelect box.

### Important Notes

You must complete the Summary Definition dialog box to proceed with the Table Wizard.

The table below describes the options and functionality on the Summary Definition dialog box:

Option	Description
Table Name	Read-only field that displays the table name.
AS FullSelect	Enter the query in box.
Summary table Options	Lets you set table options for the table you are creating. Definition Only - Lets you set copy options for the table definition. If you select this check box, you can set the copy options to include Column Defaults and Identity Column Attr. The default selection is Column Defaults. Refreshable Options - Lets you set table refresh options for data and query.

## SWITCH ONLINE

**NOTE:** The Switch Online functionality is available for IBM DB2 LUW for Linux, Unix, and Windows only.

The Switch Online dialog box lets you access a tablespace by switching it online after the parent container(s) have been recovered or restored.



## Important Notes

None

For more information, see [Completing the Switch Online Dialog Box](#).

## COMPLETING THE SWITCH ONLINE DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Switch Online.  
Rapid SQL opens the Switch Online dialog box.
- 3 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.

## TRANSFER OWNERSHIP

Rapid SQL lets an object owner or a user with SECADM authority transfer ownership of an object to another user. The new user is automatically granted the same privileges as the former owner. Ownership can be transferred on an object-by-object basis or you can transfer all objects currently owned by an individual user.

**NOTE:** This functionality is available for IBM DB2 for Linux, Unix, and Windows version 9.1 and higher.

The following table lists the object types that support ownership transfer:

Aliases	Check constraints	Event monitors	Foreign Keys
Functions	Indexes	Primary keys	Procedures
Materialized query tables	Schema	Sequences	Structured types
Tables	Tablespaces	Triggers	Unique keys
User datatypes	Views		

### To transfer ownership of an object

- 1 Connect to the data source that contains the object or objects for which ownership is being transferred. For more information, see [Datasource Management](#).
- 2 On the *Datasource Explorer*, expand the data source nodes until the target object type node is visible and then select that node. Rapid SQL displays objects of that type in the right pane of the Explorer window.

**NOTE:** When you select a non-user object type, the transfer ownership operation will apply to the object or objects you select. If you select a user, the operation will apply to all objects owned by that user.

- 3 In the right pane of the data source explorer, select one or more objects, right-click and select Transfer Ownership from the context menu. Rapid SQL opens the Transfer Ownership wizard.
- 4 Use the following table as a guide to understanding and modifying settings in the wizard:

Step	Settings and tasks
<b>Action options</b>	<p>If you initiated the ownership transfer against a user, select the specific objects belonging to that user that are to be assigned to a new user.</p> <p>Use the <b>New Owner</b> dropdown to select a defined user as the new owner of the selected object or objects.</p>
<b>Preview</b>	<p>Preview the DDL generated for the operation and when ready, click <b>Execute</b>. For information on the other options, see <a href="#">Preview</a> and <a href="#">Scheduling</a>.</p>

## TRUNCATE

The Truncate dialog box lets you quickly delete the rows of a table. For details, see the following topics:

[Truncating IBM DB2 LUW Objects](#)

[Truncating Microsoft SQL Server Objects](#)

[Truncating Oracle Objects](#)

[Truncating IBM DB2 LUW for OS/390 and z/OS Objects](#)

[Truncating Sybase ASE Objects](#)

### Important Notes

If you truncate a table, Rapid SQL deletes all the rows. These rows are not logged as individual drops and cannot be recovered from a transaction log or other type of log.

For more information, see [Completing the Truncate Dialog Box](#).

## COMPLETING THE TRUNCATE DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Truncate.  
Rapid SQL opens the Truncate dialog box.
- 3 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.

For more information, see:

[Truncating IBM DB2 LUW for Linux, Unix, and Windows Objects](#)

[Truncating Microsoft SQL Server Objects](#)

[Truncating Oracle Objects](#)

[Truncating IBM DB2 LUW for OS/390 and z/OS Objects](#)

[Truncating Sybase ASE Objects](#)

## TRUNCATING IBM DB2 LUW OBJECTS

The Truncate dialog box lets you quickly delete the rows of a table.

**NOTE:** If you truncate a table, Rapid SQL deletes all the rows. These rows are not logged as individual drops and cannot be recovered from a transaction log or other type of log.

For more information, see:

[Truncate](#)

[Completing the Truncate Dialog Box](#)

## TRUNCATING MICROSOFT SQL SERVER OBJECTS

The Truncate dialog box lets you quickly delete the rows of a table.

**NOTE:** If you truncate a table, Rapid SQL deletes all the rows. These rows are not logged as individual drops and cannot be recovered from a transaction log or other type of log.

For more information, see:

[Truncate](#)

[Completing the Truncate Dialog Box](#)

## TRUNCATING ORACLE OBJECTS

The Truncate dialog box lets you truncate tables and clusters. Truncating a table or cluster is a quick and efficient way to delete all of a table's or a cluster's existing rows. You can truncate any table or cluster in their schema or, if you have the DROP ANY TABLE system privilege, you can truncate any table in any schema.

When you truncate a table or cluster, you can specify whether space currently allocated for the table is returned to the containing tablespace or if it is returned to the system. The table below includes notes on the objects you can truncate:

Object	Notes
Cluster	<p>Before truncating a cluster containing a parent key, disable all referencing foreign keys existing in different tables.</p> <p>Truncating a cluster does not generate any rollback information and commits immediately.</p> <p>Oracle alters the storage parameter NEXT to the size of the last extent deleted from the segment.</p>

Object	Notes
Table	<p>Before truncating a table containing a parent key, disable all referencing foreign keys existing in different tables.</p> <p>Truncating a table does not generate any rollback information and commits immediately.</p> <p>Oracle alters the storage parameter NEXT to the size of the last extent deleted from the segment.</p> <p>Oracle automatically deletes all data in the table's indexes and any materialized view direct-load INSERT information associated with a truncated table.</p> <p>If the table is not empty, all associated nonpartitioned indexes and all partitions of associated global partitioned indexes are marked unusable.</p> <p>You cannot truncate a hash cluster nor can you truncate individual tables in a hash cluster or an index cluster.</p>

The table below describes the options and functionality on the Truncate dialog box.

Option	Description
Storage Handling	<p>Drop Storage - Select if you want the freed extents returned to the system where they can be used by other objects.</p> <p>Reuse Storage - Select if you want the space to remain allocated to the table or cluster you have just truncated.</p>

For more information, see:

[Truncate](#)

[Completing the Truncate Dialog Box](#)

## TRUNCATING IBM DB2 LUW FOR OS/390 AND Z/OS OBJECTS

The Truncate dialog box lets you quickly delete the rows of a table.

**CAUTION:** If you truncate a table, Rapid SQL deletes all the rows. These rows are not logged as individual drops and cannot be recovered from a transaction log or other type of log.

Truncating a table is a faster alternative to deleting all of its rows. You cannot recover a truncated table because TRUNCATE TABLE is an unlogged command.

When truncating an IBM DB2 z/OS table, the **Action Options** tab offers the following settings:

Setting	Description
<b>Reuse Storage</b>	Corresponds to the DROP STORAGE/REUSE STORAGE clause of a TRUNCATE TABLE statement. This setting specifies whether storage currently allocated to the table is reused or dropped.
<b>Restrict When Delete Triggers</b>	Corresponds to the RESTRICT WHEN DELETE TRIGGERS/IGNORE DELETE TRIGGERS clause of a TRUNCATE TABLE statement. When enabled, an error is returned if triggers are defined for the table. When disabled, triggers defined for the table are not activated by the Truncate operation.
<b>Immediate</b>	Corresponds to the IMMEDIATE clause of a TRUNCATE TABLE statement. If enabled, the truncate operation is executed immediately and cannot be undone. If disabled, a Rollback can undo the Truncate operation.

For more information, see:

[Truncate](#)

[Completing the Truncate Dialog Box](#)

## TRUNCATING SYBASE ASE OBJECTS

The Truncate dialog box lets you truncate single or multiple tables. Truncating a table is a faster alternative to deleting all its rows. You cannot recover a truncated table because TRUNCATE TABLE is an unlogged command.

**NOTE:** You cannot truncate a table referenced by a foreign key constraint. Instead, use a DELETE statement without a WHERE clause.

**TIP:** When you truncate a table, Sybase ASE removes all rows from the target table, but retains the table structure (its indexes, columns, constraints, etc.). The counter used by an identity for new rows is reset to the seed for the column. To retain the identity counter, use a DELETE statement instead of TRUNCATE. To remove the target table definition and its data, use a DROP TABLE statement.

For more information, see:

[Truncate](#)

[Completing the Truncate Dialog Box](#)

## UPDATE STATISTICS

The Update Statistics dialog box lets you update the statistics for an active table or index. As indexes grow and shrink in response to data modification, the accuracy of their statistics can deteriorate.

**NOTE:** The Update Statistics dialog box is available for [IBM DB2 LUW for Linux, Unix, and Windows](#), [Microsoft SQL Server](#), and [Sybase ASE](#).

## Important Notes

None

For more information, see [Completing the Update Statistics Dialog Box](#).

## COMPLETING THE UPDATE STATISTICS DIALOG BOX

To complete this dialog box, do the following:

- 1 On the *Database Explorer*, select the target object node.  
Rapid SQL displays the target objects in the Database Explorer.
- 2 On the Database Explorer, right-click the target object, and then select Update Statistics.  
Rapid SQL opens the Update Statistics dialog box.
- 3 Select dialog box options.
- 4 Do one of the following:

Button	Description
Preview	Opens the <a href="#">Preview dialog box</a> .
Schedule	Opens Embarcadero Job Scheduler, or (if unavailable) another scheduling application. For more, see <a href="#">Scheduling</a> .
Execute	Executes the task.

For more information, see:

[Updating IBM DB2 LUW for Linux, Unix, and Windows Object Statistics](#)

Updating Microsoft SQL Server Object Statistics

Updating Sybase ASE Object Statistics

## UPDATING IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS OBJECT STATISTICS

The table below describes the options and functionality on the Update Statistics dialog box.

Tab	Option	Description
Table Options	Update table statistics	Updates table statistics.
	Do not collect column statistics	Column statistics provide information that the optimizer uses to choose the best access plans for queries.

Tab	Option	Description
	Collect column statistics on key columns only	<p>Collects column statistics on columns that make up all the indexes defined on the table.</p> <p>Column statistics provide information that the optimizer uses to choose the best access plans for queries.</p>
	Collect column statistics on all columns	<p>Collects column statistics for all columns.</p> <p>Column statistics provide information that the optimizer uses to choose the best access plans for queries.</p>
	Do not collect distribution statistics	<p>Does not collect basic statistics or distribution statistics on the columns.</p> <p>For efficiency both of RUNSTATS and subsequent query-plan analysis, you might collect distribution statistics on only the table columns that queries use in WHERE, GROUP BY, and similar clauses. You might also collect cardinality statistics on combined groups of columns. The optimizer uses such information to detect column correlation when it estimates selectivity for queries that reference the columns in the group.</p>
	Collect distribution statistics on key columns only	<p>Collects both basic statistics and distribution statistics on key columns only.</p> <p>For efficiency both of RUNSTATS and subsequent query-plan analysis, you might collect distribution statistics on only the table columns that queries use in WHERE, GROUP BY, and similar clauses. You might also collect cardinality statistics on combined groups of columns. The optimizer uses such information to detect column correlation when it estimates selectivity for queries that reference the columns in the group.</p>



Tab	Option	Description
	Collect distribution statistics on all columns	Collects both basic statistics and distribution statistics on all columns.  For efficiency both of RUNSTATS and subsequent query-plan analysis, you might collect distribution statistics on only the table columns that queries use in WHERE, GROUP BY, and similar clauses. You might also collect cardinality statistics on combined groups of columns. The optimizer uses such information to detect column correlation when it estimates selectivity for queries that reference the columns in the group.
	Frequency	Lets you specify the maximum number of frequency values to collect, between 1 and 32767.
	Quantiles	Lets you specify the maximum number of distribution quantile values to collect, between 1 and 32767.
	Exclude XML columns	Lets you collect statistics on non-XML columns only. XML columns are not included in statistics collection.
Index Options	Update index statistics	Updates index statistics.
	Collect extended index statistics	Collects extended index statistics, the CLUSTERFACTOR and PAGE_FETCH_PAIRS statistics that are gathered for relatively large indexes.
	Collect sample statistics	Rapid SQL uses a CPU sampling technique when compiling the extended index statistics. If the option is not specified, every entry in the index is examined to compute the extended index statistics.
	Select Indexes	Lets you select the indexes.
Access Options	Allow read only access during collection	Allows read only access while Rapid SQL updates the statistics.
	Allow read/write access during collection	Allows read and write access while Rapid SQL updates the statistics.

For more information, see:

[Update Statistics](#)

[Updating Table Statistics](#)

[Collecting Column Statistics](#)

[Collecting Distribution Statistics for Tables](#)

[Setting Access Levels While Updating Statistics for Indexes](#)

[Collecting Extended Index Statistics](#)

[Collecting Sample Statistics for Indexes](#)

#### UPDATING TABLE STATISTICS

To update table statistics, do the following:

#### COLLECTING COLUMN STATISTICS

To collect column statistics, do the following:

#### COLLECTING DISTRIBUTION STATISTICS FOR TABLES

To collect distribution statistics for tables, do the following:

#### SETTING ACCESS LEVELS WHILE UPDATING STATISTICS FOR INDEXES

To set access levels while updating statistics for an index, do the following:

#### COLLECTING EXTENDED INDEX STATISTICS

To collect extended index statistics, do the following:

#### COLLECTING SAMPLE STATISTICS FOR INDEXES

To collect sample statistics for indexes, do the following:

### UPDATING MICROSOFT SQL SERVER OBJECT STATISTICS

You can update statistics so that Microsoft SQL Server performs the most efficient query possible. This feature updates statistical information on your database so that the query processor can determine the optimal strategy for evaluating a query. These statistics record the key values used for distribution in an database.

You can use the Update Statistics dialog box if there is significant change in the key values in the database, if a large amount of data in an indexed column has been added, changed, or removed, or if a table has been truncated causing significant changes in the distribution of key values.

The Update Statistics dialog box lets you specify tables and indexes for validation. This dialog box offers different update options depending on your version of Microsoft SQL Server.

**TIP:** Avoid updating statistics on your target tables during busy access periods. Microsoft SQL Server locks remote tables and indexes while reading data for update statistics.

For Microsoft SQL Server version 7 or later, the Update Statistics dialog box lets you specify a full or a percentage of a full scan to be used for updating table or index statistics. It also lets you enable or disable future automatic recomputations of statistics. These recomputations are made at Microsoft SQL Server's discretion. When updating statistics for tables, this dialog box also lets you specify the type of statistics you require.

## Updating Objects

The table below describes the options and functionality on the Update Statistics dialog box for objects.

Option	Description
Scan Range	Full Scan - Select when you want index statistics on every available row.  Sample Scan - Select when database size prohibits a full scan and you can afford to rely on statistics extrapolated from a sample of all available rows.
Statistics Type	Index - Select if you only require statistics on the target tables' indexed columns.  Columns - Select if you require statistics on the target tables in their entirety.  All existing statistics - Select if you require statistics on the whole database.  NOTE: Only available for updating statistics on tables.
Statistics Recompute	Select if you want Microsoft SQL Server to recompute and update the statistics for the index as part of its normal internal maintenance.  Deselect if you want the scheduling of future recomputations to be solely your responsibility.

## Updating Databases

The table below describes the options and functionality on the Update Statistics dialog box for databases.

Option	Description
Database(s) to Update	Lets you specify the database(s) to update.

For more information, see:

[Update Statistics](#)

## UPDATING SYBASE ASE OBJECT STATISTICS

The Update Statistics dialog box lets you specify tables and indexes for validation. This dialog box offers different update options depending on your version of Sybase ASE.

**TIP:** Avoid updating statistics on your target tables during busy access periods. Sybase ASE locks remote tables and indexes while reading data for update statistics.

For more information, see:

[Update Statistics](#)**Important Notes**

None

# SQL SCRIPTING

Rapid SQL incorporates a powerful SQL scripting environment, the ISQL Editor. The ISQL Editor lets you write, debug, test and deploy solid SQL code for your database applications. The scripting environment lets you:

- Open multiple interactive script windows.
- Execute multiple scripts on the same desktop.
- Capture multiple result sets on the same desktop.

Rapid SQL's scripting environment is comprised of different windows:

- [ISQL Editor](#)
- [DDL Editor](#)
- [Results Editor](#)

These windows are context sensitive to the type of script you are opening or extracting. For example, if you extract the schema for a table, Rapid SQL opens a DDL Window containing the script. If you execute a script, a result window displays containing the results of your execution statement.

**TIP:** Since you must drop a database object before you can recreate it, you can set the DDL Editor to automatically include DROP statements for specified objects.

**TIP:** You can have multiple ISQL windows open at the same time, with each running separate queries.

## ISQL EDITOR

The ISQL Editor includes the ISQL Window and DDL Editor.

The ISQL Window lets you:

- Insert files and open files.
- Rename and configure query tabs.
- Find and replace with regular expressions.

- Mail your script files.

**TIP:** To enlarge or reduce (zoom) the display font size, press Ctrl+= or Ctrl-- (plus or minus on the numeric pad), or hold the Ctrl key while scrolling the mouse wheel. To return to the default size, press Ctrl-/ (on the numeric pad).

**TIP:** To toggle to the next SQL window, press CTRL +T.

**TIP:** The row limit option lets you display only the first 'n' rows from any submitted query.

**TIP:** For Oracle, Rapid SQL displays REF CURSOR contents in the ISQL Window and [Results Tab](#).

**TIP:** For IBM DB2 LUW for Linux, Unix, and Windows and Oracle, you can access the Code Completion functionality with the CTRL+K shortcut.

## Related Topics

- [Toolbar Options](#)
- [Opening ISQL Windows](#)
- [Opening DDL Editors](#)
- [Opening Script Files](#)
- [Inserting Files into an ISQL Window](#)
- [Splitter Windows](#)
- [Find and Replace in an ISQL Window](#)
- [Regular Expressions](#)
- [Navigating in an ISQL Window](#)
- [Scheduling](#)
- [Sending SQL Scripts](#)
- [Renaming and Closing Query Window Tabs](#)
- [Printing a Script](#)
- [Saving and Closing Scripts](#)
- [SQL Preprocessor](#)
- [Automated Error Detection and Coding Aid Features](#)

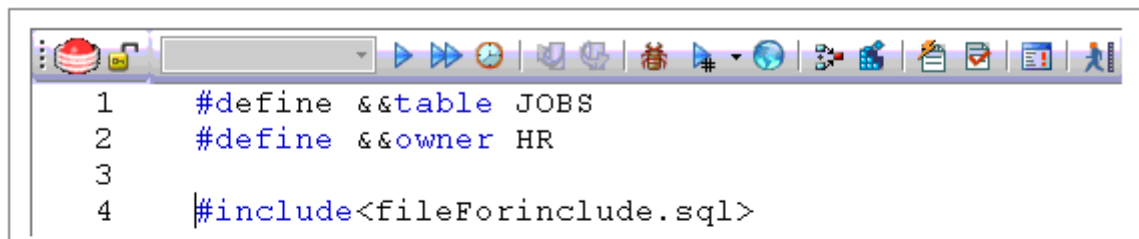
## PREPROCESSING #DEFINE AND #INCLUDE DIRECTIVES

Rapid SQL provides SQL preprocessing similar to functionality provided by C language compiler directives. The ISQL Editor supports a simplified version of the following directives.

- **#include** provides a means to include the contents of a file in a script at the location of the directive
- **#define** provides a simple, global search and replace function within a script

The following figure illustrates the result of basic preprocessing of a script, nesting of **#define** directives/references, and the notations required by the ISQL editor. The original script includes two **#define** directives and a **#include** reference to a one-line file named **fileForInclude.sql**. The referenced file includes two identifiers to be replaced with **#define** processing.

Script content before SQL preprocessing

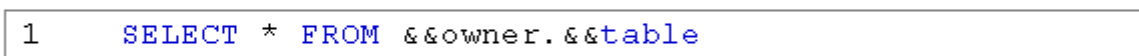


```

1  #define &&table JOBS
2  #define &&owner HR
3
4  #include<fileForInclude.sql>

```

Contents of fileForInclude.sql

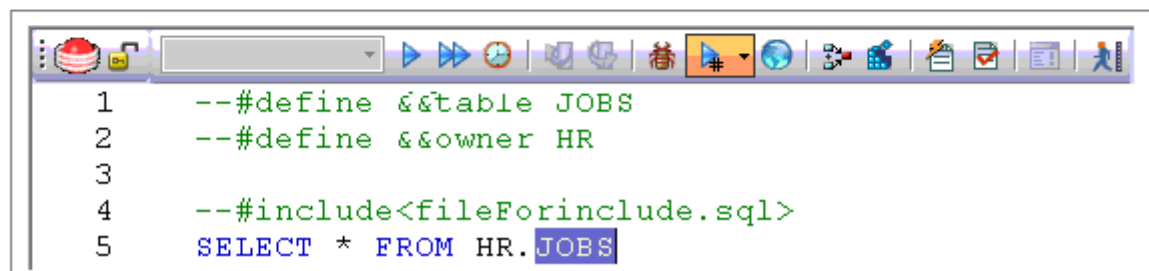


```

1  SELECT * FROM &&owner.&&table

```

Script content produced by SQL preprocessing



```

1  --#define &&table JOBS
2  --#define &&owner HR
3
4  --#include<fileForInclude.sql>
5  SELECT * FROM HR.JOBS

```

The key steps in working with SQL preprocessing are:

- **Preparing the ISQL editor for #include and #define preprocessing** - Prior to using SQL preprocessing features, you should set the paths that will be searched in processing **#include** directives. For details, see [Setting Up Rapid SQL to Preprocess #include Directives](#).
- **Using #define and #include directives in scripts** - While the supported directives approximate typical C language #define and #include functionality, there are differences in functionality and required syntax. For example, identifiers in the **#define** directive and in all instances to be replaced, must be prefixed with two ampersand characters (**#define &&PI 3.14159**). For detailed information, see [#include Functionality and Syntax](#) and [#define Functionality and Syntax](#).

- **Preprocessing and executing scripts containing #define and #include directives** - The ISQL Editor offers two preprocessing options. You can have a script preprocessed without being executed, opening the processed script in a new editor window with all **#define** and **#include** substitutions made. This lets you view the processed script before execution or continue working with the processed SQL. Alternatively, you can have the script preprocessed and executed in a single step. For details, see [Preprocessing and Executing Scripts Containing #define and #include Directives](#).

#### SETTING UP RAPID SQL TO PREPROCESS #INCLUDE DIRECTIVES

In ISQL editor processing of a **#include** directive, Rapid SQL searches the following locations, in the following order, for the specified file:

- 1 The location specified on the **Datasource Properties** tab of the Datasource Registration Wizard/Editor. For details, see [Registering or editing datasources](#).
- 2 The location specified on the Directories tab of the Options editor. For details, see [Directories Options](#).

**NOTE:** For detailed information on setting Rapid SQL options, see [Specifying Rapid SQL application and feature options](#).

Before using SQL preprocessing, ensure that server side or server side and local search paths for include files are specified.

#### #INCLUDE FUNCTIONALITY AND SYNTAX

Rapid SQL support for the **#include** directive provides a means to include the contents of a file in a script at the location of the directive. For example, if a script contains the following:

```
#include mydeclarations.sql
```

then on preprocessing of the script, there are two effects:

- The line containing the **#include** directive is commented out before the script is sent to the database
- The text in the file mydeclarations.sql is placed in the script following the commented out line with the **#include** directive.

The **#include** directive is supported for simple file names only. Supported syntax of the **#include** directive for use in the ISQL editor, Procedure Object Editor, or Package Body Object Editor, is as follows:

```
#include <filename.ext>
```

where:

- *filename.ext* is a simple filename and extension

**NOTE:** For those familiar with C compiler functionality, angle bracket and quoted forms are supported only indirectly. While **#include <filename.ext>** and **#include "filename.ext"** forms are valid, they are functionally equivalent to the **#include filename.ext**. Using the angle bracket or quoted forms has no effect on locations searched for the target file.



Searches are performed in the locations specified in the setup for this feature. For details, see [Setting Up Rapid SQL to Preprocess #include Directives](#).

Error processing is as follows:

- If the preprocessor fails to include the specified file, it displays an error message noting the reason for the failure (such as the file does not exist, insufficient permissions on the file, or file too large ). Preprocessing or execution of the script cannot continue until the error is corrected.
- If the file is found in the first search location specified in [Setting Up Rapid SQL to Preprocess #include Directives](#) but cannot be opened (permission denied for example), no attempt will be made to locate the file in the second specified search location.

#### #DEFINE FUNCTIONALITY AND SYNTAX

The **#define** directive provides a simple, global search and replace function within a script. For example, if a script contains the following:

```
#define &&PI 3.14159
```

then on execution of the script, there are two effects:

- All instances of **&&PI** in the script would be replaced by **3.14159** on execution of the script
- The line containing the **#define** directive is commented out before the script is sent to the database

The **#define** directive is supported for symbolic constants only. Supported syntax of the #define directive for use in the ISQL editor, Procedure Object Editor, or Package Body Object Editor, is as follows:

```
#define &&Identifier Replacement_text
```

where:

- *Identifier* is any character string appearing in the script
- *Replacement\_text* is the string that will replace all instances of the string specified by the *Identifier* argument. Valid values are strings, numbers or combinations consisting of the digits **0-9**, characters **a-z**, characters **A-Z**, and the underscore character.

**NOTE:** In addition to the actual **#define** directive appearing in a script, the ampersand notation is also required in all references that are to be replaced. References that are not prefixed with ampersand characters are not processed.

Nested **#define** directives are also supported. For example if a script contains the following:

```
#define &&myTable Clients

#define &&embtClients Embarcadero
#define &&tempTable New&&myTable
#define &&embtTempTable &&embtClients&&myTable

Select * from &&tempTable;
Select * from &&embtTempTable
```

then after preprocessing, the contents of the script would be as follows:

```
Select * from NewClients;
Select * from EmbarcaderoClients
```

## PREPROCESSING AND EXECUTING SCRIPTS CONTAINING #DEFINE AND #INCLUDE DIRECTIVES

The ISQL editor offer two preprocess/execute modes:

- Rapid SQL can preprocess the script, opening the script in a new tab, with all **#define** and **#include** substitutions made. This lets you view your preprocessed script or continue with edits after preprocessing, before executing it in the new tab.
- Rapid SQL can preprocess and execute the script in a single step.

Preprocessing or preprocessing/executing the script consists of selecting a mode and the executing the script. Preprocessing mode is controlled by the preprocessing dropdown on the ISQL editor toolbar.



### To preprocess a script and have the preprocessed script opened in a new tab:

- 1 From the Preprocess dropdown, select **Pre-Process Only**. The Preprocess dropdown icon takes on a distinctive appearance to indicate **Pre-Process Only** mode.



- 2 On the ISQL toolbar, click the Execute button.



The script with all #define and #include replacements made, opens in a new ISQL editor.

**NOTE:** For details on error processing and specific handling of directives, see [#include Functionality and Syntax](#) and [#define Functionality and Syntax](#).

### To preprocess and execute a script in a single step:

- 1 From the Preprocess dropdown, select **Pre-Process and Execute**.
- 2 On the ISQL toolbar, click the Execute button.

The script executes. While the script sent to the server for execution includes all **#define** and **#include** substitutions, the new **Query** tab contains the original, unprocessed script.

For related information, see the following topics:

- See [Executing Scripts](#) for other script execution options
- See [Preprocessing #define and #include Directives](#) for an introduction to SQL preprocessing in the ISQL editor

## TOOLBAR OPTIONS

The table below describes the options of the ISQL Editor toolbar:

Option	Description
Lock	Lets you lock an ISQL Window to a specific data source connection. When locked, the ISQL Window does not respond to data source change events in other windows, and only uses the current connection for subsequent execution.
Execute	Executes the script.
Step Execute	Initiates step execution of the script. For more information, see <a href="#">Step Execution Facility</a> .
Schedule	Opens a dilaog that lets you schedule an SQL job. For more information, see <a href="#">Scheduling</a> .
SQL Begin Transaction	Toggles the Begin Transaction status on and off.
SQL Rollback	Performs a rollback on the current transaction.
SQL Commit	Commits the current transaction.
Debug	Opens the Embarcadero SQL Debugger.
Preprocess	Lets you set a preprocessing state for execution of the script.  If you select <b>Pre-process Only</b> , subsequent use of the Execute button will result in the script being preprocessed without beng executed. The script opens n a new editor window with all #define and #include substitutions made.  If you select <b>Pre-process and Execute</b> , all #include and #define directives are carried out and the script is executed. No new editor window opens.  For more information, see <a href="#">Preprocessing #define and #include Directives</a> .
Script Execution Facility	Provides a shortcut to opening the Script Execution Facility. For more information, see <a href="#">Executing Scripts Against Multiple Datasources</a> .
Query Plan	Activates and deactivates Query Plan mode. For more information, see <a href="#">Using the Query Plan Facility</a> .
Query Options	Provides a shortcut to opening the Query Options dialog. For more information, see <a href="#">Query Options</a> .
Format	Transforms spaghetti-style written SQL code into an easier read format.
Syntax Check	Initiates a syntax check. This is only necessary if SQL parsing is disabled and autmatic syntax checking is therefore disabled. For more information, see <a href="#">Syntax Checking</a> .
Errors	Closes the Errors tab.

## ISQL WINDOWS

A script is a collection of SQL statements used to perform a task. Scripts, stored as files, form the foundation of most application projects. Definitions of tables, indexes, stored procedures, as well as reports and batch jobs, are usually stored in script files and shared by groups of developers using version control software. You can manage SQL scripts in SQL windows. Rapid SQL lets you open multiple SQL windows in one or more workspaces.

## Opening ISQL Windows

To open the ISQL Window, do the following:

- 1 On the *File* menu, click New ISQL.

OR

On the Main tool bar, click New ISQL.

Rapid SQL opens an SQL window in your current workspace.

For more information, see

[ISQL Window Status Bar](#)

[ISQL Editor](#).

### ISQL WINDOW STATUS BAR

The ISQL window Status bar lets you view:

- Auto commit status (Oracle) - Automatically commits SQL statements as soon as the statements are run.
- Begin Transaction ON/OFF (SQL Server and Sybase)

**TIP:** For Microsoft SQL Server and Sybase, to set Begin Transaction status to "Yes", on the ISQL Window toolbar, click the SQL Begin Tran button.

**TIP:** For Oracle, you can apply auto commit status changes to all open ISQL windows. You can modify the Oracle Auto Commit status and ISQL tab of the Options Editor. For details, see [ISQL Options](#).

### AUTOMATED ERROR DETECTION AND CODING AID FEATURES

DBArtisan analyzes code as you add content to an ISQL editor session and offers the following automated features:

- [Syntax Checking](#) automatically flags syntactical errors in your scripts
- [Semantic Validation](#) detects object name references to objects not present in the datasource index.
- [Code Complete](#) lets you insert or replace object names, selected from suggestion lists, as you edit a script.

### SYNTAX CHECKING

The ISQL editor can perform on-the-fly or manually-initiated syntax checking. DBArtisan checks the syntax of the current contents of an ISQL Editor window against the SQL dialect native to the datasource to which the ISQL editor session is connected. Syntax checking can be performed regardless of whether an ISQL editor session is locked to a datasource.

For each syntax error detected, DBArtisan annotates your script and offers assistance as follows:

- Line numbers for lines containing syntax errors are flagged with an error icon.
- Hovering the mouse over an error icon displays a tooltip with an error message
- The specific error in the line of code is underlined in red

```

22 BEGIN
23 v_start_year := TO_NUMBER(TO_CHAR(p_start_date, 'Y
24      Encountered unexpected token "v_start_year"
25
26 p_years_elapsed := v_end_year - v_start_year;
27
28 END YEARS_ELAPSED;

```

**NOTE:** A line-by-line listing of individual syntax errors is also available in the ISQL editor's Error pane. For more information, see [About the Error pane](#).

DBArtisan can be configured to perform syntax checks automatically or to require a syntax check to be initiated manually:

- **Automatic** - Automatic syntax checking is enabled from the Options Editor. For details, see [ISQL Options - Code Assist Tab](#).

If the **Enable Real-time syntax checking** check box is selected, a syntax check is performed whenever there is an interval of 1.5 seconds or more between key strokes. Syntax error annotations persist until the error is corrected and the next automatic syntax check is executed.


- **Manual** - If automatic syntax checking is disabled, a syntax check is only performed when you click the Syntax button on the ISQL editor toolbar:



When automatic syntax checking is disabled, syntax error annotations persist until the error is corrected and you explicitly run another syntax check.

## SEMANTIC VALIDATION

Semantic validation is an on-the-fly ISQL editor feature that verifies that object names are correctly specified. It ensures that the names of supported object types (columns, tables, synonyms, and views) present in a script match those for the connected datasource. Minimizing errors associated with typographical errors or references to obsolete object definitions, DBArtisan analyzes the names of supported object types as you type, and raises an error condition when it detects a name not present on the datasource.



A semantic error can indicate one of the following conditions:

- The object name as specified in the script is not present on the database
- The referenced object is hidden by a SQL filter
- The schema information that Rapid SQL obtains to implement this feature is outdated. This can happen if other users or applications modify the schema during your ISQL editor session. For information on refreshing this information, see [ISQL Options - Code Assist Tab](#).

When making use of this feature, always consider the following points

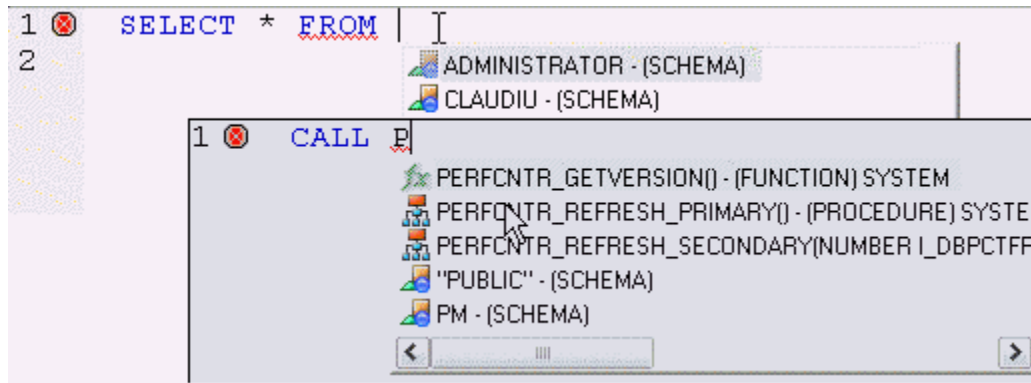
- For SQL Server and Sybase datasources, statements that contain temporary table names (that is, tables whose name begins with the # character) are ignored by the semantic validation feature.
- Semantic errors in a statement are not displayed for a statement that currently has outstanding syntax errors. For more information, see [Syntax Checking](#).
- For SQL Server and Sybase datasources, system table names (that is tables whose name begins with the # character) are ignored by the semantic validation feature.
- While semantic errors can be flagged with an ERROR severity level, they can be more easily distinguished from syntax errors by setting a severity level of WARNING. Similarly, if you do not want semantic errors flagged, you can set the severity level to IGNORE. For details, see [ISQL Options - Code Assist Tab](#).
- Ensure that the ISQL Editor window is connected to a datasource. Rapid SQL requires a datasource name in order to resolve an object name.

In the case of SQL Server and Sybase, you must also select a database. The exception to this requirement is scripts that have a USE statement, specifying a database, appearing before any object name references.

**NOTE:** A line-by-line listing of individual semantic errors is also available in the ISQL editor's Error pane. For more information, see [About the Error pane](#).

## Code Complete

Code Complete lets you quickly and accurately write DML and call/execute statements by providing a fast and intuitive lookup for database objects. Code Complete lets you select from a suggestion box that lists objects appropriate at the cursor location within a statement. Code Complete offers intelligent suggestions in providing simple object names or in constructing fully-qualified names.



This feature becomes available at all points in INSERT, UPDATE, DELETE, SELECT, CALL, and EXEC statements, including all legal clauses, where a reference to the name of one of the following object types is valid:

- Columns
- Tables
- Views
- Functions
- Procedures
- Packages
- Synonyms

When the Code Complete feature is enabled, it can be invoked in two ways:


- Automatically, if auto-activation is enabled. A Code Complete suggestion box is offered when you start typing an object name in a relevant location and then stop typing for an interval that exceeds the specified auto-activation delay.
- Manually, if auto-activation is disabled or if the specified auto-activation delay is sufficiently large, by pressing CONTROL+SPACE. or typing a period (.).

**NOTE:** By default, Code Complete is enabled and configured for automatic invocation. For information on enabling or disabling this feature and configuring auto-activation and other Code Complete preferences, see [ISQL Options - Code Assist Tab](#).

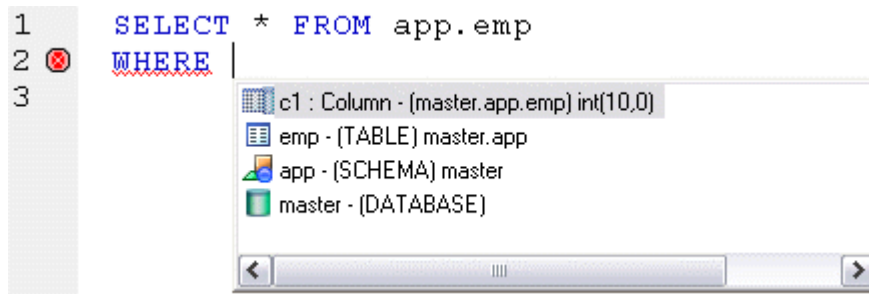
Code Complete offers context sensitive suggestions based on the cursor location within the statement when Code Complete is invoked. For example:

- **When invoked without any preceding text** - Code Complete will supply basic suggestions appropriate to that clause.

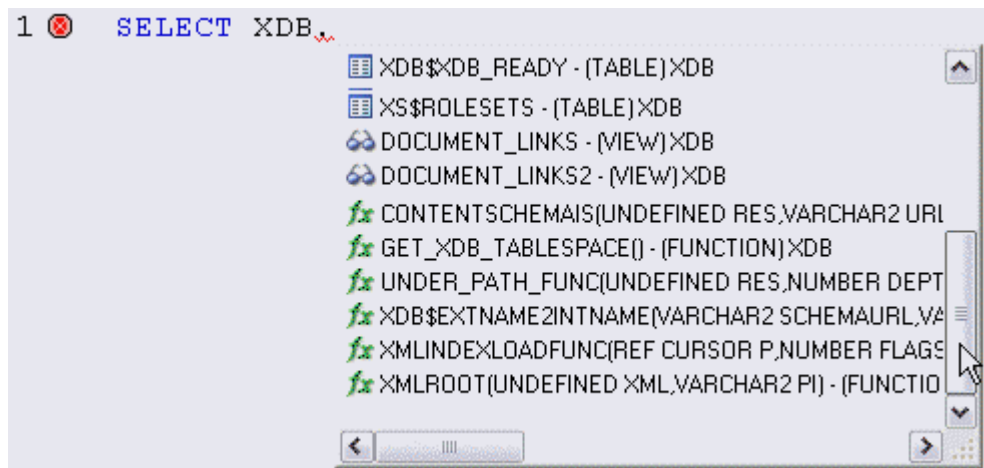
For example, suggestions for the table name in a FROM clause would help a user construct a qualified table or view name. Against a SQL Server datasource for example, this would include schemas and databases.



Suggestions for the column name in a WHERE clause in the same statement would contain the schemas, databases and tables\views referenced in the FROM clause as well as a listing of the applicable columns based on the referenced tables\views.




- **After the delimiter within an object name** - Code Complete offers a list that lets you select the next element to build a fully qualified object name.








- **After typing one or more characters of an object name** - Code Complete offers a listing of the names of all object name elements, relevant to the cursor location within the object name, that start with the typed characters.



- **At any point within a partially-specified or complete object name** - Code Complete offers all selections appropriate to the location of the cursor.



In addition, you can type additional characters while the suggestion list is available to further filter the items appearing in the list.

#### To insert or replace a suggestion:

- 1 Double-click an object name or object name element in the list. Alternatively, you can use the arrow keys and press ENTER.

#### To dismiss the suggestion list without making a selection:

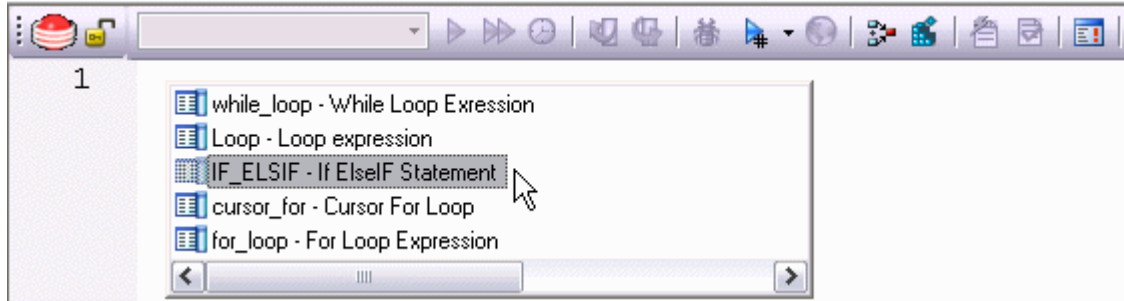
- 1 Click anywhere outside the suggestion list or press ESC.

**NOTE:** The schema information that Rapid SQL obtains to implement this feature can be outdated by the actions of other users or applications during an ISQL editor session. For information on refreshing this information, see [ISQL Options - Code Assist Tab](#).

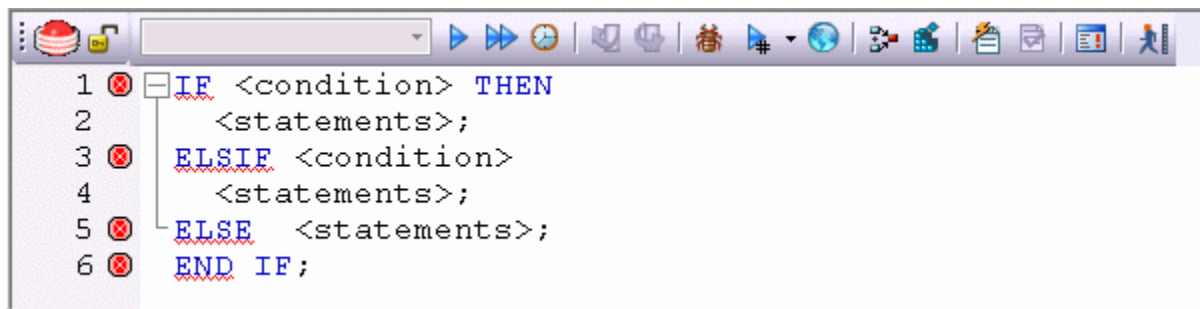
### USING CODE TEMPLATES

Code templates are complete code blocks that can be easily added to open windows or scripts with a few keystrokes. Templates let you define standard or common code units such as loops, comment blocks, query structures, and so on, add them to your code and edit them accordingly.

When you activate Code Assist by typing CTRL+SPACE, the Code Assist menu opens, offering any defined code templates for the currently connected DBMS platform.



Selecting the short name for the code template and pressing RETURN adds the template at the cursor position.



Rapid SQL loads a default set of code templates at startup but you can also add and delete templates from the currently available set. In addition, you can save sets of templates to file and subsequently load specific template sets, allowing you to customize your templates to different platforms or development projects.


Before making use of this feature, you should be familiar with the following tasks:

- Enabling the code templates feature. For details, see [Enabling the Code Templates and Auto Replace Features](#).
- Maintaining the list of available templates. For details, see [Maintaining Code Template Definitions](#) and [Loading and Saving Custom Code Workbench Settings](#)


## USING AUTO REPLACE EXPRESSIONS

Rapid SQL lets you define shortcuts consisting of a few characters that represent longer character strings. Instances of these Auto Replace expressions are automatically replaced by the replacement string on activation events such as typing SPACE, TAB, or RETURN. This feature is useful for creating shortcuts for one-line commands or SQL statement subsets, or even to detect and fix common typographical errors such as **teh** for **the**.

For example, consider an Auto Replace definition with an expression of **sel** to represent **Select \* From:**



If the associated activation event includes a SPACE, then on typing **sel** followed by typing SPACE, the following replacement occurs.



Rapid SQL loads a default set of Auto Replace definitions at startup but you can also add, edit, and delete Auto Replace definitions. In addition, you can save sets of definitions to file and subsequently load specific Auto Replace definitions, allowing you to customize your templates to different platforms or development projects.

Before making use of this feature, you should be familiar with the following tasks:

- Enabling the Auto Replace feature. For details, see [Enabling the Code Templates and Auto Replace Features](#).
- Creating and maintaining the list of available Auto Replace definitions. For details, see [Maintaining Auto Replace Expressions](#) and [Loading and Saving Custom Code Workbench Settings](#)

## DDL EDITORS

**NOTE:** This functionality is available for all platforms.

Rapid SQL lets you open a new DDL Editor when you want to create a script that is tied to an object type in your database. The DDL Editor opens containing a template script for the selected object type. Because the DDL Editor is directly tied to a database, database warning messages can be issued. For example, if you have a create table script which includes a DROP TABLE statement, the DDL Editor warns you about the existence of this statement and that you could lose existing table data.

## Important Notes

None

The table below describes the options and functionality on the Create New Database Object dialog box.

Option	Description
Object Type	Lets you select the object type to which you want to attach the script.
Owner	Lets you type the name of the object owner for the object. The name of the owner connected to the current datasource is used as the default.
Object Name	Lets you type the name of the object type.

Rapid SQL opens a DDL Editor containing a template for the object type you selected.

For more information, see [Opening DDL Editors](#).

## OPENING DDL EDITORS

To open DDL Editors, do the following.

- 1 On the File menu, click *New DDL Editor*.

OR

On the Main tool bar, click the Down arrow on New, and then click DDL Editor.

- 2 In Object Type select the object type to which you want to attach the script.
- 3 In Owner type the name of the object owner for the object. The name of the owner connected to the current datasource is used as the default.
- 4 In Object Name type the name of the object type.
- 5 Click OK.

Rapid SQL pastes the DDL into the ISQL Editor.

For more information, see [ISQL Editor](#).

## OPEN FILES

**NOTE:** This functionality is available for all platforms.

The Open Files dialog box lets you open existing files.

**Important Notes**

None

The table below describes the options and functionality on the Open Files dialog box.

Option	Description
File Name	Lets you type a script file name. You can open multiple script files simultaneously by selecting multiple file names in the file list.
Files of Type	Lets you control the files displayed in the file list by changing the selection. Rapid SQL displays multiple default file extensions.
Open Button	Click to open one or more files into the current workspace.

**Completing the Open Files Dialog Box**

To complete the Open Files dialog box, do the following:

- 1 On the File menu, click Open.  
OR  
On the Main tool bar, click Open.
- 2 Select a script file.  
OR
- 3 In File Name, type a script file name.
- 4 In Files of type, select types of files to display.
- 5 Click Open to open one or more files into the current workspace.  
Rapid SQL pastes the script into the ISQL Editor.

For more information, see [ISQL Editor](#).

**WHAT TYPE OF FILE**

**NOTE:** This functionality is available for all platforms.

The What Type of File Is dialog box lets you select options for unknown file types.

**Important Notes**

None

The table below describes the options and functionality on the What Type of File Is dialog box:

Option	Description
The file is a general purpose SQL script	Select if the unknown file is a SQL script.
The file contains the DDL to create a database object of the file contains Oracle Anonymous PL/SQL.	Lets you select object type, type the owner, and object name.
Always open unknown files into a SQL window without prompting.	Select to hide What Type of File Is Dialog Box for future unknown file types.

For more information, see [Completing the What Type of File Is Dialog Box](#).

#### COMPLETING THE WHAT TYPE OF FILE IS DIALOG BOX

To complete the What Type of File Is dialog box, do the following:

- 1 On the *File* Menu, click *Open*.  
Rapid SQL opens the Open File(s) dialog box.
- 2 In the *Open File(s)* dialog box, select the target script, and then click *Open*.  
Rapid SQL opens the What type of file dialog box.
- 3 Select options, and then click OK.  
Rapid SQL opens the target script in an SQL Editor.

For more information, see [ISQL Editor](#).

#### INSERT FILE INTO CURRENT FILE

**NOTE:** This functionality is available for all platforms.

The ISQL Editor facilitates the reuse of SQL scripts by letting you insert an existing file into another script.

The table below describes the options and functionality on the Insert File into Current File dialog box.

Option	Description
File Name	Lets you type a script file name. You can open multiple script files simultaneously by selecting multiple file names in the file list.

Option	Description
Files of Type	Lets you control the files displayed in the file list by changing the selection. Rapid SQL displays multiple default file extensions.
Open Button	Click to insert one or more files into the current workspace.

## Important Notes

None

For more information, see [Completing the Insert File into Current File Dialog Box](#).

### COMPLETING THE INSERT FILE INTO CURRENT FILE DIALOG BOX

To complete the Insert File into Current File dialog box, do the following:

- 1 On the Edit menu, click Edit Insert File.  
OR  
On the *Edit* tool bar, click Edit Insert File.  
Rapid SQL opens the Insert File into Current File dialog box.
- 2 Select a script file.  
OR
- 3 In File Name, type a script file name.
- 4 In Files of type, select types of files to display.
- 5 Click Open to insert one or more files into the current workspace.  
Rapid SQL inserts the file.

For more information, see [ISQL Editor](#).

## SPLITTING WINDOWS

You can split a SQL window into four different screens so that you can view different sections of a document simultaneously. You have the option to split the ISQL Window horizontally, vertically or into quadrants.

### Splitting the ISQL Window Horizontally

To split the ISQL Window horizontally, do the following:

- 1 Point to the split box at the top of the vertical scroll bar on the right side of the SQL window.
- 2 When the pointer changes, drag it to the desired position.

## Splitting the ISQL Window Vertically

To split the ISQL Window vertically, do the following:

- 1 Point to the split box at the top of the horizontal scroll bar on the bottom left side of the SQL window.
- 2 When the pointer changes, drag it to the desired position.

## Removing Splits from an ISQL Window

Double-click the appropriate split bar to return the SQL window to its normal state.

For more information, see [ISQL Editor](#).

## FIND

**NOTE:** This functionality is available for all platforms.

The Find dialog box lets you search text in your SQL scripts.

## Important Notes

None

The table below describes the options and functionality on the Find dialog box.

Option	Description
Find What	Lets you type your search string.
Match whole word only	Select to search for only the complete word.
Match Case	Select to make the search case-sensitive.
Regular expression	Select if you are looking for a <a href="#">regular text expression</a> .
Wrap around search	Lets you search from the end of the script and back to the insertion point.
Direction	Lets you specify the direction you want to search. Click the Up or Down option button.
Find Next Button	Click to find the next occurrence of your search string.
Mark All Button	Click to place a small blue dot next to every line number in the script which meets the required search string criteria.

For more information, see [Completing the Find Dialog Box](#).



## COMPLETING THE FIND DIALOG BOX

To complete the Find dialog box, do the following

- 1 On the Edit menu, click Find.
- OR  
On the *Edit* tool bar, click Find.
- 2 In Find What, type your search string.
- 3 Select Match whole word only to search for only the complete word.
- 4 Select Match Case to make the search case-sensitive.
- 5 Select Regular expression to search for a regular text expression.
- 6 Select Wrap around search to search from the end of the script and back to the insertion point.
- 7 In Direction, click Up or Down.
- 8 Click Find Next to find the next occurrence of your search string.
- 9 Click Mark All to place a small blue dot next to every line number in the script which meets the required search string criteria.

The ISQL Editor highlights the object name if it matches the search criteria.

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

## REPLACE

**NOTE:** This functionality is available for all platforms.

The Replace dialog box lets you search and replace text in your SQL scripts.

The table below describes the options and functionality on the Replace dialog box.

Option	Description
Find What	Lets you type your search string.
Replace With	Lets you type the replacement text.
Match Case	Select to make the search case-sensitive.
Regular expression	Select if you are looking for a <a href="#">regular text expression</a> .
Wrap around search	Lets you search from the end of the script and back to the insertion point.
Direction	Lets you specify the direction you want to search. Click the Up or Down option button.

Option	Description
Find Next Button	Click to find the next occurrence of your search string.
Replace Button	Click replace the current selection.
Replace All Button	Click to automatically find and replace all occurrences of your search string within the current window.

For more information, see [Completing the Replace Dialog Box](#).

#### COMPLETING THE REPLACE DIALOG BOX

To complete the Replace dialog box, do the following:

- 1 On the Edit menu, click Replace.
- OR
- On the Edit tool bar, click Replace.
- 2 In Find What, type your search string.
- 3 In Replace With, type the replacement text.
- 4 Select Match whole word only to search for only the complete word.
- 5 Select Match Case to make the search case-sensitive.
- 6 Select Regular expression to search for a regular text expression.
- 7 Select Wrap around search to search from the end of the script and back to the insertion point.
- 8 In Direction, click Up or Down.
- 9 Click Find Next to find the next occurrence of your search string.
- 10 Click Replace to replace the current selection.
- 11 Click Replace All to automatically find and replace all occurrences of your search string within the current window.

For more information, see [ISQL Editor](#).

## REGULAR EXPRESSIONS

Regular Expressions are offered as an optional search criteria in the SQL windows search facility. Regular Expressions serve as powerful notation for describing string matching patterns. Special characters are used to denote certain match criteria on which the ISQL Editor should conduct its search. The table below describes the special characters and their meanings:

Character	Meaning	Example
^	Circumflex - Constrains the search to the start of a line	^Rap -- Match lines beginning with Rap ^Emb -- Match lines beginning with Emb
\$	Dollar - A dollar as the last character of the string constrains matches to the end of lines.	if\$--Match lines ending with if ^end\$-- Match lines consisting of the single word end
.	Period - A period anywhere in the string matches any single character.	T.l -- Matches Tool, Till, Tail etc. H.w -- Matches Huw, How, Haw etc. ^Sin.ers -- Matches lines beginning with Sindes, Sinners etc.
*	Asterisk - An expression followed by an asterisk matches zero or more occurrences of that expression.	to* -- Matches t, to, too etc. 00* -- matches 0, 00, 000, 0000 etc.
+	Plus - An expression followed by a plus sign matches one or more occurrences of that expression.	to+ -- Matches to, too etc. 10+ -- Matches 10, 100, 1000, 10000 etc. /(:d+)/ -- Matches (0), (12464), (12) etc.
?	Question mark - An expression followed by a question mark optionally matches that expression.	for? -- Matches f and for 10? -- Matches 1 and 10
()	Brackets - Brackets can be used to group characters together prior to using a * + or?.	Rap(id)? -- Matches Rap and Rapid B(an)*a -- Matches Ba, Bana and Banana
[ ]	Square brackets - A string enclosed in square brackets matches any character in that string, but no others. If the first character of the string is a circumflex, the expression matches any character except the characters in the string. A range of characters can be specified by two characters separated by a -. These should be given in ASCII order (A-Z, a-z, 0-9 etc.).	{[0-9]} -- Matches {0}, {4}, {5} etc. /([0-9]+)/ -- Matches (100), (342), (4), (23456) etc. H[uo]w -- Matches Huw and How Gre[^py] -- Matches Green, Great etc. but not Grep, Grey etc. [z-a] -- Matches nothing ^[A-Z] -- Match lines beginning with an upper-case letter

Character	Meaning	Example
\	Backslash - A backslash quotes any character. This allows a search for a character that is usually a regular expression specifier.	\\$ -- Matches a dollar sign \$ \+ -- Matches a +

For more information, see [ISQL Editor](#).

## GOTO

**NOTE:** This functionality is available for all platforms.

The Goto dialog box lets you move to a specific line or column in your script.

### Important Notes

None

The table below describes the options on the Goto dialog box:

Option	Description
Line Number	Lets you type or select the target line number.

### Completing the Goto Dialog Box

To complete the Goto dialog box, do the following:

- 1 On the Edit menu, click *Goto*.  
Rapid SQL opens the Go To dialog box.
- 2 In Line Number, type or select the target line number.
- 3 Click OK.

Rapid SQL moves the cursor to the target line.

For more information, see [ISQL Editor](#).

## SENDING SQL SCRIPTS

If you have MAPI-compliant E-mail software installed on your computer, then you can send SQL scripts to other users.

To send a SQL script, do the following:

- 1 On the File menu, click Send.

OR

On the Main tool bar, click Send.

Rapid SQL opens your E-mail application.

- 2 In the *Address* box, type the name of the addressee(s) and any other options.

**NOTE:** The ISQL Editor automatically attaches a file containing your SQL script to the e-mail message.

- 3 Click *Send*.

Rapid SQL sends the result set to the specified addressee(s).

For more information, see [ISQL Editor](#).

## RENAMING AND CLOSING QUERY WINDOW TABS

SQL windows are tabbed windows that can be closed or renamed.

To rename a Query Tab, you need an open SQL window that includes an executed script. For information on how to execute scripts, see:

- [Executing Scripts](#)
- [Script Execution Facility](#)

### Renaming a Query Window Tab

To rename a Query Window Tab, do the following:

- 1 Right-click the Query Tab on the *SQL window*, and then click *Rename*.

Rapid SQL opens the Rename Tab dialog box.

**NOTE:** The Query Tab can be located on the top or bottom of the SQL window. You can set the location of the tab when configuring Datasource options.

- 2 In the *New Name* box, type the name of the new Query window.
- 3 Click *OK*.

Rapid SQL changes the name and closes the Rename Tab dialog box. The new name of the tab displays directly on the Query Window Tab at the top of the window.

## Closing a Query Window Tab

To close a Query Window Tab, do the following:

- 1 At the top of the *ISQL window*, right-click the Query Tab, and then click Close or Close All.

Rapid SQL closes the Query.

For more information, see [ISQL Editor](#).

## PRINT

**NOTE:** This functionality is available for all platforms.

The Print dialog box lets you can specify a range of panels to print, or print the contents of a script window to a file.

## Important Notes

None

The table below describes the options and functionality on the Print dialog box.

Tab	Option	Description
Scope	Printer	Displays default printer.
	Print Range	Lets you select the appropriate print range.
	Number of Copies	Lets you click the Up or Down arrow or type the number of copies you want.
Page Setup	Header	Lets you type header type to display at the top of the page.
	Footer	Lets you type header type to display at the bottom of the page.
	Header/Footer not within Margins	Select to position header and footer outside the margins.
	Margins	Lets you specify margins in either inches or centimeters.
Options	Chromacoding	Lets you select Use Color if you have a color printer. Lets you select Use Font if script includes italics and bold fonts.
	Line Numbering	Lets you specify the interval between numbered lines.
	Other Options	Lets you select other options.
Documents	Document Box	Lets you select documents to print.
	Clear	Click to clear list.
	Invert	Click to switch printing order.

Tab	Option	Description
Configurations	New Configuration Name	Lets you type a new configuration which saves you current settings, and then click Create.
	Delete	Lets you delete an existing configuration.
	Load	Lets you load an existing configuration.
	Update	Lets you update an existing configuration.

The table below describes the buttons on the Print dialog box:

Button	Description
Save Settings	Lets you save settings.
Font	Lets you open the Font dialog box and select fonts.
Setup	Lets you open the Print Setup dialog box and select printer options.
Preview	Lets you open the Print Preview dialog box that lets you preview the document before you print it.
Print	Lets you print the document.

For more information, see [Completing the Print Dialog Box](#).

#### COMPLETING THE PRINT DIALOG BOX

To complete the Print dialog box, do the following:

- 1 On the File menu, click Print.  
OR  
On the Main tool bar, click Print.
- 2 On the tabs, select options.
- 3 Click Print.

For more information, see [ISQL Editor](#).

#### SAVING AND CLOSING SCRIPTS

Untitled scripts are named SQL1...SQLn by default, and all scripts are saved with the SQL extension. You can save your files with any valid name.

For more information, see:

[Saving a Script](#)

[Closing a Script](#)

[About the Error pane](#)

[ISQL Editor](#)

#### SAVING A SCRIPT

To complete the Save dialog box, do the following:

- 1 On the File menu, click Save.

OR

On the Main tool bar, click Save.

Rapid SQL opens the Save As dialog box.

- 2 If this is a new file, in the File Name box, type the name of the file.
- 3 If this is an existing file and you want to use save as, on the Main menu, click Save As, and in the File Name box, type the name of the file.
- 4 Click Save.

Rapid SQL closes the Save As dialog box.

For more information, see [Saving and Closing Scripts](#).

#### SAVING MODIFIED DOCUMENTS

The table below describes the options and functionality on the Save Modified Documents dialog box:

Option	Description
Save All	Click to save documents in all open ISQL windows.
Save Selected	Click to save selected documents.
Invert	Click to clear selection.
Save None	Click to not save documents and close the application.

#### Important Notes

None

For more information, see [Saving and Closing Scripts](#).



## CLOSING A SCRIPT

To close a script, do the following:

- 1 On the Main menu, click Close.

OR

On the *ISQL window* tool bar, click Close.

OR

In the upper right corner of the window, double-click the *System* menu icon.

Rapid SQL starts closing the script.

- 2 If you have not saved your script, Rapid SQL prompts you to save the file. Click Yes to save and No to close without saving.

For more information, see [Saving and Closing Scripts](#).

## ABOUT THE ERROR PANE

You can view errors currently present in an ISQL editor window in a separate pane. The error pane provides a dynamic, line-by-line listing of the following error types:

- **Semantic** - errors in specified object names. Semantic problems persist in the Error pane until they are corrected. For more information, see [Semantic Validation](#).
- **Syntactic** - syntax problems associated with the SQL dialect of the data source to which the ISQL Editor session is connected. Syntax problems persist in the Error pane until they are corrected. For more information, see [Syntax Checking](#).
- **Execution** - errors detected the last time the script was executed in the current ISQL editor window. Execution errors persist in the Error pane until they are refreshed by a subsequent execution.

A single menu command toggles the error pane open and closed.

### To open or close the error pane:

- 1 On the Query menu, select *Show Errors*.

For more information, see [ISQL Editor](#).

## EDITING SCRIPTS

The ISQL Window incorporates a powerful editor that lets you edit your SQL scripts. The ISQL Window includes editing features such as:

- Uppercase or lowercase character conversion.
- Commenting and uncommenting blocks of text.
- Selecting text.

- Inserting PL/SQL tags.
- Setting and navigating with bookmarks.

The ISQL Editor provides [Paste SQL Syntax](#) and [Paste SQL Statements](#) utilities from which you can paste syntax for SQL commands and functions directly into an ISQL Window.

## Editing a Script

To edit a script, do the following:

- 1 In the SQL window, type your changes or additions.
- 2 When you are finished with your changes, on the ISQL window tool bar, click [Execute](#) to compile the script.

For more information, see [ISQL Editor](#).

## PASTE SQL SYNTAX

**NOTE:** This functionality is available for all platforms.

The Paste SQL Syntax facility lets you paste SQL syntax without having to refer to documentation to find syntax for SQL commands. You can also paste the SQL directly into an ISQL window. The Paste SQL Syntax facility that includes SQL syntax for:

- Commands
- Functions
- XML
- Other object syntax

## Important Notes

None

The table below describes the options and functionality on the SQL Syntax dialog box.

Option	Description
SQL Statements	Lets you select the target command, functions, or optimizer hint.
Syntax	Displays the syntax.
Paste Button	Click to paste the SQL statement into your ISQL Window.

**NOTE:** You must change the placeholders (e.g., expression) in the statements to reflect the specific objects in your database. For assistance with basic SQL statements, such as Select, Insert, Update, or Delete, use the Paste SQL Statement facility instead.

The table below describes options for each RDBMS platform:

Oracle	Sybase ASE	Microsoft
SQL Commands	SQL Commands	SQL Commands
Number Functions	Aggregate Functions	Aggregate Functions
Character Functions	Datatype Conversion Functions	Datatype Conversion Functions
Date Functions	Date Functions	Date Functions
Conversion Functions	Mathematical Functions	Mathematical Functions
Group Functions	String Functions	String Functions
Other Functions (User, NVL, etc.)	System Functions	System Functions
Optimizer Hints	System Diagnostics	Text/Image Functions

For more information, see [Completing the SQL Syntax Dialog Box](#).

#### COMPLETING THE SQL SYNTAX DIALOG BOX

To complete the SQL Syntax dialog box, do the following:

- 1 Place your insertion point in an open SQL window.
  - 2 On the Edit menu, click Paste SQL Syntax.
- OR
- On the *ISQL Window* tool bar, click Paste SQL Syntax.
- 3 In SQL Statements, select the target command, functions, or optimizer hint.
  - 4 To paste the SQL statement into your ISQL Window, click Paste.

For more information, see [Editing Scripts](#).

#### PASTE SQL STATEMENTS

**NOTE:** The functionality is available for all platforms.

The Paste SQL Statement facility lets you create Select, Insert, Update, and Delete SQL statements. The Paste SQL Statement window displays all available database object names from which you can choose to construct SQL statements. These statements are dependent on the object type you choose.

#### Important Notes

None

The table below describes the options and functionality on the Paste SQL dialog box.

Option	Description
Datasource	Lets you select the target datasource.
Database	Lets you select the target database.
Owner	Lets you select the owner.
Object Type	Lets you select the target object type.
Select	Click if you want a SELECT SQL statement.
Update	Click if you want an UPDATE SQL statement.
Insert	Click if you want an INSERT SQL statement.
Delete	Click if you want a DELETE SQL statement.
Object Type Box	<p>Rapid SQL displays a column of objects, given your selections of datasource, owner, and object type.</p> <p>Lets you select the check box next to the target object type, or click All to select all.</p> <p>Paste - Click to paste this object into your Editor window.</p>
Middle Box	<p>Rapid SQL displays attributes appropriate to the type of object you selected. For example, if you select Tables, Rapid SQL fills the second column with the columns of the table you select in the first column. If you select Procedures, Rapid SQL fills the second column with available parameters for the procedure you select in the first column.</p> <p>Lets you click the target object type properties or click All to select all. For example, if you selected a Table, then select one or more columns in the middle box; or, if you selected a Procedure, then select parameters in the middle box.</p> <p>Paste - Click to paste the object properties into your Editor window, under the Object Type box.</p>
Right Box	<p>Displays the SQL statement.</p> <p>Paste - Click when you are satisfied with the entire SQL statement. Rapid SQL pastes the SQL statement into your Editor window.</p>

For more information, see [Completing the Paste SQL Dialog Box](#).

#### COMPLETING THE PASTE SQL DIALOG BOX

To complete the Paste dialog box, do the following

- 1 Place your insertion point in an open *Editor* window.
- 2 On the Edit menu, click Paste SQL.

OR

On the *Editor* window tool bar, click Paste SQL.

- 3 In Datasource, select the target datasource.
- 4 In Database, select the target database.
- 5 In Owner, select the owner.
- 6 In Object Type, select the target object type.
- 7 Click Select if you want a SELECT SQL statement.
- 8 Click Update if you want an UPDATE SQL statement.
- 9 Click Insert if you want an INSERT SQL statement.
- 10 Click Delete if you want a DELETE SQL statement.
- 11 In the Object Type Box, select the check box next to the target object type, or click All to select all.
- 12 Click Paste to paste this object into your Editor window.
- 13 In the Middle Box, click the target object type properties or click All to select all.
- 14 Click Paste to paste the object properties into your Editor window, under the Object Type box.

Rapid SQL the SQL statement in the right box.

- 15 In the Right Box, click Paste when you are satisfied with the entire SQL statement.

Rapid SQL pastes the SQL statement into your Editor window.

For more information, see [Editing Scripts](#).

## SELECTING TEXT

The ISQL Editor lets you select a single word, a single line, or a block of text.

### Selecting a Single Word

To select a single word, do the following:

- 1 In the *ISQL Editor* window, position the pointer in the word and double-click.

### Selecting a Line of Text

Rapid SQL offers two ways to select a line of text:

- 1 In the *ISQL Editor* window, click the line number listed on the left side of the window.

OR

- 2 Position the pointer at the beginning of the line of text, hold down the SHIFT key, and then click the end of the line of text.

## Selecting a Block of Text

To select a block of text, do the following:

- 1 In the *ISQL Editor* window, drag until the block of text is selected.

For more information, see [Editing Scripts](#).

## MOVING AND COPYING TEXT IN AN ISQL WINDOW

You can move or copy information anywhere in an Editor window or into a different Editor window by dragging it.

### Moving Text

To move text, do the following:

- 1 In the *ISQL Editor* window, select the text you want to move.
- 2 Drag to the target location.

### Copying Text

To copy text, do the following:

- 1 In the *ISQL Editor* window, select the text you want to copy.
- 2 On the *Edit* menu, click *Copy*.
- 3 Drag to the target location.

**NOTE:** If you are dragging between different Editor windows, arrange the Editor windows so that the source and destination windows are open and visible. You must be able to see both the original and target locations.

For more information, see [Editing Scripts](#).

## COMMENTING AND UNCOMMENTING SCRIPTS

For most developers and development teams, documenting code is a fundamental part of the coding process. Besides contributing explanatory material for others who need to analyze or maintain your code later, the ability to comment and uncomment code can be very useful for bypassing statements in a procedure during compilation.

## Commenting Code

To commenting out code, do the following:

- 1 On the Edit menu, click Comment Out.

OR

Select a line of code, and on the Edit tool bar, click Comment Out.

Rapid SQL comments code.

## Uncommenting Code

To uncomment code, do the following:

- 1 On the Edit menu, click *Undo Comment Out*.

OR

Select a line of code, and on the Edit tool bar, click Undo Comment Out.

Rapid SQL uncomments code.

For more information, see [Editing Scripts](#).

## CHANGING CASE

When writing SQL scripts, you can change the letter casing of certain statements or lines of code. You can change case from lowercase to uppercase, or from uppercase to lowercase, using the case functions.

## Changing Case

To change case, do the following:

- 1 Select one or more letters in your script.
- 2 On the Edit menu, click Upper Case.

OR

On the *Edit* menu, click *Lower Case*.

OR

On the *Edit* tool bar, click Upper Case.

OR

On the Edit tool bar, click Lower Case.

Rapid SQL changes the case.

For more information, see [Editing Scripts](#).

**CUTTING, COPYING AND PASTING TEXT AND COLUMNS IN AN ISQL WINDOW**

The Editor window incorporates Cut, Copy and Paste text, and whole columns functions. You can move the selected text or columns to and from the Windows clipboard.

For more information, see:

[Copying and Pasting Text](#)

[Cutting and Pasting Text](#)

[Copying and Pasting a Column](#)

[Cutting and Pasting a Column](#)

**COPYING AND PASTING TEXT**

To copy and paste text, do the following:

- 1 In the *Editor* window, select the target text.

- 2 On the Edit menu, click Copy.

OR

On the *Editor* tool bar, click Copy.

Rapid SQL copies the text.

- 3 On the Edit menu, click Paste.

OR

On the *Editor* tool bar, click Paste.

Rapid SQL pastes the text.

For more information, see [Cutting, Copying and Pasting Text and Columns in an ISQL Window](#).

**CUTTING AND PASTING TEXT**

To cut and paste text, do the following:

- 1 In the *Editor* window, select the target text.

- 2 On the Edit menu, click Cut.

OR

On the *Editor* tool bar, click Cut.

Rapid SQL cuts the text.

- 3 On the Edit menu, click Paste.

OR

On the *Editor* tool bar, click Paste.

Rapid SQL pastes the text.



For more information, see [Cutting, Copying and Pasting Text and Columns in an ISQL Window](#).

#### COPYING AND PASTING A COLUMN

To copy and paste a column, do the following:

- 1 In the *Editor* window, position the pointer in front of the column of the target text.
- 2 Press *ALT* and simultaneously drag the pointer over the target column.
- 3 On the *Edit* menu, click *Copy*.

OR

On the *Editor* tool bar, click *Copy*.

OR

Right-click the text, and then click *Copy*.

Rapid SQL copies the column.

- 4 In the *Editor* window, position the pointer where you want to paste the column.
- 5 On the *Edit* menu, click *Paste*.

OR

On the *Editor* tool bar, click *Paste*.

OR

Right-click the text, and then click *Paste*.

Rapid SQL pastes the column.

For more information, see [Cutting, Copying and Pasting Text and Columns in an ISQL Window](#).

#### CUTTING AND PASTING A COLUMN

To cut and paste a column, do the following:

- 1 In the *Editor* window, position the pointer in front of the column of the target text.
- 2 Press *ALT* and simultaneously drag the pointer over the target.
- 3 On the *Edit* menu, click *Cut*.

OR

On the *Editor* tool bar, click *Cut*.

OR

Right-click the text, and then click *Cut*.

Rapid SQL cuts the column.

- 4 In the *Editor* window, position the pointer where you want to paste the column.

5 On the *Edit* menu, click Paste.

OR

On the *Editor* tool bar, click Paste.

OR

Right-click the text, and then click *Paste*.

Rapid SQL pastes the column.

For more information, see [Cutting, Copying and Pasting Text and Columns in an ISQL Window](#).

## SETTING BOOKMARKS

Bookmarks are useful tools for navigating throughout an Editor window. You can jump back and forth between bookmarks easily, and there is no practical limit to the number of bookmarks you can set.

The table below describes the options for setting bookmarks:

Feature	Description
Bookmarks:	Bookmarks are valuable navigation aids for jumping from one portion of a script to another. You can add bookmarks in important areas of your scripts, then jump back and forth between bookmarks.

## Setting a Bookmark

To set a bookmark, do the following:

- 1 In the Editor window, position the pointer in front of the line you want to bookmark.
- 2 On the Editor window tool bar, click Bookmark.

Rapid SQL inserts a blue dot in the gutter next to the line you have book marked.

For more information, see:

[Moving Between Bookmarks](#)

[Clearing Bookmarks](#)

## MOVING BETWEEN BOOKMARKS

You use the Next Bookmark and the Previous Bookmark buttons to move back and forth between bookmarks.

## Goto the Next Bookmark

To goto the next bookmark, do the following:

- 1 In the Editor window, position the pointer where you have set bookmarks, and then click Bookmark.

Rapid SQL jumps to the next bookmark.

## Goto the Previous Bookmark

To goto the previous bookmark, do the following:

- 1 In the Editor window, position the pointer where you have set bookmarks, and then click Goto Bookmark.

Rapid SQL jumps to the previous bookmark.

For more information, see [Setting Bookmarks](#).

## CLEARING BOOKMARKS

To clear bookmarks, do the following:

- 1 On the Editor tool bar, click Clear Bookmark.

Rapid SQL clears all bookmarks.

**NOTE:** This does not clear any error indicators (red dots) that might be showing in the gutter.

For more information, see [Setting Bookmarks](#).

## EXECUTING SCRIPTS

The ISQL Editor lets you execute all or part of your SQL scripts. Unless you are executing large scripts that have multiple statements within them, or you need to view optimizer statistics, you execute most of your SQL scripts without options.

**NOTE:** For Oracle Client 8.0.5, if you execute a multi-line SQL statement with an error in the syntax, the cursor does not jump to the line of the error.

The ISQL Editor incorporates multiple features to refine and ease script execution. The table below describes these features:

Feature	Description
<a href="#">Script Execution Facility</a>	The Script Execution Facility lets you execute scripts in parallel against multiple datasources. The facility also lets you schedule a job to perform the script execution at the appointed time, which saves development time and ensures accuracy and integrity across your databases.
<a href="#">Step Execution Facility</a>	The Step Execution Facility processes batches from one delimiter to the next.

Feature	Description
<a href="#">Query Plan</a>	The Query Plan Facility provides a graphical display that lets you identify the execution path that your SQL follows.  Rapid SQL's Query Plan window now displays data for the estimated costs, number of rows, and bytes returned by each plan step.
<a href="#">Query Options Dialog Box</a>	MICROSOFT SQL SERVER and SYBASE ASE ONLY:  The Query Options dialog box lets you customize what you see when you execute your query.

## Executing a Script

To execute a script, do the following:

- 1 On the Editor window tool bar, click Execute.
- OR
- In the Editor window, right-click, and then click Execute.
- Rapid SQL starts executing the script.

## Executing Partial Scripts

To execute a partial script, select the portion of the script to be executed and follow the steps outlined above. This feature is helpful while debugging large scripts and stored procedures.

### SCRIPT EXECUTION FACILITY

Rapid SQL has a Script Execution Facility that lets you run parallel queries against multiple datasources. This facility is also a stand-alone utility that is available from the utilities tool bar. If used as a stand-alone utility, you can directly type your script or copy and paste a script in an additional tab. Once you have selected the desired options, Rapid SQL establishes multiple threads and database connections to run the script simultaneously against the selected target datasources. Additionally, ANSI SQL scripts can run against multiple datasources from different DBMS vendors.

The Script Execution Facility runs with two output modes:

- Graphical
- File

Graphical output mode opens multiple result windows, one for each target datasource.

File output mode sends all output to a specified file or files. This feature allows for the execution of scripts against a large number of target datasources in a single operation and enables scheduling capabilities.

Once the scripts have finished executing, you have the option to send a notification message indicating that the script execution has completed via e-mail or Net Send. This message includes attachments of all the output files generated.

For more information, see [Executing Scripts Against Multiple Datasources](#).

#### EXECUTING SCRIPTS AGAINST MULTIPLE DATASOURCES

To execute scripts against multiple datasources, do the following:

- 1 On the *Utilities* menu, click Script Execution Facility.

OR

On the Utilities tool bar, click Script Execution Facility.

OR

On the Editor window tool bar, click Script Execution Facility.

The table below describes the options and functionality on the Script Execution Facility dialog box.

**NOTE:** This option is not available for a DDL Editor or PL/SQL Editor. To execute a script from a DDL or a PL/SQL Editor, use the stand-alone Script Execution Facility.

- 2 If you selected File output:

Option	Description
Script	Lets you type, copy and paste, or cut and paste a script.
Target	<p>Only Connected Datasources - Select to see only your currently connected datasources.</p> <p>All DBMS Types - Select to see all DBMS types.</p> <p>In the Datasource grid, select the check boxes next to the datasources against which you want to run your query, or click Select All to select all datasources.</p> <p>Database - Lets you type the name of the target database.</p>
Output	<p>Specify how you want the output to be handled</p> <p>Graphical Output - Select for graphical output.</p> <p>File Output - Select for file output.</p> <p>In the Directory box, type the full path and directory name in which you want to place the output file, or click Browse File icon to locate the directory.</p> <p>Click the File Type list to specify the file type you want to apply to the output file.</p> <p>Include column titles when saving - Select to include column titles in your saved file.</p> <p>Open files with registered applications - Select to open the files with registered application.</p>

Option	Description
Notify	<p>Job Description - Lets you type a job description to supply a subject in your e-mail message.</p> <p>E-mail addresses - Lets you type the e-mail address(es) separated by semi-colons.</p> <p>Net Send User Names - Lets you type the Net Send user name(s) separated by semi-colons.</p>
Schedule Button	This option is available when you select the File Output Option button and have the Microsoft Task Scheduler installed. For more information regarding scheduling, see <a href="#">Scheduling</a> .

For more information, see [Script Execution Facility](#).

#### COMPILATION ERROR TAB

Rapid SQL displays any errors encountered when a package, function or procedure was last compiled by a user in a separate tab of those object editors. The Errors Tab provides the line number, column position and error message for each compilation error.

For more information, see [Executing Scripts](#).

#### COMMAND HISTORY

The ISQL Editor is equipped with a command history box, which lets you keep a history of previously used SQL commands. The Command History is a list that is available in the [Results](#) window. Command history lets you go back and run previously used commands from the list rather than opening or inserting a script.

#### Executing from Command History

To execute from Command History, do the following:

- 1 In the Results window, click the Command History list, and then click the command you want to execute.

The Query window is automatically populated with the selected command.

- 2 To execute the command, click Execute.

For more information, see [Executing Scripts](#).

#### CANCELING A QUERY

The ISQL Editor lets you cancel a query while the rows are still being returned.

## Canceling a Query

To cancel a query, do the following:

- 1 On the Editor window tool bar, click Cancel.

**NOTE:** This button is only enabled after a script has begun executing.

For more information, see [Executing Scripts](#).

## STEP EXECUTION FACILITY

Step execution of scripts is an invaluable method to debug your scripts. The Step Execution facility lets you step through each batch in your script. While some batches can include many lines of code, some batches can consist of one line. Rapid SQL parses the statements and moves from batch to batch during the step execution process, each step controlled by you clicking the step execution button.

The table below describes how Rapid SQL provides a number of useful tools for working with step execution of your scripts:

Feature	Description	
ISQL Window Gutter	The ISQL Window Gutter is a vertical gray display bar located to the left of the ISQL window. It includes location indicators such as line numbers, error indicators, and bookmarks. The gutter is a quick visual cue to your current position in the script.	
Script Line Numbers	Line numbers are included to let you navigate large scripts. Error messages in the output window indicate the line number where the error occurred.	
Automatic Error Flagging	When using the Step Execution mode, Rapid SQL flags errors with a red dot in the ISQL window gutter. The errors are flagged sequentially as they are stepped into.	
Point and Click Error Navigation	Rapid SQL displays errors in the output window at the bottom of the screen and selects the errors as they occur. You can click each error and Rapid SQL takes you directly to the line where that error occurred.	
Step		Click the Step button to step into the next batch of code.
Step Back		Click the Step Back icon to step back to the most previous batch of code.
Step Over		Click the Step Over icon to jump over a batch to the next batch.
Run to Cursor		Click the Run to Cursor icon to execute all of the code between the beginning of the script to wherever you have inserted the pointer.
Cancel Step Execution		Click the Cancel Step Execution icon to change to regular execution mode.

For more information, see [Using the Step Execution Facility](#).

### USING THE STEP EXECUTION FACILITY

Rapid SQL offers three ways to use the step execution facility:

- 1 Open a script.
  - 2 On the *Query* menu, click *Step-Execute*.
- OR
- On the Editor window tool bar, click Step-Execute.
- OR
- Right-click, and then click, Step-Execute.

Rapid SQL starts the procedure and displays errors at the bottom of the Editor window.

**NOTE:** Rapid SQL indicates each executing line with a yellow arrow in the gutter of the Editor window. The gutter is that gray area between the line numbers on the left and the code window. As you step through your batches, Rapid SQL moves the arrow to indicate your current location.

- 3 To continue stepping through the script, on the Editor window tool bar, click Step, which displays in the Editor window after you have started the Step Execute procedure.
- 4 To step back, on the Editor window tool bar, click Step Back, which displays in the Editor window after you have started the Step Execute procedure.
- 5 To step over a batch and execute the next batch, on the Editor window tool bar, click Step Over, which displays in the Editor window after you have started the Step Execute procedure.
- 6 To stop *Step Execution* mode, on the Editor window tool bar, click Stop Step Execute, which displays in the Editor window after you have started the Step Execute procedure.
- 7 To *Run to Cursor*, on the Editor window tool bar, click Run to Cursor, which is available when the pointer is in the Editor window.

For more information, see [Executing Scripts](#).

### USING THE QUERY PLAN FACILITY

Each RDBMS platform lets you view the execution path that your SQL follows. For details, see the following topics:

- [Viewing a tree-based Query Plan \(IBM DB2 for LUW, SQL Server, Sybase\)](#)
- [Viewing a graphical Query Plan \(Oracle\)](#)

### VIEWING A TREE-BASED QUERY PLAN (IBM DB2 FOR LUW, SQL SERVER, SYBASE)

For IBM DB2 for Windows, Unix, and Linux, Microsoft SQL Server, and Sybase ASE, Rapid SQL lets you view a tree-based representation of a query plan. The Query Plan toolbar button is a toggle. Set it to enable the Show Plan mode.



Rapid SQL's Query Plan window displays data for the estimated costs, number of rows, and bytes returned by each plan step.

**NOTE:** For IBM DB2 for Linux, Unix, and Windows, Rapid SQL includes a tree view of statements and associated costs.

### To view a tree-based representation of a query plan

- 1 Open a script.
- 2 On the *Query* menu, click Query Plan.  
Rapid SQL starts the Show Plan mode.
- 3 To generate the Show Plan in a separate result window, click Execute.

For more information, see [Query Options](#).

### VIEWING A GRAPHICAL QUERY PLAN (ORACLE)

For Oracle, Rapid SQL lets you view a graphical representation of a query plan. You can toggle the view between the graph-based view and a tree-based view, view details for each step, and work with a number of viewing options.

The Query Plan toolbar button is a toggle. Set it to enable the Show Plan mode.

### To view a graphical representation of a query plan

- 1 Open a script.
- 2 On the *Query* menu, click Query Plan.  
Rapid SQL starts the Show Plan mode.
- 3 To generate the Show Plan in a separate result window, click Execute.

Options when working with the Query Plan facility include:

- Hovering the mouse over an execution step node to display detailed cost details for that step
- Right-clicking and choosing **Find Node** or **Find Next Node** to search large plans for nodes whose label contains a specified text string
- Right-clicking and choosing **Zoom In** or **Zoom Out**
- Right-clicking and choosing an **Orientation** sub-menu command to change the orientation of the view
- Right-clicking and choosing **Overview Window** to open a small window showing the entire plan
- Right-clicking and choosing **Save to File** to open a dialog that lets you save the graphical plan as a graphics file

- Clicking the Query Plan button to toggle between the graphical view and a tree-based view
- For multiple plans created from the same script/ISQL window, using the dropdown at the top of the plan to change the plan displayed

### QUERY OPTIONS

Rapid SQL offers a set of options corresponding to DBMS-specific query options. They let you customize the execution environment for an ISQL session with query-handling directives in areas such as performance, logging/reporting, and error handling. The following tables lists the query options available for each DBMS platform:

Minimally, query options are sent to the server on an ISQL session's first execution. Whether Rapid SQL needs to be sent to the server on subsequent executions depends on whether the session is locked or unlocked and whether you have changed the query option settings.

If you take no action, Rapid SQL will use a set of default query option settings to send to the server, one set per DBMS platform. You can, however save a set of query option settings for each DBMS, and have that set of settings used as the default for ISQL sessions against that platform. As well, during any ISQL session, you can modify the query option settings for the current session.

For more information on understanding and working with query options, see the following topics:

- [Setting query options](#)
- [Query option summary by DBMS](#)
- [Detailed query option descriptions](#)
- [When query options are sent to the server](#)

### Setting query options

The query options you select apply only to the current ISQL window instance. In addition, non-default query option settings you specify are not saved. You can however, save a set of query option settings to a file and subsequently open that file in other ISQL editor sections to have those settings applied to that session.

To specify query options:

- 1 With the ISQL Editor open, select Query Options from the **Query** menu.

Once open, you have the following options:

- Set each option manually using the associated check box. For detailed descriptions of query options, see [Query option summary by DBMS](#).
- Use the **Save** button to open a dialog that lets you save the current option query settings to an XML file. The dialog also has a **Load these query options as default** option that lets you use the settings in the saved file as the default query option settings for the current platform. Exercising this option sets the **Load Query Options** feature on the DBMS-specific **ISQL** tab of the Options editor to specify the file you saved. For more information, see [ISQL Options](#).
- Use the **Load** button to open a dialog that lets you locate and open a previously-saved XML file containing query option settings.
- Use the **Reset** button to restore the Rapid SQL query option defaults.

- 2 Click **OK** to set query options for the current ISQL window session.

**NOTE:** For information on conditions for which query options are sent to the server, see [When query options are sent to the server](#).

### Query option summary by DBMS

The following lists the specific query options offered for each platform:

IBM DB2 for Linux, Unix, and Windows	<a href="#">Batch Delimiter</a>	<a href="#">Check syntax when executing</a>	<a href="#">Create Explain plan tables if required</a>	<a href="#">Create explain plan tables on the SYSTOOLS schema</a>
	<a href="#">Isolation Level</a>	<a href="#">Max Errors Before Aborting</a>	<a href="#">Row Count</a>	<a href="#">Run Script with batch execution</a>
IBM DB2 for z/OS	<a href="#">Batch Delimiter</a>	<a href="#">Check syntax when executing</a>	<a href="#">Max Errors Before Aborting</a>	<a href="#">Row Count</a>
	<a href="#">Run Script with batch execution</a>			
Interbase/Firebird	<a href="#">Batch Delimiter</a>			
Microsoft SQL Server	<a href="#">Abort On Overflow</a>	<a href="#">Batch Delimiter</a>	<a href="#">Force Plan</a>	<a href="#">Ignore Overflow</a>
	<a href="#">Isolation Level</a>	<a href="#">Max Errors Before Aborting</a>	<a href="#">No Count</a>	<a href="#">No Exec</a>
	<a href="#">Parse Only</a>	<a href="#">Row Count</a>	<a href="#">Run Script with batch execution</a>	<a href="#">SET ... Options</a>
	<a href="#">Show Plan</a>	<a href="#">Statistics I/O</a>	<a href="#">Statistics Time</a>	<a href="#">Text Size</a>

MySQL	<a href="#">Batch Delimiter</a>	<a href="#">Big Tables</a>	<a href="#">Client Character</a>	<a href="#">Insert ID</a>
	<a href="#">Interactive Timeout</a>	<a href="#">Last Insert ID</a>	<a href="#">Max Errors Before Aborting</a>	<a href="#">Query Cache Type</a>
	<a href="#">Row Count</a>	<a href="#">SET Other Variables</a>	<a href="#">SQL Auto IS NULL</a>	<a href="#">SQL Big Selects</a>
	<a href="#">SQL Big Tables</a>	<a href="#">SQL Buffer Result</a>	<a href="#">SQL Log Bin</a>	<a href="#">SQL Log Off</a>
	<a href="#">SQL Log Update</a>	<a href="#">SQL Low Priority Updates</a>	<a href="#">SQL Max Join Size</a>	<a href="#">SQL Quote Show Create</a>
	<a href="#">SQL Safe Updates</a>	<a href="#">SQL Select</a>	<a href="#">SQL Warnings</a>	<a href="#">Transaction Isolation</a>
	<a href="#">Unique Checks</a>			
Oracle	<a href="#">Batch Delimiter</a>	<a href="#">Check syntax when executing</a>	<a href="#">Execution Information</a>	<a href="#">I/O Activity</a>
	<a href="#">Index Activity</a>	<a href="#">LONG Size Bytes</a>	<a href="#">Max Errors Before Aborting</a>	<a href="#">Network Activity</a>
	<a href="#">Parse Activity</a>	<a href="#">Row Count</a>	<a href="#">Run Script with batch execution</a>	<a href="#">Sort Activity</a>
	<a href="#">Table Activity</a>			
Sybase	<a href="#">Abort On Overflow</a>	<a href="#">Abort On Truncation</a>	<a href="#">ANSI NULL</a>	<a href="#">Batch Delimiter</a>
	<a href="#">Chained</a>	<a href="#">Force Plan</a>	<a href="#">Ignore Overflow</a>	<a href="#">Index Selection</a>
	<a href="#">Isolation Level</a>	<a href="#">Join Selection</a>	<a href="#">Last Insert ID</a>	<a href="#">LONG Size Bytes</a>
	<a href="#">Max Errors Before Aborting</a>	<a href="#">Network Activity</a>	<a href="#">No Count</a>	<a href="#">No Exec</a>
	<a href="#">Output to Execution Window</a>	<a href="#">Output to Server Error Log</a>	<a href="#">Parse Activity</a>	<a href="#">Parse Only</a>
	<a href="#">Prefetch</a>	<a href="#">Row Count</a>	<a href="#">Run Script with batch execution</a>	<a href="#">Set Quoted Identifier</a>
	<a href="#">Show Plan</a>	<a href="#">Statistics I/O</a>	<a href="#">Statistics Subduer Cache</a>	<a href="#">Statistics Time</a>

### Detailed query option descriptions

The following table, provides an alphabetical listing of options, option descriptions, and the platform to which the option applies:

Option/Variable	Description	Platform
<b>Abort On Overflow</b>	If set to TRUE, queries will be aborted on encountering a value greater than the specified size.	SQL Server, Sybase
<b>Abort On Truncation</b>	Specifies behavior following a loss of scale by an exact numeric datatype during an implicit datatype conversion. When set to TRUE, a statement that causes the error is aborted but other statements in the transaction or batch continue to be processed. When set to FALSE, Rapid SQL truncates the query results and continues processing.	Sybase

Option/Variable	Description	Platform
<b>ANSI NULL</b>	When set to TRUE, controls results of logical operations with NULL values.	Sybase
<b>Batch Delimiter</b>	The batch separator must be a viewable character and not a space, new line, or tab. The defaults have been selected to ensure compatibility with the features of Rapid SQL and the respective platform, but can be customized. Note: A custom delimiter works only from within an ISQL window and can't be used for extraction operations. Oracle: "/" (forward slash) DB2: ";" (semicolon) Sybase: "go" SQL Server: "go" MySQL: ";" (semicolon)	DB2 LUW, DB2 z/OS, InterBase/Firebird, MySQL, Oracle, SQL Server, Sybase
<b>Big Tables</b>	When set to TRUE, allows big result sets by saving all temporary sets to file. This can slow queries.	MySQL
<b>Chained</b>	Invokes a begin transaction before the following statements: delete, insert, open, fetch, select, and update. You must still explicitly close the transaction with a commit.	Sybase
<b>Check syntax when executing</b>	TRUE/FALSE	DB2 LUW, DB2 z/OS, Oracle
<b>Client Character</b>	Default character set.	MySQL
<b>Create Explain plan tables if required</b>	If set to TRUE, Explain Plan tables are created, as necessary. If set to FALSE and you don't manually create tables, Explain Plan operations will fail.	DB2 LUW
<b>Create explain plan tables on the SYSTOOLS schema</b>	If set to TRUE, Explain Plan tables are created on the SYSTOOLS schema. If the tables already exist in the user's default schema, Rapid SQL continues to use those tables. Refer to DB2 documentation for a listing of Explain Plan tables that must be deleted in order to use the SYSTOOLS option. If set to FALSE, Explain Plan tables are created under the user's default schema.	DB2 LUW
<b>Execution Information</b>	True/False	Oracle
<b>Force Plan</b>	When set to TRUE, processes a join in the same order as tables appear in the FROM clause of a SELECT statement only.	SQL Server, Sybase
<b>I/O Activity</b>	True/False	Oracle
<b>Ignore Overflow</b>	When set to TRUE, Rapid SQL will ignore any overflow caused by a resulting value that is larger than a column's specified size.	SQL Server, Sybase
<b>Index Activity</b>	True/False	Oracle
<b>Index Selection</b>	Valuable when tuning query performance.	Sybase
<b>Insert ID</b>	Choose a value to be used a following INSERT or ALTER TABLE statement when you supply an AUTO_INCREMENT value.	MySQL
<b>Interactive Timeout</b>	28800 default	MySQL
<b>Isolation Level</b>	Lets you set DB2 Isolation Levels of UNCOMMITTED READ, RESET, CURSOR STABILITY, REPEATABLE READ, or READ STABILITY.	DB2 LUW

Option/Variable	Description	Platform
<b>Isolation Level</b>	<p>Read Committed: Microsoft SQL Server default transaction isolation level. Specifies that shared locks are held while data is read to avoid dirty reads. You can change the data before the end of the transaction, resulting in nonrepeatable reads or phantom data.</p> <p>Read Uncommitted: The lowest level of transaction isolation. Transactions are isolated to ensure that physically corrupt data is not read. Applies dirty read, or isolation level 0 locking, which ensures that no shared locks are issued and no exclusive locks are honored.</p> <p>If set, it is possible to read uncommitted or dirty data values in the data can be changed and rows can appear or disappear in the data set before the end of the transaction.</p> <p>Repeatable Read: Places locks on all data used in a query, preventing other users from updating the data. Other users can insert new phantom rows into the data and are included in later reads in the current transaction.</p> <p>Concurrency is lower than Read Committed. Use this option only when necessary.</p> <p>Serializable: The highest level of transaction isolation. Transactions are completely isolated from one another.</p> <p>Places a range lock on the data set, preventing other users from updating or inserting rows into the data set until the transaction is complete.</p> <p>Concurrency is lower than Repeatable Read. Use this option only when necessary.</p>	SQL Server
<b>Isolation Level</b>	<p>0</p> <p>1: Sybase default isolation level.</p> <p>Prevents dirty reads.</p> <p>2: Prevents dirty and non-repeatable reads.</p> <p>3: Prevents dirty and non-repeatable reads and phantoms. This level is equivalent to performing all selects with holdlock.</p>	Sybase
<b>Join Selection</b>	Valuable when tuning query performance.	Sybase
<b>Last Insert ID</b>	Set the value to be stored in the binary log when you use LAST_INSERT_ID() in a statement that updates a table.	MySQL
<b>LONG Size Bytes</b>	8,192 is the default	Oracle
<b>Max Errors Before Aborting</b>	Select the maximum number of errors encountered before Rapid SQL aborts a script. Setting this value to zero disables the feature.	DB2 LUW, DB2 z/OS, MySQL, Oracle, SQL Server, Sybase
<b>Network Activity</b>	True/False	Oracle
<b>No Count</b>	Terminates the message indicating the number of rows affected by a Transact-SQL statement from being returned as part of the results.	SQL Server, Sybase
<b>No Exec</b>	When set to TRUE, compiles each query without executing it.	SQL Server, Sybase
<b>Output to Execution Window</b>	TRUE/FALSE	Sybase

Option/Variable	Description	Platform
<b>Output to Server Error Log</b>	TRUE/FALSE	Sybase
<b>Parse Activity</b>	TRUE/FALSE	Oracle
<b>Parse Only</b>	When set to TRUE, checks the syntax of each Transact-SQL statement and returns any error messages without compiling or executing the statement.  When TRUE, makes Microsoft SQL Server only parse the statement. When FALSE, makes Microsoft SQL Server compile and execute the statement.  Do not use Parse Only in a stored procedure or a trigger.	SQL Server, Sybase
<b>Prefetch</b>	When set to TRUE, enables large I/Os to the data cache. When set to FALSE, disables large I/Os to the data cache.	Sybase
<b>Query Cache Type</b>	The query is cached for ON or DEMAND.	MySQL
<b>Row Count</b>	When set to TRUE, Rapid SQL terminates the query after returning the specified number of rows.	DB2 LUW, DB2 z/OS, MySQL, Oracle, SQL Server, Sybase
<b>Run Script with batch execution</b>	TRUE/FALSE	DB2 LUW, DB2 z/OS, Oracle, SQL Server, Sybase
<b>SET ... Options</b>	The <b>Send Set Options</b> setting dictates whether the remaining ANSI SQL Options in this category are sent to the server. The default for this option is set using the <b>Enable SET query options</b> setting on the ISQL tab of the Options editor. For details, see <a href="#">ISQL Options</a> .  If the <b>Send Set Options setting</b> is enabled, the remaining settings in this category let you specify the specific ANSI SQL options that are sent to the server: <b>ansi_defaults</b> , <b>ansi_nulls</b> , <b>ansi_null_dflt_on</b> , <b>QUOTED IDENTIFIER</b> , <b>cursor_close_on_commit</b> , <b>ANSI_PADDING</b> , <b>ANSI WARNINGS</b> , and <b>IMPLICIT_TRANSACTIONS</b> . The initial default values are hard-coded, not obtained from server settings.	SQL Server
<b>SET Other Variables</b>	Variables to be set at runtime.	MySQL
<b>Set Quoted Identifier</b>	TRUE/FALSE	Sybase
<b>Show Plan</b>	When set to TRUE, reports data retrieval methods chosen by the Microsoft SQL Server query optimizer.	SQL Server, Sybase
<b>Sort Activity</b>	TRUE/FALSE	Oracle
<b>SQL Auto IS NULL</b>	When set to TRUE, enables you to find the last inserted row for a table.	MySQL
<b>SQL Big Selects</b>	When set to TRUE, SELECT statements likely to take a very long time to execute will be aborted (i.e., where the number of rows examined exceeds the max join size)	MySQL
<b>SQL Big Tables</b>	TRUE/FALSE	MySQL
<b>SQL Buffer Result</b>	When set to TRUE, forces results from SELECT statements into temporary tables.	MySQL

Option/Variable	Description	Platform
<b>SQL Log Bin</b>	When set to TRUE, allows logging to the binary log.	MySQL
<b>SQL Log Off</b>	When set to TRUE, no logging is done to the general query log.	MySQL
<b>SQL Log Update</b>	When set to TRUE, allows logging to the binary log.	MySQL
<b>SQL Low Priority Updates</b>	When set to TRUE, gives table modifying operations lower priority than SELECT operations	MySQL
<b>SQL Max Join Size</b>	When set to TRUE, you can catch SELECT statements where keys are not used properly and that would probably take a long time. Set it if your users tend to perform joins that lack a WHERE clause, that take a long time, or that return millions of rows.	MySQL
<b>SQL Quote Show Create</b>	When set to TRUE, table and column names will be quoted.	MySQL
<b>SQL Safe Updates</b>	When set to TRUE, the query aborts UPDATE or DELETE statements that do not use a key in the WHERE clause or a LIMIT clause. This makes it possible to catch UPDATE or DELETE statements where keys are not used properly and that would probably change or delete a large number of rows	MySQL
<b>SQL Select</b>	The maximum number of records that should be returned from SELECT statements.	MySQL
<b>SQL Warnings</b>	Defines whether or not single row insert statements generate an information string in the event of a warning.	MySQL
<b>Statistics I/O</b>	Lets you display information regarding the amount of disk activity generated by Transact-SQL statements.	SQL Server, Sybase
<b>Statistics Subduer Cache</b>	Displays the number of cache hits, misses, and the number of rows in the subquery cache for each subquery.	Sybase
<b>Statistics Time</b>	Displays the number of milliseconds required to parse, compile, and execute each statement.	SQL Server, Sybase
<b>Table Activity</b>	True/False	Oracle
<b>Table Count</b>	Sets the number of tables that Sybase ASE considers at one time while optimizing a join.	Sybase
<b>Text Size</b>	8,192 is the default	SQL Server, Sybase
<b>Transaction Isolation</b>	Repeatable Read is the default. Read Committed, Read Uncommitted, and Serializable are the other options. Refer to MS SQL Query Options dialog box for an explanation.	MySQL
<b>Unique Checks</b>	Performs uniqueness checks for secondary indexes of MyISAM tables.	MySQL



### When query options are sent to the server

Rapid SQL optimizes sending query options to the server for locked connections. Since an unlocked connection results in a new connection for each execution, query options must be sent with each execution. The following table outlines the specific cases:

<b>Session unlocked</b>	<p>When you open an ISQL window, if you not lock the connection, all query options are sent to the server each time you execute a script, immediately before executing the actual script. As long as the connection remains unlocked, all query options are sent to the server at each execution.</p> <p>Similarly if you unlock a currently locked session, all query options are sent to the server at each subsequent execution.</p> <p>For information on locking sessions, see <a href="#">Toolbar Options</a>.</p>
<b>Session locked</b>	<p>When you open an ISQL and lock the connection, all query options are sent to the server on the first execution. On subsequent executions, if you do not modify query option settings, no query options are sent to the server on subsequent executions.</p> <p>If you open <b>Query Options</b> dialog and modify options, on the next execution only the modified settings are sent to the server. For more information, see <a href="#">Setting query options</a>.</p> <p>These rules apply as long as the connection remains locked.</p>

### RESULT OPTIONS

The Result Options dialog box lets you set the SQL Results Window options. The table below describes Results options:

**NOTE:** The options you select only apply to the current window. To make options permanent, set the options in the Options Editor.

Interface Element	Option	Description	Default
Result Window	Single Window	Displays all results in one tabbed result window. Multiple result sets are appended together in the window. Single Window and Multiple Windows options are mutually exclusive.	Selected
	Multiple Windows	Displays multiple result sets one result set per window. Single Window and Multiple Windows options are mutually exclusive.	Not selected
Placement	Attached to Editor	Used in conjunction with Single Window option or Multiple Window option. Indicates that results appear as tabbed windows attached to the ISQL Window. Attached to Editor and Unattached options are mutually exclusively.	Selected

Interface Element	Option	Description	Default
	Unattached	Used in conjunction with Single Window option or Multiple Windows option. Indicates that results appear in windows separate from the ISQL Window. Attached to Editor and Unattached options are mutually exclusive.	Not Selected
Format	Standard Grid	Displays all result sets in a standard grid format. Result sets are only displayed in grid format in SQL Editors that are opened after you have selected this option. It does not apply to SQL Editors that are already open.	Selected
	HTML	Displays all result sets as HTML tables. Result sets are only displayed in HTML format in SQL Editors that are opened after you have selected this option. It does not apply to SQL Editors that are already open.	Not selected
	ASCII Text	Displays all result sets as ASCII Text. Result sets are only displayed in ASCII Text format in SQL Editors that are opened after you have selected this option. It does not apply to SQL Editors that are already open.	Not selected

#### USING THE TOKEN MATCHING CAPABILITY

When you are working with a large script with multiple levels of embedded steps, compare left and right parentheses, curly braces, square brackets and BEGIN/END pairs to make sure that you have delimited your code properly.

The Token Matching capability of Rapid SQL helps you achieve clean code.

#### Finding Matching Tokens

To find matching tokens, do the following:

- 1 Position the insertion pointer just to the left of the first token in a token pair you want to locate.
- 2 Click Match Token.

Rapid SQL jumps the pointer to the next available token.

For more information, see [Query Plan Facility](#).

## RESULTS EDITOR

The results of your SQL queries are displayed in the Results Tab of each Editor Window, which captures result sets in a virtual data grid that accommodates large result sets. The data grid offers many of the features of a basic Windows spreadsheet, giving you a great deal of flexibility in editing the worksheet and formatting its contents.

**TIP:** For Oracle, Rapid SQL displays REF CURSOR contents in the ISQL Window and Results Tab.

You have many choices for navigating and viewing your SQL query results. The Results window includes functionality to set result window options, find and replace, export data to other products such as Microsoft Excel, and mail your results files.

For more information, see [Configuring Result Set Windows](#).

## CONFIGURING RESULT SET WINDOWS

Result set windows can be configured in a variety of ways. You can configure your result set windows to present result sets in multiple or single panels, attached or detached from the corresponding ISQL window. These options can be set globally to save you the hassle of setting options for each result window. Additionally, Result windows can be torn off and dropped onto an open area of the workspace to create separate windows. These are known as Tear Off Tab Windows.

You can set the Result Window to display result sets in any of the following modes:

- [Single result sets in one attached tab window.](#)
- [Multiple result sets in one attached tab window.](#)
- [Single result sets in one separate unattached tab windows.](#)
- [Multiple result sets in one separate unattached tab windows.](#)
- Tear Off Tab windows.

## VIEWING MULTIPLE RESULT SETS IN SEPARATE UNATTACHED WINDOWS

To set the option to have multiple result sets displayed in multiple unattached windows, do the following:

- 1 On the *File* menu, click Options.  
OR  
On the Main tool bar, click Options.  
Rapid SQL opens the Options Editor.
- 2 In the Options Editor, click the list, and then click Results or click the Results Tab.
- 3 In the Results Window grid, click the Multiple Windows option.

- 4 In the Results Window grid, click the Unattached option.
- 5 Click OK.

Rapid SQL saves the settings and closes the Options Editor.

For more information, see [Configuring Result Set Windows](#).

#### VIEWING MULTIPLE RESULT SETS IN ONE ATTACHED TAB WINDOW

To set the option to have multiple result sets displayed in one attached tab window, do the following:

- 1 On the *File* menu, click Options.  
OR  
On the Main tool bar, click Options.  
Rapid SQL opens the Options Editor.
- 2 In the Options Editor, click the list, and then click Results or click the Results Tab.
- 3 In the Results Window grid, click the Multiple Windows option.
- 4 In the Results Window grid, click the Attached to Editor option.
- 5 Click OK.

Rapid SQL saves the settings and closes the Options Editor.

For more information, see [Configuring Result Set Windows](#).

#### VIEWING SINGLE RESULT SETS IN ONE SEPARATE UNATTACHED WINDOW

To set the option to have a single result set displayed in one unattached window, do the following:

- 1 On the *File* menu, click Options.  
OR  
On the Main tool bar, click Options.  
Rapid SQL opens the Options Editor.
- 2 In the Options Editor, click the list, and then click Results or click the Results Tab.
- 3 In the Results Window grid, click the Single Window option.
- 4 In the Results Window grid, click the Unattached option.
- 5 Click OK.

Rapid SQL saves the settings and closes the Options Editor.

For more information, see [Configuring Result Set Windows](#).

## VIEWING SINGLE RESULT SETS IN ONE ATTACHED TAB WINDOW

To set the option have all SQL query results displayed in a single attached tab window, do the following:

- 1 On the *File* menu, click Options.

OR

On the Main tool bar, click Options.

Rapid SQL opens the Options Editor.

- 2 In the Options Editor, click the list, and then click Results or click the Results Tab.
- 3 In the Results Window grid, click the Multiple Windows option.
- 4 In the Results Window grid, click the Attached to Editor option.
- 5 To re-use the same result set window for subsequent result sets, select the Reuse window check box. This option is only valid for single, attached result windows.
- 6 Click OK.

Rapid SQL saves the settings and closes the Options Editor.

For more information, see [Configuring Result Set Windows](#).

## EXPORTING DATA TO OTHER PRODUCTS

You can export data from a result set to traditional spreadsheet products, such as Microsoft Excel. You can copy part or all of a result set and paste it into your traditional spreadsheet product by way of the Microsoft Windows Clipboard function. You can also save your result sets as tab delimited files, comma separated files or as HTML tables. HTML tables can be opened in your default Internet browser. Tab delimited files and CSV files can be opened in any spreadsheet or word processing program.

**NOTE:** Rapid SQL supports pipe delimited ( '|' ) files when you save result sets.

For more information, see [Results Editor](#).

## SETTING RESULT WINDOWS TO READ ONLY MODE

To set your result windows to read only mode to keep anyone from accidentally editing or altering a result set, do the following:

- 1 Select a *Result* window that you want to make read only.
- 2 On the Edit menu, click *Read Only*.

Rapid SQL places a check-mark next to read only and sets the command.

**NOTE:** The Read Only command is a toggle. When it is set, the formatting buttons on the Edit menu are not available.

For more information, see [Results Editor](#).

## RESULT WINDOW STATUS BAR

The Result Window Status Bar displays information about each Result window at the bottom of each window. You can display or hide the status bar by setting the Status Bar preference. This preference applies to all result windows.

For more information, see [Results Editor](#).

## MAILING RESULT SETS

If you have MAPI-compliant electronic mail software installed on your computer, then you can mail result sets to other users.

### Mailing a Result Set

To mail a result set, do the following:

- 1 Open the Message dialog box:
- 2 On the *File* menu, click Send.

OR

On the File tool bar, click Send.

Rapid SQL opens the open Message dialog box.

- 3 In the *Address* box, type the name of the addressee(s) and any other options.

The ISQL Editor automatically attaches a file containing your result set to the mail message.

- 4 Click *Send*.

Rapid SQL sends the result set to the specified addressee(s).

For more information, see [Results Editor](#).

## CLOSING AND RENAMING RESULT WINDOW TABS

Rapid SQL lets you close or rename tabbed Result set windows.

### Renaming a Result Window Tab

To rename a Result Window Tab, do the following:

- 1 Place your pointer over the Results Tab.
- 2 Right-click the Results Tab, and then click *Rename*.

Rapid SQL opens the Rename Tab Editor.

- 3 In the *New Name* box, type the new name.
- 4 Click OK.

Rapid SQL changes the name and closes the Rename Tab Editor.

### Closing a Result Window Tab

To close a Result Window Tab, do the following:

- 1 On the *Result Window Tab* tool bar, click Close.

OR

Right-click the Results Tab, and then click Close.

Rapid SQL closes the Result Window Tab.

For more information, see [Results Editor](#).

## SAVING AND CLOSING RESULT SETS

You can save your result sets using the standard Save and Save As functions. You can save multiple copies of the same result set and modify each copy to specific formatting requirements.

### Saving Results

To save results, do the following:

- 1 On the *File* menu, click Save.

OR

On the Main tool bar, click Save.

Rapid SQL opens the Save Results dialog box.

- 2 In the File name box, type the name of the result set.

- 3 In Save as type, select the file type.

**TIP:** You can now save data in Excel 2000 or later .xls, user-specified delimited .txt, and XML formats.

- 4 To include column titles, select Include column titles when saving.
- 5 For delimited results, in User Specified Delimiter type the delimiter.
- 6 Click Save.

Rapid SQL saves the file and closes the Save As dialog box.

### Closing a Result Set

To close a result set, do the following:

- 1 On the *Main* menu, click Close.

OR

On the *Result* tool bar, click Close.

OR

In the upper-right corner of the window, double-click the System menu icon.

Rapid SQL closes the Result Set.

- 2 If you have not saved your result set, Rapid SQL prompts you to save the file. Click Yes to save and No to close without saving.

For more information, see [Results Editor](#).

## EDITING RESULT SETS

The Results Editor provides many ways to edit and customize your result windows. The Data Grid offers a host of features for manipulating, sorting and formatting data.

### Topics

- [Cutting, Copying, and Pasting Cell Contents](#)
- [Cutting, Copying, and Pasting Rows](#)
- [Cutting, Copying, and Pasting Columns](#)
- [Adding and Inserting Rows](#)
- [Adding and Inserting Columns](#)
- [Deleting Rows and Column](#)
- [Resizing Rows and Columns](#)



- [Sorting Data](#)

## CUTTING, COPYING, AND PASTING CELL CONTENTS

The Result window supports standard cut, copy and paste functionality.

### Cutting Cell Contents

To cut cell contents, do the following:

- 1 In the Results window, double click or tab to the target cell. A double line bounds the selected cell. You can also select text using standard text selection techniques.
- 2 On the *Edit* menu, click Cut.

OR

On the *Result* tool bar, click Cut.

Rapid SQL cuts the cell.

### Copying Cell Contents

To copy cell contents, do the following:

- 1 In the Results window, double click or tab to the target cell. A double line bounds the selected cell. You can also select text using standard text selection techniques.
- 2 On the *Edit* menu, click Copy.

OR

On the *Result* tool bar, click Copy.

Rapid SQL copies the cell.

### Pasting Cell Contents

To paste cell contents, do the following:

- 1 In the *Results* window, double click or tab to the target cell. A double line bounds the selected cell.
- 2 On the *Edit* menu, click Paste.

OR

On the *Result* tool bar, click Paste.

Rapid SQL pastes the cell.

For more information, see [Editing Result Sets](#).

## CUTTING, COPYING, AND PASTING ROWS

You can perform standard cut, copy, and paste functions on rows, just as you can on individual cells.

### Cutting Rows

To cut a row, do the following:

- 1 In the Results window, click the numbered row heading on the left side of the row.
- 2 On the *Edit* menu, click Cut.

OR

On the *Result* tool bar, click Cut.

Rapid SQL cuts the row.

### Copying Rows

To copy rows, do the following:

- 1 In the Results window, double click or tab to the target row. A double line bounds the selected row. You can also select text using standard text selection techniques.
- 2 On the *Edit* menu, click Copy.

OR

On the *Result* tool bar, click Copy.

Rapid SQL copies the row.

### Pasting Rows

To paste rows, do the following:

- 1 In the Results window, double click or tab to the target row. A double line bounds the selected row.
- 2 On the *Edit* menu, click Paste.

OR

On the *Result* tool bar, click Paste.

Rapid SQL pastes the row.

For more information, see [Editing Result Sets](#).

## CUTTING, COPYING, AND PASTING COLUMNS

You can perform standard cut, copy, and paste functions on columns, just as you can on rows.

### Cutting Columns

To cut columns, do the following:

- 1 In the Results window, click the column heading above the first row.
- 2 On the *Edit* menu, click Cut.

OR

On the *Result* tool bar, click Cut.

Rapid SQL cuts the column.

### Copying Columns

To copy columns, do the following:

- 1 In the Results window, click the column heading.
- 2 On the *Edit* menu, click Copy.

OR

On the *Result* tool bar, click Copy.

Rapid SQL copies the column.

### Pasting Columns

To paste columns, do the following:

- 1 In the Results window, click the column heading above the first row to select the target column.
- 2 On the *Edit* menu, click Paste.

OR

On the *Result* tool bar, click Paste.

Rapid SQL pastes the column.

For more information, see [Editing Result Sets](#).

## ADDING AND INSERTING ROWS

You can add or insert rows to expand or rearrange your result sets.

### Adding a Row

To add a row, do the following:

- 1 To add a row as the last row of the result set, position the pointer inside the result set.
- 2 On the Edit menu, click *Add Row*.

- 3 To add a row inside the result set, click the numbered row heading where you want to add a row.
- 4 On the Edit menu, click *Add Row*.

### Inserting a Row

To insert a row, do the following:

- 1 To insert a row as the last row of the result set, position the pointer inside the result set.
- 2 On the Edit menu, click Insert Row.
- 3 To insert a row inside the result set, click the numbered row heading where you want to insert a row.
- 4 On the Edit menu, click Insert Row.

For more information, see [Editing Result Sets](#).

### ADDING AND INSERTING COLUMNS

You can add or insert columns to expand or rearrange your result sets.

#### Adding a Column

To add a column, do the following:

- 1 Position the pointer inside the result set.
- 2 Select Edit, Add Column from the main menu. The new column is added as the last column of the result set.

#### Inserting a Column

To insert a column, do the following:

- 1 Select the column where you want to insert a column.
- 2 Select Edit, Insert Column from the main menu. The new column is inserted to the left of the column that you selected.

For more information, see [Editing Result Sets](#).

### DELETING ROWS AND COLUMNS

You can delete entire rows and columns to edit your result sets.

**Deleting a Row**

To delete a row, do the following:

- 1 Select the target row of data to delete.
- 2 On the *Edit* menu, click Delete Row.

OR

Right-click the row, and then click Delete Row.

Rapid SQL deletes the row.

**Deleting a Column**

To delete a column, do the following:

- 1 Select the target column of data to delete.
- 2 On the *Edit* menu, click Delete Column.

OR

Right-click the column, and then click Delete Column.

Rapid SQL deletes the column.

For more information, see [Editing Result Sets](#).

**RESIZING ROWS AND COLUMNS**

Resizing rows and columns can aid readability of the spreadsheet contents or condense space for editing and formatting purposes.

**Resizing Rows to Their Default Height**

To resize rows to their default height, do the following:

- 1 Select one or more rows by clicking on the numbered row headings to the left of the rows.
- 2 Right-click the selected rows, and then click *Resize Rows*.

**Resizing Rows to a New Height**

To resize rows to a new height, do the following:

- 1 Select one or more rows by clicking the numbered row headings to the left of the rows.
- 2 Change the pointer to a solid horizontal bar with arrows on top and bottom by moving it to one of the upper or lower borders of the row heading.
- 3 Click and grab the row border and drag the pointer to enlarge or shrink the height of the row.

### Resizing Columns to Their Default Widths

To resize columns to their default widths, do the following:

- 1 Select one or more columns by clicking the column headings.
- 2 Right-click the selected columns, and then click *Resize Columns*.

### Resizing Columns to a New Width

To resize columns to a new width, do the following:

- 1 Select one or more columns by clicking the column headings.
- 2 Change the pointer to a solid horizontal bar with arrows on top and bottom by moving it to one of the upper or lower borders of the column heading.
- 3 Click and grab the column border and drag the pointer to enlarge or shrink the height of the column.

For more information, see [Editing Result Sets](#).

## SORTING DATA

To order and organize data in a coherent manner, you can sort columns alphanumerically in single result sets.

### Sorting Data

To sort data, do the following:

- 1 In the data grid, select the column(s) you want to sort.
- 2 Double click the column header(s) to sort the data in the column in ascending, alphanumeric order. To sort the column in descending order, double click the column header again.

**NOTE:** This option is not valid for multiple result sets.

For more information, see [Editing Result Sets](#).

## FORMATTING RESULT SETS

The ISQL Editor provides flexibility in formatting result sets, for analysis and reporting, from within a Result Window. Result sets can be formatted to best suit your purpose, whether it be sending via e-mail, printing, or exporting a file to other programs, such as Microsoft Excel. Some of these features change only the screen display of the results, while others allow you to format for printing.

### Topics

- [Setting Alignment Properties](#)

- [Setting Border Properties](#)
- [Setting Result Set Display Properties](#)
- [Setting Fonts](#)
- [Setting Colors](#)

## SETTING ALIGNMENT PROPERTIES IN RESULT WINDOWS

You can use the format menu or the shortcut menu to alter result set alignment properties. Selecting the Alignment command displays the Format Styles tabbed dialog box:

### Setting Alignment Properties

To set alignment properties, do the following:

- 1 On the *Format* menu, click Alignment.  
OR  
Right-click the Result data grid, and then click Alignment.  
Rapid SQL opens the Format Styles dialog box.
- 2 Click the Align Tab.
- 3 To change *Horizontal* properties, click the appropriate *option* button in the *Horizontal* box.
- 4 To enable *Wrap Text*, *Allow Enter* and/or *Auto Size*, select the appropriate check boxes.
- 5 To disable *Wrap Text*, *Allow Enter* and/or *Auto Size*, select the appropriate check boxes and deselect these options.

The table below describes alignment properties:

Property	Option	Description
Horizontal	Standard	Aligns data in cells based on their default datatype alignment properties. For example, numbers are right aligned, whereas text is left aligned.
	Left Aligned	Forces all data to be left aligned, regardless of datatype
	Center	Forces centering of all data, regardless of datatype
	Right Aligned	Forces all data to be right aligned, regardless of datatype
Vertical	Standard	Aligns data in cells based on their default datatype alignment properties. For example, numbers are right aligned, whereas text is left aligned.
	Top	Forces all data to be aligned to the top of each cell
	Center	Forces all data to be aligned in the center of each cell
	Bottom	Forces all data to be aligned at the bottom of each cell

Property	Option	Description
Wrap Text		Specifies that data exceeding the length of the cell should wrap to additional lines in the same cell
Allow Enter		Specifies that the contents of cells can be edited or not
Auto Size		Specifies whether or not rows should size automatically to accommodate the longest data item in a column

For more information, see [Formatting Result Sets](#).

#### FORMAT BORDER STYLES

You can use the format menu or the shortcut menu to alter border properties. Selecting the Border command displays the Format Styles tabbed dialog box.

#### Completing the Format Styles Dialog Box

To complete the Format Styles dialog box, do the following:

- 1 On the *Format* menu, click Borders.  
OR  
Right-click the Result data grid, and then click Borders.  
Rapid SQL opens the Format Styles dialog box.
- 2 Click the Borders Tab.
- 3 On the *Border* box, you can indicate whether or not a border should appear on the top, bottom, right, left, or for a range of cells by clicking the corresponding boxes.
- 4 To set a range of cells apart by setting a particular border style around that range, select the range in the result set before opening the *Format Styles* dialog box. To select the *Range* property, click the range box.
- 5 In the *Type* box, you can select the type of line you want to border the cell or cells by clicking the corresponding boxes.
- 6 To select a color, click the *Color* list, and then click the border color.
- 7 Click OK.

Rapid SQL saves your changes and closes the Format Styles dialog box.

For more information, see [Formatting Result Sets](#).

#### DISPLAY SETTINGS

You can use the format menu or the shortcut menu to alter display properties. Selecting the Properties command displays the Display Setting dialog box.



The Display Settings dialog box lets you set a number of format properties for the result set window:

- How titles and grid lines are displayed.
- How lines in the grid are displayed.
- How the result set window background is displayed.
- How current cells, rows and columns are indicated.

### Completing the Display Settings Dialog Box

To complete the Display Settings dialog box, do the following:

- 1 On the *Format* menu, click *Properties*.

OR

Right-click the Result data grid, and then click *Properties*.

Rapid SQL opens the Display Settings dialog box.

- 2 In the *Titles and Gridlines* box, select the target display properties check boxes. The change is displayed in the right-pane *Preview* window of the *Properties* dialog box.
- 3 In the *Color* box, click the property you want to change, and then click the new color. Changes are propagated automatically when you change them.
- 4 To alter the user properties, in the *User Properties* list, click the attribute you want applied to the current cell.
- 5 To save your changes to the current configuration, select the *Save settings to profile* check box.
- 6 Click *OK*.

Rapid SQL saves your changes and closes the *Properties* dialog box.

For more information, see [Formatting Result Sets](#).

### FORMAT FONT STYLES

You can use the format menu or the shortcut menu to alter font properties. Selecting the *Font* command displays the *Format Styles* tabbed dialog box.

## Completing the Format Styles Dialog Box

To complete the Format Styles dialog box, do the following:

- 1 On the *Format* menu, click *Font*.

OR

Right-click the Result data grid, and then click *Font*.

*Rapid SQL* opens the Format Styles dialog box.

- 2 Click the *Font Tab*.
- 3 In the *Font Combo* box, type or click the font you want to use.
- 4 In the *Outline Combo* box, type or click the outline you want to use.
- 5 In the *Size Combo* box, type or click the size you want to use.
- 6 To make a line cross through the length of the text, in the *Effects* box, select the *Strikeout* check box.
- 7 To underline the text, in the *Effects* box, select the *Underline* check box.
- 8 To change the text color, click the *Text Color* list, and then click the new color.

The Preview box displays the sample text of your selections

For more information, see [Formatting Result Sets](#).

## FORMAT COLOR STYLES

You can use the format menu or the shortcut menu to alter color properties. Selecting the *Color* command displays the Format Styles tabbed dialog box.

The table below describes the options and functionality on the Format Styles dialog box:

Option	Description
Interior	<p>To change the result set pattern, in the box, click the list, and then click a new pattern.</p> <p>To set the foreground color of the result set, in the box, click a foreground group color button.</p> <p>To set the background color of the result set, in the box, click a background group color button.</p>

Option	Description
3D-Effect	<p>To change the 3D-Effect, in the box, click the appropriate option button:</p> <p>The Normal option button provides a standard flat appearance.</p> <p>The Raised option button provides a raised 3D appearance in the selected cells.</p> <p>The Inset option button provides a depressed appearance in the selected cells. The default is typically set to Normal.</p>

### Completing the Format Styles Dialog Box

To complete the Format Styles dialog box, do the following:

- 1 On the *Format* menu, click *Color*.

OR

Right-click the Result data grid, and then click *Color*.

Rapid SQL opens the Format Styles dialog box.

For more information, see [Formatting Result Sets](#).

### NOTES ON XML TYPES AND UNICODE DISPLAY IN THE RESULTS EDITOR

When viewing data in the Results grid, keep the following in mind:

- XML data types are supported for IBM DB2 for Windows, Unix, and Linux, Microsoft SQL Server, and Oracle. In the Results grid, XML data types are displayed as LOB content.
- Support for display of Unicode characters is provided as follows:
  - IBM DB2 for Windows, Unix, and Linux V8 and V9: **character**, **clob**, **varchar**, and **longvarchar** types
  - SQL Server 2000: for **nchar**, **nvarchar**, and **ntext** types
  - SQL Server 2005: **nchar**, **nvarchar**, **ntext**, and **nvarchar(max)** types
  - Oracle 8i, 9i, and 10g: **NCHAR**, **NVARCHAR2** and **NCLOB** for non-Unicode UTF8 Character Set Instances and **NCHAR**, **NVARCHAR2**, **CHAR**, **VARCHAR2**, **LONG**, **NCLOB** and **CLOB** for Unicode UTF8 Character Set Instances
  - Sybase 12.5 and 15.2: **UNICHAR**, **UNIVARCHAR** and **UNITEXT** for non-Unicode UTF8 Character Set Instances and **UNICHAR**, **UNIVARCHAR**, **UNITEXT**, **NCHAR**, **NVARCHAR**, **CHAR**, **VARCHAR** and **TEXT** for Unicode UTF8 Character Set Instances



# PERMISSIONS MANAGEMENT

Rapid SQL includes permissions management capabilities that include:

- [Explicit Permissions](#)
- [Cascading Permissions](#)
- [Using Roles to Grant Permissions and Privileges](#)
- [Using Roles to Revoke Permissions and Privileges](#)
- [Grant Privilege\(s\) To](#)
- [Revoke Privilege\(s\) From](#)
- Deny Privileges From

## EXPLICIT PERMISSIONS

Explicit permissions are permissions specifically granted by one user to another. Granting a permission is an integral part of database security and management. Rapid SQL tracks explicit permission changes across your enterprise. Multiple explicit permissions can be consolidated in roles. Different groups and roles can share permissions for an object. Rapid SQL lets you grant permissions with the Roles Editor. Embarcadero lets you edit permissions on the Privileges tabs of the object editors.

**TIP:** The Admin option is similar to the grantable option for privileges. If the user has the admin option, they can grant that role to other people.

**TIP:** In most cases, you will want to make all roles granted, default roles. If you do not make a role default for a user, they will have to enable it with a SQL command. To avoid this complication, assign roles as default unless business rules specify otherwise.

**TIP:** The Grantable option gives the user the ability to grant that privilege to other users. Use the Grantable option SPARINGLY.

For more information, see [Permissions Management](#).

## CASCADING PERMISSIONS

Cascading permissions are the path of privileges granted from one user to another user, group, or role. Using cascading permissions is a way to manage multiple sets of permissions and privileges for groups of users. When you drop a user with a revoke cascade command, all permissions and privileges granted by the dropped user are also revoked. Embarcadero lets you edit and set grant permission parameters with the Permissions Editor.

Once a user with grant permission privileges is dropped with cascade, reinstate permissions for all affected users.

For more information, see [Permissions Management](#).

## USING ROLES TO GRANT PERMISSIONS AND PRIVILEGES

Roles are sets of user privileges you associate with access to objects within a database. Roles streamline the process of granting permissions. You can use roles to grant sets of permissions and privileges to users and groups. Rapid SQL lets you grant permissions to roles with the Roles Editor.

For more information, see [Permissions Management](#).

## USING ROLES TO REVOKE PERMISSIONS AND PRIVILEGES

Roles can be effective in revoking permissions and privileges from users. Rather than individually revoke permissions from users, you can revoke groups of permissions from multiple users simultaneously using roles. Rapid SQL lets you revoke permissions with the Roles Editor.

Rapid SQL lets you identify existing users and their database permissions, and to detect and manage changes to user permissions by rolling change back or forward. Rapid SQL lets you manage database permissions in a cross-platform database environment, and gives you the ability to compare objects/permissions and migrate them to other instances. Using the compare functionality for permissions management, Rapid SQL gives you:

- Enhanced permissions management across the enterprise (Oracle, Sybase ASE, Microsoft SQL Server, and IBM DB2 LUW for Linux, Unix, and Windows).
- Faster detection of changed user rights.
- Ability to detect database accounts (users/logins) that are invalid.
- Rollback of invalid permissions in a single click.
- Archives of users, groups, roles and their permissions.

For more information, see [Permissions Management](#).

## GRANT PRIVILEGE(S) TO

The Grant Privilege(s) To dialog box lets you select or clear the check boxes corresponding to the target privileges.

Permissions and privileges allow or limit access to system database administration and database objects. To manage databases, it is important to monitor who can access the enterprise structure and change the underlying schema. It is also important to monitor who can grant and revoke permissions and privileges in the enterprise. Rapid SQL lets you manage permissions and privileges in the Grant Privilege(s) To dialog box.

On the Privileges Tab of the editor, you can see whether a privilege was inherited from a role or group, or if it was granted explicitly by a user. Rapid SQL also shows if the privilege is grantable (granted with the GRANT OPTION.) The table below describes the icons:

**NOTE:** The available privileges depend on the target DBMS and object.

The table below describes the options and functionality on the Grant Privilege(s) From Dialog Box:

Option	Description
Privilege	Lets you select or clear the check boxes corresponding to the target privileges.
Grantable	Select No to prevent cascading the privileges to other users.

The table below describes the object permissions:

Object	Permission(s)
Index	CONTROL
Packages	BIND, CONTROL, EXECUTE
Schema	ALTERIN, CREATIN, DROPIN (w/GRANT OPTION)
Tables	ALTER, CONTROL, DELETE, INDEX, INSERT, REFERENCES (& on column), SELECT, UPDATE (& on column) (w/GRANT OPTION)
Tablespaces	USE (w/GRANT OPTION)
Views	CONTROL, DELETE, INSERT, SELECT, UPDATE (& on column) (w/GRANT OPTION)

For more information, see:

[Permissions Management](#)

## REVOKE PRIVILEGE(S) FROM

The table below describes the options and functionality on the Revoke Privilege(s) From dialog box:

Option	Description
Privilege	Lets you select or clear the check boxes corresponding to the target privileges.
Cascade	Select No to prevent cascading the revocation privileges to other users.

For more information, see:

[Permissions Management](#)

## DENY PRIVILEGES FROM

**NOTE:** Deny Privileges is available for Microsoft SQL Server only.

The Deny Privileges From dialog box lets you deny a privilege from a security account in the current database.

The table below describes the options and functionality on the Deny Privilege(s) From dialog box:

Option	Description
Privilege	Lets you select or clear the check boxes corresponding to the target privileges.
Cascade	Select No to prevent cascading the deny privileges to other users.

**TIP:** On the editor, the Deny privilege can be Revoked, just as a Grant permission can be revoked.

For more information, see:

[Permissions Management](#)



## PROJECT MANAGEMENT

Rapid SQL database project management facilities help you organize, alter, and keep track of changes to database objects or SQL scripts. The project management facilities act as a repository to maintain all source code for a database project. Rapid SQL also incorporates version control functions and build management facilities to help you manage and build projects. Once a project has been created, Rapid SQL lets you:

- Review a file's history.
- Return to earlier versions of a file.
- Develop concurrently.

A project generally includes SQL script files that you can maintain and create in unison with your database administration and development cycle. Projects can also contain subfolders (for easy categorization) as well as other projects (subprojects.) You can create projects manually by inserting existing SQL script files. You can also create projects automatically by reverse-engineering a database schema or an existing version control project.

When administering or developing SQL database schemas, SQL source code files are the foundation of effective change management for database objects. The DDL commands used to create database objects on SQL database servers are often in a constant state of flux as new columns and constraints are added to table schemas or stored procedure logic is changed. The ability to track and store these changes to files directly from Rapid SQL alleviates any disruption in the development environment.

Lengthy script files containing the DDL to compile stored procedures and triggers on a database server such as Oracle or Sybase Adaptive Server go through constant revisions similar to C program files and word processing documents. Rapid SQL's Project Management facilities let you effectively monitor revisions for your database servers.

For more information, see:

[Create a New Project](#)

[Working With Projects](#)

## CREATE A NEW PROJECT

Creating a project can be a very time-consuming task. Rapid SQL eases this task by providing step-by-step wizards to guide you through the process. You have a variety of options to choose from when you create a new project. Rapid SQL lets you create new projects:

- [From a database.](#)
- [From existing files.](#)
- [From a version control project.](#)
- [Without initialization.](#)

The [New Project dialog box](#) lets you select how you want to create a new project. Using the New Project dialog box, you can name, specify a file location, and provide a description of the project.

## NEW PROJECT DIALOG BOX

The New Project dialog box lets you select how you want to create a new project. Using the New Project dialog box, you can name, specify a file location, and provide a description of the project.

### Important Notes

None

The table below describes the options and functionality of the New Project dialog box:

Option	Description
Name	Lets you enter a name for the project.
Location	Lets you type or browse and locate a directory for the project. NOTE: This is set as the working directory for your project.
Description	OPTIONAL: Lets you enter a description of the project.
Initialize New Project	Lets you select how you want to create a new project. <a href="#">From Database</a> - Lets you reverse engineer an existing database. <a href="#">From Existing Files</a> - Lets you create a project from scratch, or use an existing project. <a href="#">From Version Control</a> - Lets you reverse-engineer a project from existing version control system projects. This is helpful for users who have already created database projects in their version control systems. <a href="#">Do Not Initialize</a> - Lets you create an empty project where you can later add files and/or database objects.

For more information, see [Completing the New Project Dialog Box](#).

### COMPLETING THE NEW PROJECT DIALOG BOX

To complete the New Project Dialog Box, do the following:

- 1 On the File menu, select New, and then Project.  
Rapid SQL opens the New Project Dialog Box.
- 2 In Name, enter a name for the project.
- 3 In Location, type or browse and locate a directory for the project.
- 4 OPTIONAL: In Description, lets you enter a description of the project.

- 5 In Initialize New Project, select how you want to create a new project.
- 6 Click OK.

## CREATE A NEW PROJECT FROM A DATABASE

Rapid SQL offers a context-sensitive wizard that reverse-engineers all or part of an existing database. When extracting database schema, Rapid SQL captures the complete definition of the following objects:

- Tables
- Views
- Stored Procedures
- Functions

Rapid SQL then creates a project containing files based on the relevant object types. Rapid SQL creates a separate file for each database object.

Reverse engineering is a powerful tool for analyzing, controlling, and documenting existing database objects. You can use the extracted SQL source code for archival and reference purposes.

### Important Notes

- You must create a New Project from a Database in order to access the New Project Reverse Engineering Wizard.
- Rapid SQL automatically discovers and sets the project script file build order by referencing the system catalog to determine dependencies.
- In the [New Project dialog box](#), you must select the From Database option.

For more information, see [Creating a New Project from a Database](#).

## CREATING A NEW PROJECT FROM A DATABASE

To create a new project, do the following:

- 1 On the *File* menu, click *New*, and then click *Project*.  
OR  
On the *Main* toolbar, click the drop-down arrow to the right of the *New* button.  
OR  
On the *Project* toolbar, click the *New Project*.  
OR  
In the workspace, right-click *New*, and then click *Project*.  
Rapid SQL opens the *New Project* dialog box.
- 2 Complete the [New Project](#) dialog box and select the *From Database* option.
- 3 Click *OK*.  
Rapid SQL opens the *New Project Reverse Engineering Wizard*.
- 4 Complete the [New Project Reverse Engineering Wizard](#).

## COMPLETING THE NEW PROJECT REVERSE ENGINEERING WIZARD

**NOTE:** You must create a [New Project from a Database](#) in order to access the *New Project Reverse Engineering Wizard*.

**NOTE:** Rapid SQL automatically discovers and sets the project script file build order by referencing the system catalog to determine dependencies.

To complete the *New Project Reverse Engineering Wizard*, do the following:

- 1 Complete the *New Project* dialog box, and then click *OK*.  
**NOTE:** You must select the *From Database* option.  
Rapid SQL opens the *New Project Reverse Engineering Wizard*.
- 2 Complete the wizard.
- 3 To make changes, click *Back* to the appropriate panels of the Wizard and make changes.
- 4 When you are satisfied with the definition, click *Execute*.  
Rapid SQL starts creating project files from database.

For more information, see [New Project Reverse Engineering Wizard - Panel](#).

**NEW PROJECT REVERSE ENGINEERING WIZARD - PANEL 1**

The first panel of the *New Project Reverse Engineering Wizard* lets you specify the datasource that you want to reverse engineer into a project.

The table below describes the options and functionality of the first panel of the New Project Reverse Engineering Wizard:

Option	Description
Select a Datasource	Lets you select a datasource to reverse engineer. Select a datasource.

For more information, see [Completing the New Project Reverse Engineering Wizard](#).

**NEW PROJECT REVERSE ENGINEERING WIZARD - PANEL 2**

The second panel of the New Project Reverse Engineering Wizard lets you specify the target database on the server.

The table below describes the options and functionality of the second panel of the New Project Reverse Engineering Wizard:

Option	Description
Select Database to reverse-engineer	Lets you select a database to reverse engineer. Select a database.

For more information, see [Completing the New Project Reverse Engineering Wizard](#).

**NEW PROJECT REVERSE ENGINEERING WIZARD - PANEL 2**

The third panel of the New Project Reverse Engineering Wizard lets you specify the database object owners, the database object types, and general script types of the objects you want to extract.

The table below describes the options and functionality of the third panel of the New Project Reverse Engineering Wizard:

Option	Description
Select Database Object Types	Owner list - Lets you select the owner of the objects you want to extract. Objects - Lets you select the objects types you want to extract. TIP: To select or deselect all object types right-click to Select or Unselect all.
Extract Scope	Lets you select the objects to extract from the database. All Objects - Lets you extract all objects for each object type that you select. Opens <a href="#">Panel 4 for All Objects</a> . Selected Objects Only - Lets you specify the specific objects that you want to extract. Opens <a href="#">Panel 4 for Selected Objects</a> . NOTE: Rapid SQL continues to a different Panel 4 depending on which option you choose for the Extract Scope.

Option	Description
Include Indexes with Table DDL	If selected, includes indexes when extracting table DDL.
Include FKs with Table DDL	If selected, includes foreign keys when extracting table DDL.
Include Drop Statement	If selected, includes drop statements when extracting database SQL.
Include Object Privileges	If selected, includes object privileges when extracting database SQL.

For more information, see [Completing the New Project Reverse Engineering Wizard](#).

### New Project Reverse Engineering Wizard - Panel 3 for All Objects

This panel of the New Project Reverse Engineering Wizard lets you specify the ownership.

The table below describes the options and functionality of the fourth panel of the New Project Reverse Engineering Wizard for All Objects:

Option	Description
Retain	Select to include the existing object owners' names in the CREATE statements.
Exclude	Select to exclude the object owners' names in the CREATE statements.
Transfer	Select to include the designated owner's name in the CREATE statements. Select a user.

For more information, see [Completing the New Project Reverse Engineering Wizard](#).

### NEW PROJECT REVERSE ENGINEERING WIZARD - PANEL 4 FOR SELECTED OBJECTS

This panel of the New Project Reverse Engineering Wizard only applies if you have chosen to generate a customized script. You can select specify objects and set script options for each object type.

The table below describes the options and functionality of the fourth panel of the New Project Reverse Engineering Wizard for selected objects.

**NOTE:** You must select options and statements for each object type you select.

Option	Description
Object Type	Lets you select an object type.
Object list	Lists the owner and names of objects for the selected type in the database.
Options for Objects	Lets you select an option for the object type. NOTE: Options vary depending on the object type selected.

Option	Description
Statements	Lets you select a script option for the particular object type. NOTE: Statements vary depending on the option and object type selected.

For more information, see [Completing the New Project Reverse Engineering Wizard](#).

#### NEW PROJECT REVERSE ENGINEERING WIZARD - PANEL 4

This panel of the New Project Reverse Engineering Wizard lets you review the Reverse Engineering definition to verify its accuracy.

**NOTE:** If you have specified to integrate Rapid SQL with an underlying version control system during installation, Rapid SQL asks to add the new files to version control. Click Yes to add the files.

For more information, see [Completing the New Project Reverse Engineering Wizard](#).

## CREATE A NEW PROJECT FROM EXISTING FILES

Rapid SQL lets you create a new project from:

- Scratch.
- Existing files.

To create a new project from an existing file, do the following:

- 1 On the *File* menu, click *New*, and then click *Project*.

OR

On the *Main* toolbar, click the drop-down arrow to the right of the *New* button.

OR

On the *Project* toolbar, click the *New Project*.

OR

In the workspace, right-click *New*, and then click *Project*.

Rapid SQL opens the *New Project* dialog box.

- 2 Complete the [New Project dialog box](#) and select the *From Existing Files* option.
- 3 Click *OK*.

Rapid SQL opens the *Add File(s) to Project* dialog box.

- 4 In the *Add File(s) to Project* dialog box, type the files that you want to add in the *File* name field. Use the following to help locate the files:
  - In *Directories*, browse by clicking a directory.
  - In the *List files of type* list, select a file type.
  - In the *Drives* list, select a drive.
- 5 Click *Add* to add a file or click *Add All* to add all the files.
- 6 When you finish, click *OK*.

Rapid SQL creates the project.

**NOTE:** If you have specified to integrate Rapid SQL with an underlying version control system during installation, Rapid SQL asks to add the new files to version control. Click *Yes* to add the files.

## CREATE A NEW PROJECT FROM A VERSION CONTROL PROJECT

Rapid SQL lets you reverse-engineer a project from existing version control system projects. This is helpful if you have already created database projects in a version control system.

**NOTE:** The Intersolv PVCS API does not support the creation of a project from a source code control project.

Rapid SQL also lets you create a project from a version control project that contains sub-directories, while including files from those sub-directories in the Rapid SQL project.

To create a new project from a Version Control project, do the following:

- 1 On the *File* menu, click *New*, and then click *Project*.

OR

On the *Main* toolbar, click the drop-down arrow to the right of the *New* button.

OR

On the *Project* toolbar, click the *New Project*.

OR

In the workspace, right-click *New*, and then click *Project*.

Rapid SQL opens the *New Project* dialog box.

- 2 Complete the [New Project dialog box](#) and select *From Version Control Project*.
- 3 Click *OK*.

Rapid SQL opens the *Choose project from (Version Control name)* dialog box.

- 4 Select the project or files you want to include, and then click *OK*.

Rapid SQL opens the *Files to be included in (project path)* dialog box.



- 5 Select the project of files you want to include, and then click OK.

Rapid SQL creates the project.

## CREATE A NEW PROJECT WITHOUT INITIALIZATION

Rapid SQL gives you the option to not initialize a project. Rapid SQL creates a project tab with the name you specify in the New Project dialog box. This tab functions as a shell to:

- [Add files.](#)
- [Add database objects.](#)

To create a new project without initialization, do the following:

- 1 On the *File* menu, click *New*, and then click *Project*.

OR

On the *Main* toolbar, click the drop-down arrow to the right of the *New* button.

OR

On the Project toolbar, click the *New Project*.

OR

In the workspace, right-click *New*, and then click *Project*.

Rapid SQL opens the New Project dialog box.

- 2 Complete the [New Project dialog box](#) and select the *Do Not Initialize* option.
- 3 Click *OK*.

Rapid SQL creates the Project Tab.

## WORKING WITH PROJECTS

A project is similar to a file system. Both are a collection of files that you create and maintain. Because projects are hierarchical, you can place a [subproject](#) under another project, and a subproject under a subproject. Once you have created a new project, Rapid SQL provides many functions to help you maintain and modify the project.

### Related Topics

[New Project Dialog Box](#)

[Opening an Existing Project](#)

[Opening a Recent Project](#)

[Closing a Project](#)

[Build Project](#)

[Set Build Order](#)

[Add Database Object File\(s\) to Project Wizard](#)

[Execute Project Files](#)

[Add File\(s\) to a Project](#)

[Open a File from a Project](#)

[Subprojects](#)

[Project Properties](#)

[Confirm Delete Dialog Box](#)

## OPENING AN EXISTING PROJECT

To open an existing project, do the following:

- 1 On the *File* menu click *Open Project*.  
Rapid SQL opens the Open Project dialog box.
- 2 In *File name*, type the name and location of the project or use browse to locate the project.  
**NOTE:** Project files are designated with a \*.epj extension.
- 3 Click Open.  
Rapid SQL opens the Project Tab containing the project.

For more information, see [Working with Projects](#).

## OPENING A RECENT PROJECT

To open a recent project, do the following:

- 1 On the *File* menu select *Recent Project*, and then select a project.  
Rapid SQL opens the Project Tab containing the project.

For more information, see [Working with Projects](#).

## CLOSING A PROJECT

To close a project, do the following:

- 1 On the *File* menu, click *Close Project*.
- 2 On the *Project Tab*, right-click and then click *Close Project*.

Rapid SQL closes the project.

For more information, see [Working with Projects](#).

## BUILD PROJECT

The Build Project dialog box lets you:

- Generate a project build script and display it in a SQL window.
- Execute the project build immediately on build target.
- Schedule a project build.

### Important Notes

None

The table below describes the Build Project dialog box:

Option	Description
Datasource	Displays the target datasource.
Database	Displays the target database.
Build subprojects	Select to include subprojects.
Generate a Project Build Script and Display it in a SQL window.	Select to generate the project build script and have Rapid SQL display it in a SQL window.
Execute Project Build Immediately on Build Target	Select to execute build immediately.
Schedule Project Build for a Later Time	Select to schedule the build and then specify optional options.

For more information, see [Completing the Build Project Dialog Box](#).

## COMPLETING THE BUILD PROJECT DIALOG BOX

To complete the Build Project dialog box, do the following:

- 1 On the *File* menu click *Open Project*.

Rapid SQL opens the Open Project dialog box.

- 2 In *File name*, type the name and location of the project or use browse to locate the project.

**NOTE:** Project files are designated with a \*.epj extension.

- 3 Click Open.

Rapid SQL opens the Project Tab containing the project.

- 1 On the Project Tab, right-click the target project, and then select Build.

Rapid SQL opens the Build Project dialog box.

- 2 Select options.

- 3 Click OK.

Rapid SQL builds the project according to your specifications.

## SET BUILD ORDER

You can specify the order in which you want Rapid SQL to build your project files in the Set Build Order dialog box. If you created your project manually or from a version control project, you must specify a build order, otherwise the files are built in the order that they appear in the tree of the Project Tab.

### Setting Build Order

- 1 On the *File* menu click *Open Project*.

Rapid SQL opens the Open Project dialog box.

- 2 In *File name*, type the name and location of the project or use browse to locate the project.

**NOTE:** Project files are designated with a \*.epj extension.

- 3 Click Open.

Rapid SQL opens the Project Tab containing the project.

- 4 On the *Project* menu, click Build Order.

OR

On the *Project* toolbar, click Build Order.

OR

On the Project Tab, right-click a project file, and then click Build.

Rapid SQL opens the Set Build Order dialog box.

- 5 Click the files you want to move and then click the *Up* and *Down* to change the order to build the files.

- 6 When you finish specifying the order, select OK.

Rapid SQL sets the build order.

**NOTE:** The next time you build the project, Rapid SQL uses this build order.

## ADD DATABASE OBJECT FILE(S) TO PROJECT WIZARD

You can add database objects to an existing project using a simple wizard that can reverse-engineer an entire database or any portion of it. This lets you keep your project in sync with databases where objects are constantly being created and updated.

For more information, see:

[Working with Projects](#)

[Completing the Add Database Object File\(s\) to Project Wizard](#)

### COMPLETING THE ADD DATABASE OBJECT FILE(S) TO PROJECT WIZARD

To add database objects to a project, do the following:

- 1 On the *File* menu click *Open Project*.

Rapid SQL opens the Open Project dialog box.

- 2 In *File name*, type the name and location of the project or use browse to locate the project.

**NOTE:** Project files are designated with a \*.epj extension.

- 3 Click Open.

Rapid SQL opens the Project Tab containing the project.

- 4 On the *Project* menu, click Add Database Objects.

OR

On the *Project* toolbar click Add Database Objects.

OR

On the Project Tab, right-click and then click Add Database Objects.

Rapid SQL opens the [first panel of the Add Database Object File\(s\) to Project Wizard](#).

For more information, see [Add Database Objects to a Project Wizard](#).

### ADD DATABASE OBJECT FILE(S) TO PROJECT WIZARD - PANEL 1

The first panel of the Add Database Object File(s) to Project Wizard lets you specify the datasource that you want to reverse engineer into a project.

The table below describes the options and functionality of the first panel of the Add Database Object File(s) to Project:

Option	Description
Select a Datasource	Lets you select a datasource to reverse engineer.

For more information, see [Completing the Add Database Object File\(s\) to Project Wizard](#).

#### ADD DATABASE OBJECT FILE(S) TO PROJECT WIZARD - PANEL 2

The second panel of the Add Database Object File(s) to Project Wizard lets you specify the target database on the server.

The table below describes the options and functionality of the second panel of the Add Database Object File(s) to Project:

Option	Description
Select Database to reverse-engineer	Lets you select a database to reverse engineer.

For more information, see [Completing the Add Database Object File\(s\) to Project Wizard](#).

#### ADD DATABASE OBJECT FILE(S) TO PROJECT WIZARD - PANEL 3

The third panel of the Add Database Object File(s) to Project Wizard lets you specify the database object owners, the database object types, and general script types of the objects you want to extract.

The table below describes the options and functionality of the third panel of the Add Database Object File(s) to Project:

Option	Description
Select Database Object Types	Owner list - Lets you select the owner of the objects you want to extract. Objects - Lets you select the objects types you want to extract. TIP: To select or deselect all object types, in the Object list, right-click to Select or Unselect all.
Extract Scope	Lets you select the objects to extract from the database. All Objects - Lets you extract all objects for each object type that you select. Opens <a href="#">Panel 4 for All Objects</a> . Selected Objects Only - Lets you specify the specific objects that you want to extract. Opens <a href="#">Panel 4 for Selected Objects</a> . NOTE: Rapid SQL opens a different Panel 4 depending on which option you choose for the Extract Scope.
Include Indexes with Table DDL	Includes indexes when extracting table DDL.
Include FKs with Table DDL	Includes foreign keys when extracting table DDL.
Include Drop Statement	Includes drop statements when extracting database SQL.

Option	Description
Include Object Privileges	Includes object privileges when extracting database SQL.

For more information, see [Completing the Add Database Object File\(s\) to Project Wizard](#).

#### ADD DATABASE OBJECT FILE(S) TO PROJECT WIZARD - PANEL 4 FOR SELECTED OBJECTS

This panel of the Add Database Object File(s) to Project Wizard only applies if you have chosen to generate a customized script. You can select specify objects and set script options for each object type.

The table below describes the options and functionality of the fourth panel of the Add Database Object File(s) to Project for selected objects.

**NOTE:** You must select options and statements for each object type you select.

Option	Description
Object Type list	Lets you select an object type.
Object list	Lists the owner and names of objects for the selected type in the database.
Options for Object list	Lets you select an option for the object type. NOTE: Options vary depending on the object type selected.
Statement list	Lets you select a script option for the particular object type. NOTE: Statements vary depending on the option and object type selected.

For more information, see [Completing the Add Database Object File\(s\) to Project Wizard](#).

#### Add Database Object File(s) to Project Wizard - Panel 4 for All Objects

The fourth panel of the Add Database Object File(s) to Project Wizard lets you specify the ownership.

The table below describes the options and functionality of the fourth panel of the Add Database Object File(s) to Project for All Objects:

Option	Description
Retain	Select to include the existing object owners' names in the CREATE statements.
Exclude	Select to exclude the object owners' names in the CREATE statements.
Transfer	Select to include the designated owner's name in the CREATE statements. Select a user.

For more information, see [Completing the Add Database Object File\(s\) to Project Wizard](#).

**ADD DATABASE OBJECT FILE(S) TO PROJECT WIZARD - PANEL 5**

This panel of the Add Database Object File(s) to Project Wizard lets you retain, exclude, or transfer object ownership.

The table below describes the options and functionality on this panel:

Option	Description
Retain	Select to retain object ownership.
Exclude	Select to exclude object ownership.
Transfer ownership to following user:	Select to transfer ownership, and then select the target user from the list.

For more information, see [Completing the Add Database Object File\(s\) to Project Wizard](#).

**ADD DATABASE OBJECT FILE(S) TO PROJECT WIZARD - PANEL 6**

This panel of the Add Database Object File(s) to Project Wizard lets you review the Reverse Engineering definition to verify its accuracy.

**NOTE:** If you have specified to integrate Rapid SQL with an underlying version control system during installation, Rapid SQL asks to add the new files to version control. Click Yes to add the files.

For more information, see [Completing the Add Database Object File\(s\) to Project Wizard](#).

**EXECUTE PROJECT FILES**

You can directly execute project script files from the Project Tab using the File Execution Facility. You can also execute multiple scripts in parallel against different datasources. If a file has been placed under version control, you need to perform a Check Out operation to execute the file. Otherwise, the file opens in read-only mode.

For more information, see:

[Working with Projects](#)

[Executing Project Files](#)

**EXECUTING PROJECT FILES**

To execute project files, do the following:

- 1 On the *Project Tab*, right-click the files you want to execute, and then click *File Execution Facility*.

Rapid SQL opens the [Script Execution Facility](#) dialog box.

For more information, see [Execute Project Files](#).



## ADD FILE(S) TO A PROJECT

You can add external files to your project. This lets you manipulate and expand your project as needed.

The Add File(s) to Project dialog box lets you add files to a project.

The table below describes the options and functionality of the Add Files to Project dialog box:

Option	Description
File Name	Lets you type the files that you want to add.
Directories	Lets you browse by select a directory.
List files of type	Lets you select a file type.
Drives	Lets you select a drive.
Network	Open the Map Network Drive dialog box.
Add	Lets you add a file.
Add All	Lets you add all the files.

For more information, see:

[Working with Projects](#)

[Completing the Add Files to Project Dialog Box](#)

### COMPLETING THE ADD FILE(S) TO PROJECT DIALOG BOX

To add files to a project from, do the following:

- 1 On the *Project* menu, click Add Files.  
OR  
On the Project Tab, right-click and then click Add Files.  
Rapid SQL opens the [Add Files to Project](#) dialog box.
- 2 In *File name*, type the file name(s) or select the file(s) that you want to add to the project.
- 3 In *Directories* browse for the project. that contains the file(s) that you want to add to the project.
- 4 In the *List files of type* select the type of files you want to add to the project.
- 5 In *Drives* select a drive.
- 6 Click *Network* to open the Map Network Drive dialog box.
- 7 Click *Add* to add a file or click *Add All* to add all of the files.
- 8 Click *OK*.

Rapid SQL adds the files to the project.

## OPEN A FILE FROM A PROJECT

You can open a file directly from the Project Tab using a number of different methods.

For more information, see:

[Working with Projects](#)

[Opening a File from a Project](#)

## OPENING A FILE FROM A PROJECT

To open a file from a project, do the following:

- 1 On the *Project Tab*, right-click the file(s) you want to open, and then click *Open*.

OR

On the *Project Tab*, double-click the file(s).

Rapid SQL opens the file(s) in a new SQL window.

For more information, see [Open a File from a Project](#).

## SUBPROJECTS

Subprojects are projects within projects. you can use them to help categorize your source code files. On the Project Tab of the Database Explorer Rapid SQL lets you:

- [Create subprojects](#).
- [Delete subprojects](#).
- [Rename subprojects](#).
- [Sort subprojects](#).

For more information, see [Working with Projects](#).

## CREATING A NEW SUBPROJECT

Rapid SQL offers three ways to create a new subproject:

- 1 On the *File* menu click *Open Project*.  
Rapid SQL opens the Open Project dialog box.
- 2 In *File name*, type the name and location of the project or use browse to locate the project.

**NOTE:** Project files are designated with a \*.epj extension.

- 3 Click Open.

Rapid SQL opens the Project Tab containing the project.

- 4 On the *Project* menu, click New SubProject.  
OR  
On the *Project* toolbar, click New SubProject.  
OR  
On the Project Tab, right-click the project or subproject, and then click New SubProject.
- 5 Type the name of the new subproject and then press *Enter*.

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

#### DELETING A SUBPROJECT

To delete a project, do the following:

- 1 On the *File* menu click *Open Project*.  
Rapid SQL opens the Open Project dialog box.
- 2 In *File name*, type the name and location of the project or use browse to locate the project.  
**NOTE:** Project files are designated with a \*.epj extension.
- 3 Click Open.  
Rapid SQL opens the Project Tab containing the project.
- 4 On the Project Tab, right-click a subproject, and then click Delete.  
OR  
On the Project Tab, select a subproject, and then click Delete.  
Rapid SQL opens the Confirm Delete dialog box.
- 5 In the *Confirm Delete* dialog box, select the *Delete* local copy to delete the local copy of the subproject, and then click *OK*.  
Rapid SQL deletes the subproject.

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

#### CONFIRM DELETE DIALOG BOX

The table below describes the options and functionality on the Confirm Delete Dialog Box:

Option	Description
Delete local copy	Deletes the local copy of the subproject.

For more information, see [Deleting a Subproject](#).

### RENAMING A SUBPROJECT

To rename a subproject, do the following:

- 1 On the *File* menu click *Open Project*.

Rapid SQL opens the Open Project dialog box.

- 2 In *File name*, type the name and location of the project or use browse to locate the project.

**NOTE:** Project files are designated with a \*.epj extension.

- 3 Click Open.

Rapid SQL opens the Project Tab containing the project.

- 4 On the *Project Tab*, right-click a subproject, and then click *Rename*.

- 5 Type the name of the new subproject and then press *Enter*.

Rapid SQL renames the subproject.

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

### SORTING SUBPROJECTS

To sort a subproject, do the following:

- 1 On the *File* menu click *Open Project*.

Rapid SQL opens the Open Project dialog box.

- 2 In *File name*, type the name and location of the project or use browse to locate the project.

**NOTE:** Project files are designated with a \*.epj extension.

- 3 Click Open.

Rapid SQL opens the Project Tab containing the project.

- 4 On the *Project Tab*, right-click the directory containing the subprojects you want to sort and then click *Sort*.

Rapid SQL sorts the subproject(s).

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

### PROJECT PROPERTIES

Rapid SQL lets you view properties of Projects, Subprojects and individual files in Projects or Subprojects. The [Project Properties](#), [Subproject Properties](#), and [File Properties](#) dialog boxes display information about the Projects, Subprojects, and files.

For more information, see [Working with Projects](#).

## VIEWING PROJECT PROPERTIES

To view project properties, do the following:

- 1 On the *File* menu click *Open Project*.

Rapid SQL opens the Open Project dialog box.

- 2 In *File name*, type the name and location of the project or use browse to locate the project.

**NOTE:** Project files are designated with a \*.epj extension.

- 3 Click Open.

Rapid SQL opens the Project Tab containing the project.

- 4 On the *Project* menu, click Project Properties.

OR

On the Project Tab, right-click a project and then click Properties.

Rapid SQL opens the Project File Properties dialog box.

- 5 Complete the [Project File Properties dialog box](#).

- 6 Click OK.

For more information, see:

[Project Properties](#)

[Project File Properties Dialog Box](#)

## PROJECT FILE PROPERTIES DIALOG BOX

The table below describes the options and functionality of the Project Properties dialog box:

Option	Description
Name	Displays the name of the project.
Full Specification	Displays the project location.
Description	Displays the project description. OPTIONAL: Lets you type or edit the project description.
Associated Datasource	Lets you select a datasource for the project.
Associated Database	Lets you type the name of the database.

For more information, see [Viewing Project Properties](#).

## VIEWING SUBPROJECT PROPERTIES

To view subproject properties, do the following:

- 1 On the *File* menu click *Open Project*.  
Rapid SQL opens the Open Project dialog box.
- 2 In *File name*, type the name and location of the project or use browse to locate the project.  
**NOTE:** Project files are designated with a \*.epj extension.
- 3 Click Open.  
Rapid SQL opens the Project Tab containing the project.
- 4 On the Project Tab, select a subproject.
- 5 Right-click the subproject, and then click Properties.  
Rapid SQL opens the Subproject Properties dialog box.
- 6 Complete the [Subproject Properties dialog box](#).
- 7 Click OK.

For more information, see:

[Project Properties](#)

[Subproject Properties Dialog Box](#)

## SUBPROJECT PROPERTIES DIALOG BOX

The table below describes the options and functionality of the Subproject Properties dialog box:

Option	Description
Name	Displays the name of the subproject. You can enter or edit the subproject name.
Full Specification	Displays the subproject location.
Status	Displays the subproject status.

For more information, see [Viewing Subproject Properties](#).

## VIEWING FILE PROPERTIES

To view file properties, do the following:

- 1 Click File and then Open Project.  
Rapid SQL opens the Open Project dialog box.
- 2 Type the project name or select the project.

- 3 Click Open.

Rapid SQL opens the project on the Projects Tab.

- 4 On the Project Tab, select a file.
- 5 Right-click the file, and then click Properties.

Rapid SQL opens the File Properties dialog box.

- 6 Complete the [File Properties dialog box](#).
- 7 Click OK.

For more information, see:

[Project Properties](#)

[File Properties Dialog Box](#)

#### FILE PROPERTIES DIALOG BOX

The table below describes the options and functionality of the File Properties dialog box:

Option	Description
Name	Displays the file name.
Full Specification	Displays the file location.
Description	Displays the file description. OPTIONAL: Type or edit the file description.
Include In Build	Sets the file to be included in the build.
Last Modified	Displays the date and time of the last modification.
Size	Displays the file size.
Status	Displays the file status.
Object Type	Displays the file object type. Select an object type if the file is unspecified.





# VERSION CONTROL

Version control archives files and tracks changes to files over time. With an integrated version control in Rapid SQL, you can easily track changes to database objects.

Version control addresses the following issues:

- **Team Development** By controlling access to a file so that only one person at a time can modify, it prevents accidental replacement or loss of another user's changes.
- **Version Tracking** By archiving and tracking versions of source code files, you can retrieve them if necessary, thereby effectively creating files so that source code can be reused.
- **Safety** By adding database object scripts and files, it creates backups in case of loss, thereby ensuring a recovered version of source code.

When you create a project, Rapid SQL lets you place the project immediately into version control. You can also add projects and files to version control later.

To use integrated version control in Rapid SQL, you must have the version control client software installed on the same computer as Rapid SQL. You must also select the appropriate version control system during installation or after installation on the [Version Control Tab of the Option Editor](#).

Rapid SQL offers version control integration for the following version control systems:

- Rational ClearCase 5.0 and 6.0
- [Merant/Intersolv PVCS Version Manager version 6.0](#)
- [Microsoft Visual SourceSafe 5.0 and 6.0](#)
- [MKS Source Integrity version 7.3c](#)

For more information, see:

[Version Control Integration](#)

[Version Control Configuration](#)

[Using Version Control](#)

## VERSION CONTROL INTEGRATION

Rapid SQL offers version control integration so you can take advantage of the following version control systems:

- Rational ClearCase 5.0 and 6.0
- [Merant Version Manager version 6.0 or later](#)
- [Microsoft Visual SourceSafe 5.0 and 6.0](#)
- [MKS Source Integrity version 7.3c](#)

For more information, see:

[Integrating with Merant Version Manager](#)

[Integrating with Microsoft Visual Source Safe](#)

[Integrating with MKS Source Integrity](#)

### INTEGRATING WITH MERANT VERSION MANAGER

Rapid SQL works with Merant Version Manager version 6.0. or later.

To integrate Merant Version Manager with Rapid SQL, do the following:

- 1 Install Merant Version Manager (formerly PVCS) 8.0.
- 2 In the setup option for the client, select IDE Client.

This installs the correct DLLs and registry entries, and "Merant Version Manager" appears in the Version Control Tab of the Options Editor.

**NOTE:** To use Rapid SQL with an existing Merant project, you must import the individual files into a native Rapid SQL project. Merant does not provide the third party API support for projects that it does for basic version control operations on archived files.

For more information, see [Version Control Integration](#).

### INTEGRATING WITH MICROSOFT VISUAL SOURCESAFE

Rapid SQL works with Microsoft Visual SourceSafe versions 5.0. and 6.0. Rapid SQL uses your default Visual SourceSafe database unless you override the setting on the Version Control Tab of the Option Editor. This file is always called srcsafe.ini.

### Troubleshooting

If you are having trouble configuring Rapid SQL to use Visual SourceSafe, check the system registry to determine that the COM automation portion of Visual SourceSafe has been properly installed. To check the registry, do the following:

- 1 From the Windows *Start* button, click *Run* and then type REGEDIT to view the system registry.

Windows opens the Registry Editor.

- 2 Click the *HKEY\_CLASSES\_ROOT* directory.
- 3 Click the *SourceSafe* key.
- 4 Ensure that the Key contains the subkeys *CLSID* and *CurVer*.
- 5 If either of these keys are missing, reinstall Visual SourceSafe.

For more information, see [Version Control Integration](#).

## INTEGRATING WITH MKS SOURCE INTEGRITY

Rapid SQL works with MKS Source Integrity version 7.3c.

- 1 After installing MKS Source Integrity, install the MKS Source Integrity Extensions.

For more information, see [Version Control Integration](#).

## VERSION CONTROL CONFIGURATION

Rapid SQL lets you configure version control. You can add or remove entire projects or specific files to and from version control. The Version Control Tab of the Option Editor lets you configure version controls to your specific needs.

To use version control functions, you must have a version control system up and running on your system. You can integrate Rapid SQL with:

- Rational ClearCase 5.0 and 6.0
- [Merant Version Manager version 6.0 or later](#)
- [Microsoft Visual SourceSafe 5.0 and 6.0](#)
- [MKS Source Integrity version 7.3c](#)

### Configuring Version Control During Installation

- 1 If you have version control installed on your machine, during the Rapid SQL installation, select the option for your version control system.

**NOTE:** If you select None, Rapid SQL does not support any version control functions in Rapid SQL until you select a version control system from the application. See below.

### Configuring Version Control After Installation

- 1 On the *File* menu, click *Options*.  
Rapid SQL opens the Options Editor.
- 2 Select *Version Control* to open the *Version Control* Tab.
- 3 Select your version control system.
- 4 Click OK.

## WORKING WITH PROJECTS IN VERSION CONTROL

Projects can exist independently of version control. You can add an entire project to version control at any point. When you decide to place a project under version control, Rapid SQL creates a project on the underlying version control system that has been specified.

Once a project or file has been added to version control, the features of your version control system are available directly from the Project Tab within Rapid SQL. Any changes you make to a project or file from within Rapid SQL are simultaneously changed in your version control system.

The table below describes the options and functionality on the Add to Version Control dialog box:

Option	Description
Files to be added	Lets you select files to add to Version Control.
Comment	OPTIONAL: Lets you add a comment.
Check out immediately	Select to add file and keep it checked-out.
Store only latest	Select to add the latest version.
Remove local copy	Select to add file and remove the local copy.

For more information, see:

[Completing the Add to Version Control Dialog Box](#)

### COMPLETING THE ADD TO VERSION CONTROL DIALOG BOX

To add a project to version, do the following:

- 1 Click File and then Open Project.  
Rapid SQL opens the Open Project dialog box.
- 2 Type the project name or select the project.
- 3 Click Open.  
Rapid SQL opens the project on the Projects Tab.
- 4 On the *Project* menu, click Version Control.  
OR  
On the Project Tab, right-click a project file.  
Rapid SQL opens the Add to Version Control dialog box.
- 5 Click Add to Version Control.
- 6 Select options.

- 7 Click *OK* to add a project, and then click *Create*.

Rapid SQL adds the project to version control.

**NOTE:** Rapid SQL dims the project file icon to indicate that the project has been placed under version control.

For information on opening projects, see [Working with Projects](#).

## WORKING WITH FILES IN VERSION CONTROL

Rapid SQL lets you view and work with files stored in various projects in a version control system without including the files in a Rapid SQL project. The Add Version Control Files dialog box lets you create a list of files in a version control system, select the version control project, and then add it and its files to the [VC Files Tab](#).

The table below describes the options and functionality on the Add Version Control Files dialog box:

Option	Description
Project	Lets you type the version control project path.
Browse	Click to open Choose project from (version control name) dialog box.
List Files	Click to view the project's files in the Version control project files tree.
File Types	Click to list the project's available files, filtered by file type.
Version control project files	List the project available files.  NOTE: If the VC Files Tab is already open and contains files from the project previously selected in the Add Version Control Files dialog box, only files not already in the VC Files Tab will be listed.
Add	Click to add selected file(s) to the Files being added to version control files list box.
Add All	Click to add all file(s) to the Files being added to version control files list box.
Files being added to version control files list	Displays the files that will appear on the VC Tab.
Remove	Click to remove selected file from the Files being added to version control files list box.
Check Out	Select to automatically check out the file from version control on the VC Files Tab.
Get Latest Version	Select to automatically get latest version of the file from version control on the VC Files Tab.

For more information, see:

## Completing the Add Version Control Files Dialog Box

### COMPLETING THE ADD VERSION CONTROL FILES DIALOG BOX

To add files to version control, do the following:

- 1 Click File and then Open Version Control Files List.

Rapid SQL opens the version control system login dialog box.

- 2 Type login information.

Rapid SQL opens the Add Version Control Files dialog box.

- 3 In Project, type the project name, or click Browse to open the Choose project from (version control name) dialog box.

- 4 Select the project and click OK.

Rapid SQL displays the files in the Version control project files box.

- 5 Click File Types to list the project's available files, filtered by file type.

- 6 In List files of type select the type of files to list.

- 7 In the Version control project files box, select the target files and click Add or Add All to add the files to the version control files list.

- 8 In Options, select Check Out or Get Latest Version.

- 9 Click OK.

Rapid SQL opens the Check Out File or Get File dialog box.

- 10 In Comment, type a comment. If you select multiple files, the comment will apply to all the files.

- 11 In To, type the directory to place the file(s).

- 12 For SourceSafe, for advanced options, click Advanced.

- 13 Click OK.

Rapid SQL adds the files in the Files being added to version control files list box to the [VC Files Tab](#). If you selected the Check Out option, Rapid SQL opens the [Check Out \(Files\)](#) dialog box. If you selected the Get Latest Version option, Rapid SQL opens the [Get from Version Control](#) dialog box.

## USING VERSION CONTROL

Once you add a project or file to version control, most functions found in the underlying version control system are available directly from the Rapid SQL interface, including the [VC Files Tab](#). Basic version control procedures include:

- [Get Latest Version](#)
- [Check Out](#)
- [Check In](#)
- [Undo Check Out](#)
- [Open](#)
- [Show History](#)
- [Show Difference](#)
- [Version Control Properties](#)

**NOTE:** Your Rapid SQL version control functionality depends on your underlying version control system. For more information on version control procedures, consult the documentation included with your version control system.

## VC FILES TAB

The VC Files Tab displays version control files listed in the \*.evc (Embarcadero version control file.) The tab displays file icons indicating their current status, for example if they are checked out to the user logged in to the source control system. The files can be opened from this list, as well as operated on to manipulate their version control properties. For example, a file can be checked out or checked into the system from the VC File Tab.

The table below describes the VC Files Tab icons:

Icon	Description
File	File is not checked out by anyone.
File with single red check mark	File is checked out non-exclusively and only by the user logged into source control.
File with a single black check mark	File is checked out non-exclusively only by a one user who is not the user logged into source control.
File with a single red check mark and a red border	File is checked out exclusively and only by the user logged in to source control.
File with a single black check mark and a red border	File is checked out exclusively and only by a user who is not the user logged in to source control.
File with two red check marks	File is checked out by multiple users, including the user logged in to source control.
File with two black check marks	File is checked out by multiple users, not including the user logged in to source control.

For more information, see:

[VC Files Tab Available Functionality](#)

[Opening the VC Files Tab](#)

[Closing the VC Files Tab](#)

#### VC FILES TAB AVAILABLE FUNCTIONALITY

Once the files are on the list in the VC Files Tab, you can use the various version control functionalities.

At the VC Files level, Rapid SQL lets you:

- [Add Files](#)
- [Sort](#)
- [Get Latest Version](#)
- [Check Out](#)
- [Check In](#)
- [Undo Checkout](#)
- [Remove from Version Control](#)
- [Expand All](#)
- [Collapse All](#)
- [Refresh](#)
- [Close Files List](#)

At the Project level, Rapid SQL lets you:

- [Add Files](#)
- [Delete](#)
- [Sort](#)
- [Get Latest Version](#)
- [Check Out](#)
- [Check In](#)
- [Undo Checkout](#)
- [Remove from Version Control](#)
- [Expand All](#)
- [Collapse All](#)
- [Refresh](#)



- [Close Files List](#)

At the Directory (file type) level, Rapid SQL lets you:

- [Add Files](#)
- [Delete](#)
- [Sort](#)
- [Get Latest Version](#)
- [Check Out](#)
- [Check In](#)
- [Undo Checkout](#)
- [Remove from Version Control](#)
- [Expand All](#)
- [Collapse All](#)
- [Refresh](#)
- [Close Files List](#)

At the File level, Rapid SQL lets you:

- [Open](#)
- [Delete](#)
- [Get Latest Version](#)
- [Check Out](#)
- [Check In](#)
- [Undo Checkout](#)
- [Show History](#)
- [Show Differences](#)
- [Remove from Version Control](#)
- [Version Control Properties](#)
- [Refresh](#)
- [Close Files List](#)

For more information, see:

[Opening the VC Files Tab](#)

[Closing the VC Files Tab](#)

## [Working with Files in Version Control](#)

### OPENING THE VC FILES TAB

- 1 Select File, Open Version Control File List.

If you have files on your version control list, Rapid SQL opens the VC Files Tab. If you do not have files on your version control list, Rapid SQL opens the Add Version Control Files dialog box.

For more information, see:

[Working with Files in Version Control](#)

[Closing the VC Files Tab](#)

### CLOSING THE VC FILES TAB

- 1 Select File, Close Version Control File List.

Rapid SQL closes the VC Files Tab.

For more information, see:

[Working with Files in Version Control](#)

[Opening the VC Files Tab](#)

## VERSION CONTROL FUNCTIONALITY - OPEN

The Open functionality opens the selected file(s) with the application registered for the type(s) of the selected file(s).

For more information, see [Working with Files in Version Control](#).

## VERSION CONTROL FUNCTIONALITY - DELETE

The Delete functionality deletes the local copy of the selected item from the tree. To remove a file from version control, see [Remove From Version Control](#).

For more information, see [Working with Files in Version Control](#).

## VERSION CONTROL FUNCTIONALITY - SORT

The Sort functionality sorts the tree items alphabetically.

For more information, see [Working with Files in Version Control](#).

## VERSION CONTROL FUNCTIONALITY - GET LATEST VERSION

The Get Latest Version functionality lets you access the latest version of a file for viewing only. The Get functionality creates a local copy of the most current version of a project file in your working folder. The file is read-only, so any modifications cannot be saved.

Before working with a project, you should perform a Get on the entire project to ensure that you are working with the latest copy of the project. You should also perform a project-level, recursive Get at intervals to ensure that you have the latest version of files, in the event that they have been altered by others working on the same project.

The following table describes the option and functionality on the Get from Version Control dialog box:

Option	Description
Files to Get	Lets you select file(s) to get latest version of.
Advanced	Click to open the Advanced Get Options dialog box.

**TIP:** You can specify the file directory in the Version Control option of the Options Editor. For details, see [Version Control Options](#).

### Getting Latest Version of a Project

To get the latest version of project, do the following:

- 1 Click File and then Open Project.

Rapid SQL opens the Open Project dialog box.

- 2 Type the project name or select the project.

- 3 Click Open.

Rapid SQL opens the project on the Projects Tab.

- 4 Click the project or target files.

- 5 On the *Project* menu, click Version Control.

OR

On the *Project* toolbar, click Version Control.

OR

On the Project Tab, right-click a project or file.

Rapid SQL opens the Get From Version Control dialog box.

- 6 Click *Get Latest Version*.

Rapid SQL opens the Get From Version Control dialog box.

- 7 In the *Files to Get* box, click the project or files.

- 8 For advanced options, click Advanced.
- 9 Click OK.

Rapid SQL writes the most current version of the file to your working directory.

### Getting Latest Version of a File

To get the latest version of a file, do the following:

- 1 On the VC Tab, right-click the target file(s) and select Get Latest Version.

Rapid SQL opens the Get From Version Control dialog box.

- 2 Click *Get Latest Version*.

Rapid SQL opens the Get From Version Control dialog box.

- 3 In the *Files to Get* box, select the file(s).

- 4 For advanced options, click Advanced.

- 5 Click OK.

Rapid SQL writes the most current version of the file to the VC Tab.

For more information, see [Using Version Control](#).

## VERSION CONTROL FUNCTIONALITY - CHECK OUT

The Check Out functionality retrieves a copy of one or more selected files and creates a writable working file copy in the working directory. You must perform a Check Out to edit any file that has been placed under version control.

You can check out a single file, multiple files at once or an entire project. Rapid SQL displays a red check mark over the file icon to indicate that the file has been checked out and is writable. This does not prevent other users from performing a Get or a Check Out on the file unless you are using the exclusive Check Out feature found in MKS Source Integrity and Merant/Intersolv PVCS.

The following table describes the options and functionality on the Check Out File(s) dialog box:

Option	Description
Files to be Checked Out	Displays list of files that are eligible for check out. A black mark indicates that another user has the file(s) checked out.
Advanced	Click to open the Advanced Check Out Options dialog box.

**TIP:** You can specify the file directory in the Version Control Working Directory option of the Options Editor. For details, see [Version Control Options](#).

## Checking Out a Project

To check-out a project, do the following:

- 1 Click File and then Open Project.  
Rapid SQL opens the Open Project dialog box.
- 2 Type the project name or select the project.
- 3 Click Open.  
Rapid SQL opens the project on the Projects Tab.
- 4 Select the target project.
- 5 On the *Project* menu, click Version Control.  
OR  
On the *Project* toolbar, click Check Out.  
OR  
On the Project Tab, right-click a project.
- 6 Select *Check Out*.  
Rapid SQL opens the Check Out dialog box.
- 7 In the *Files to Be Checked Out* box select the project.
- 8 For advanced options, click Advanced.
- 9 Click OK.

Rapid SQL checks out the project or files from version control and writes the most current version of the file to your working directory.

## Check Out a File

To check out a file, do the following:

- 1 On the VC Tab, right-click the target file(s) and select Check Out.  
Rapid SQL opens the Check Out File(s) dialog box.
- 2 In the *Check Out File(s) dialog* box select the files.
- 3 For advanced options, click Advanced.
- 4 Click OK.

Rapid SQL checks out the file(s) from version control and writes the most current version of the file to the VC Tab.

For more information, see [Using Version Control](#).

## VERSION CONTROL FUNCTIONALITY - CHECK IN

After editing your files, you must Check In the revised file in order save the changes you made to the file in a project. The Check In functionality stores the new version of the updated file in the current project. The Check In functionality is only available if you have Checked Out a file. You have the option to Check In an entire project or individual files.

**TIP:** You can specify the file directory in the Version Control Working Directory option of the Options Editor. For details, see [Version Control Options](#).

The table below describes the options and functionality on the Check In dialog box:

Option	Description
Files to be checked in	Lets you specify the file(s) to check in.
Keep checked out	Adds latest version(s) of file(s) to source control but keeps the file(s) checked out.
Comment	OPTIONAL: Lets you type an optional comment.

### Checking In a Project

To check-in a project, do the following:

- 1 Click File and then Open Project.  
Rapid SQL opens the Open Project dialog box.
- 2 Type the project name or select the project.
- 3 Click Open.  
Rapid SQL opens the project on the Projects Tab.
- 4 On the *Project* menu, click Version Control, and then Check In.  
OR  
On the *Version Control* toolbar, click Check In.  
OR  
On the Project Tab, right-click a project or file, and then Check In.  
Rapid SQL opens the Check In dialog box.
- 5 In the *Files to Be Checked In* box click the project or files.
- 6 To update the version control copy but keep the project or files checked out so that you can continue working, select the *Keep Checked Out* check box.
- 7 To remove the file from the working directory and from the \*.xml file, and from the VC Files Tab, select Remove Local Copy.
- 8 OPTIONAL: In the *Comment* text box type a description of the changes.

- 9 Click *OK*.

Rapid SQL saves the modified project into version control.

### Check In a File

To check in a file, do the following:

- 1 On the VC Tab, right-click the target file(s) and select Check In.  
Rapid SQL opens the Check In File(s) dialog box.
- 2 In the *Check In File(s) dialog* box select the files.
- 3 To update the version control copy but keep the files checked out so that you can continue working, select the *Keep Checked Out* check box.
- 4 To remove the file from the working directory and from the \*.evc file, and from the VC Files Tab, select Remove Local Copy.
- 5 OPTIONAL: In the *Comment* text box type a description of the changes.
- 6 Click *OK*.

Rapid SQL saves the modified file(s) into version control.

For more information, see [Using Version Control](#).

## VERSION CONTROL FUNCTIONALITY - UNDO CHECK OUT

If you decide that you do not want to save any revisions you have made to a checked out file, you can undo the procedure that releases the lock placed on the project or file. However, you do not have the option of deleting the local copy of the file.

**TIP:** You can specify the file directory in the Version Control Working Directory option of the Options Editor. For details, see [Version Control Options](#).

The table below describes the options and functionality on the Undo Check Out dialog box.

Option	Description
Cancel the checkout for the following files	Lets you specify the files to undo checkout.
Advanced	Opens the Undo Check Out Advanced Options dialog box that lets you leave, replace, or apply the default action to your local copy.

## Undoing Checkout for a Project

To undo a checkout, do the following:

- 1 Click File and then Open Project.

Rapid SQL opens the Open Project dialog box.

- 2 Type the project name or select the project.

- 3 Click Open.

Rapid SQL opens the project on the Projects Tab.

- 4 Click the project.

- 5 On the *Project* menu, click Version Control.

OR

On the *Project* toolbar, click Undo Check Out.

OR

On the Project Tab, right-click a project or file.

- 6 Click *Undo Check Out*.

Rapid SQL opens the Undo Check Out dialog box.

- 7 In the *Cancel the check out for the following files* box of the *Undo Check Out* dialog box, click the project.

- 8 Click OK.

Rapid SQL undoes the check out the project from version control.

**NOTE:** If you Undo Check Out, you lose any changes you have made to the local copy of your project.

## Undoing Checkout for a File

To undo checkout for a file, do the following:

- 1 On the VC Tab, right-click the target file(s) and select Undo Checkout.

Rapid SQL opens the Undo Checkout dialog box.

- 2 In the *Cancel the check out for the following files* box of the *Undo Check Out* dialog box, click the project.

- 3 Click OK.

Rapid SQL undoes the check out the file from version control.

**NOTE:** If you Undo Check Out, you lose any changes you have made to the local copy of your file(s).



For more information, see [Using Version Control](#).

## VERSION CONTROL FUNCTIONALITY - SHOW HISTORY

The Show History functionality lets you view the history of version control files.

### Showing History for a Project

To show history, do the following:

- 1 Click File and then Open Project.  
Rapid SQL opens the Open Project dialog box.
- 2 Type the project name or select the project.
- 3 Click Open.  
Rapid SQL opens the project on the Projects Tab.
- 4 Click the target file.
- 5 On the *Project* menu, click Version Control, and then select Show History.

OR

On the Project Tab, right-click the a file, and then select Show History.

Rapid SQL opens the History Options dialog box.

**NOTE:** The History dialog box depends on your version control system.

### Showing History for a File

To show history, do the following:

- 1 On the VC Tab, right-click the target file(s) and select Show History.  
Rapid SQL opens the History Options dialog box.

**NOTE:** The History dialog box depends on your version control system.

For more information, see [Using Version Control](#).

## VERSION CONTROL FUNCTIONALITY - SHOW DIFFERENCES

The Show Differences functionality lets you view any differences between the current files in your working folder and the master files in the version control database. You cannot make changes to the files from this dialog box because it is used for display purposes only.

## Viewing Project Differences

To view project differences, do the following:

- 1 Click File and then Open Project.

Rapid SQL opens the Open Project dialog box.

- 2 Type the project name or select the project.

- 3 Click Open.

Rapid SQL opens the project on the Projects Tab.

- 4 Click the target file.

- 5 On the *Project* menu, click Version Control.

OR

On the Project Tab, right-click a file.

- 6 Click *Show Differences*.

Rapid SQL opens the Differences dialog box.

If the file in your working directory is the same as the one in the project, a message tells you they are identical. If there are differences, the Differences dialog box from your version control system opens and displays the two versions of the file side-by-side, highlighting any differences.

- 7 In the *Differences* dialog box, you can maneuver through the files by using the Up and Down arrows.

- 8 To set *Diff Options*, click the *Options* button.

**NOTE:** The Differences dialog box and Options depend on your version control system.

## Showing History for a File

To show history, do the following:

- 1 On the VC Tab, right-click the target file(s) and select Show History.

Rapid SQL opens the Difference Options dialog box.

- 2 In the *Difference Options* dialog box, you can maneuver through the files by using the Up and Down arrows.

**NOTE:** The Difference Options dialog box and options depend on your version control system.

For more information, see [Using Version Control](#).

## VERSION CONTROL FUNCTIONALITY - REMOVE FROM VERSION CONTROL

The Remove from Version Control functionality lets you remove entire projects or files from version control.

**NOTE:** This functionality does not destroy the file permanently from the source control system or remove the local copies of the file(s).

### Removing a Project from Version Control

To remove a project from Version Control, do the following:

- 1 Click File and then Open Project.

Rapid SQL opens the Open Project dialog box.

- 2 Type the project name or select the project.

- 3 Click Open.

Rapid SQL opens the project on the Projects Tab.

- 4 On the *Project* menu, click Version Control, and then click Remove from Version Control.

OR

On the Project Tab, right-click a project file, and then click Remove from Version Control.

Rapid SQL opens the Remove File(s) dialog box.

- 5 To permanently destroy the project select the check box, click *OK*.

Rapid SQL removes the project from version control. The project remains active in Rapid SQL.

### Removing a File from Version Control

The table below describes the options and functionality of the Remove Files dialog box:

Option	Description
Files to be removed	Lets you specify the files to be removed from version control. NOTE: Does not delete the file(s) in version control.

To remove a file from Version Control, do the following:

- 1 On the VC Tab, right-click the target file(s) and select Remove from Version Control.

Rapid SQL opens the Remove File(s) dialog box.

- 2 Specify the file(s).

- 3 Click *OK*.

Rapid SQL removes the file(s) from version control. The project remains active in Rapid SQL.

For more information, see [Working with Files in Version Control](#).

## VERSION CONTROL PROPERTIES

The Version Control File Properties dialog box displays general information and check out status, links, and paths.

### Viewing Version Control Properties for a Project

To view the version control properties, do the following:

- 1 Click File and then Open Project.

Rapid SQL opens the Open Project dialog box.

- 2 Type the project name or select the project.

- 3 Click Open.

Rapid SQL opens the project on the Projects Tab.

- 4 Click the target file.

- 5 On the *Project* menu, click Version Control, and then Version Control Properties.

OR

On the Project Tab, right-click the a file, and then select Version Control Properties.

Rapid SQL opens the Version Control Properties dialog box.

- 6 Review properties.

- 7 Click Close.

### Viewing Version Control Properties for a File

To view the version control properties, do the following:

- 1 On the VC Tab, right-click the target file(s) and select Version Control Properties.

Rapid SQL opens the Version Control Properties dialog box.

- 2 Review properties.

- 3 Click Close.

For more information, see [Using Version Control](#).

## VERSION CONTROL FUNCTIONALITY - EXPAND ALL

The Expand All functionality expands all tree items under the selected item(s).

For more information, see [Working with Files in Version Control](#).

#### VERSION CONTROL FUNCTIONALITY - COLLAPSE ALL

The Collapse All functionality collapses all tree items under a selected item(s).

For more information, see [Working with Files in Version Control](#).

#### VERSION CONTROL FUNCTIONALITY - REFRESH

The Refresh functionality obtains the current version control status for the file(s).

For more information, see [Working with Files in Version Control](#).

#### VERSION CONTROL FUNCTIONALITY - CLOSE FILES LIST

The Close Files List functionality closes the list of files.

For more information, see [Working with Files in Version Control](#).



# TOOLS

Rapid SQL incorporates a number of powerful tools to help you create, edit and manage your development environment. You can use Tools to:

- Conduct extensive database searches across multiple databases.
- Execute scripts or files across multiple databases.
- Schedule tasks.
- Identify differences in files or objects.
- Graphically build complex queries.
- Auto-generate complete procedures and packages.
- Code and test basic macros.
- Administer your ODBC data sources.

Tools is divided into sections. The table below describes each section:

Section	Description
<a href="#">Find in Files</a>	This section describes the Find in Files dialog box that lets you find a phrase or character in your files.
<a href="#">Database Search</a>	This section describes the powerful database search utility that helps you to find instances of a string across multiple databases.
<a href="#">Script Execution Facility</a>	This section describes the Rapid SQL Script Execution Facility, a stand-alone utility that establishes multiple threads and database connections letting you simultaneously execute SQL statements against multiple Oracle, Sybase Adaptive Server, Microsoft SQL Server, and IBM DB2 LUW for Open Systems datasources.
<a href="#">File Execution Facility</a>	This section describes the Rapid SQL File Execution Facility, a stand-alone utility that establishes multiple threads and database connections that lets you execute parallel queries and ANSI SQL files against multiple, cross-platform datasources.
<a href="#">Scheduling</a>	The Rapid SQL scheduling programs and utilities let you schedule and execute jobs on local datasources anytime.
<a href="#">Visual Difference</a>	Rapid SQL lets you compare two files or database objects. Using the Visual Difference Utility, you can easily synchronize and analyze database objects or files across multiple database platforms.
<a href="#">Query Builder</a>	This section describes Query Builder, a tool that lets you construct, structure, and manipulate up to five different types of queries simultaneously.

Section	Description
<a href="#">Data Editor</a>	This section describes the Data Editor to edit your tables in real-time. The Data Editor supports all editable datatypes and is an alternative way to add, edit, or delete data from your tables.
<a href="#">Code Generation Facility</a>	This section describes the Code Generation Facility that offers a quick way to generate DML statements for tables and views.
<a href="#">Import Data</a>	This section describes the Import Data Wizard that lets you create insert statements based on external files including Excel spreadsheets and text files.
<a href="#">Embarcadero Products</a>	The Tools menu lists all installed Embarcadero Technologies products. This lets you toggle to or start another Embarcadero product.
<a href="#">Code Workbench</a>	This section describes the Code Workbench that lets you enable table column look-up and population in the ISQL window, define auto replacement expressions that can be used to quickly insert commonly used blocks of SQL syntax or commands in any open window and to import and export Code Workbench specific settings for client sharing purposes.
<a href="#">Code Analyst</a>	This section describes the Code Analyst that lets you capture response time metrics on stored procedures. You can select one or multiple objects, execute them, view the results, and save those results for later viewing or comparing.

## FIND IN FILES

The Find in Files dialog box lets you find a phrase or character in your files.

### Completing the Find in Files Dialog Box

- 1 On the *Tools* menu, click *Find in Files*.

OR

On the *Tools* toolbar, click Find in Files.

Rapid SQL opens the Find in Files dialog box.



The table below describes the options and functionality on the Find in Files dialog box.:

Option	Description
Find what	Specifies the character(s) or phrase you want to find. Use the browse arrow button next to the textbox to choose options from a pop-up list.
In files/file types	Specifies the files in which to search for the character(s) or phrase. Either enter the filename(s) in the drop-down box, or click the arrow to choose a file type.
In folder	Specifies the directory where the file(s) is located. Click the browse button to view your Windows Explorer.
Match whole word only	Specifies the application to find only the entire phrase.
Match case	Specifies the application to find only the specified phrase in the case you have entered.
Regular Expression	Tells the application whether the specified character(s) is a regular expression.
Look in subfolders	Specifies the application to search the file(s) any folders located within the specified folder.
Output to Pane 2	Specifies the application to display the results in another window.

**NOTE:** You can also use the Find feature to locate a phrase or character in an ISQL window.

## DATABASE SEARCH

The powerful [database search utility](#) helps you to find instances of a string across multiple databases.

### Starting the Database Search Wizard

- 1 On the *Tools* menu, click *Database Search*.

OR

On the *Tools* toolbar, click Database Search.

Rapid SQL opens the [first panel of the Database Search Wizard](#).

## DATABASE SEARCH WIZARD - PANEL 1

The first panel of the Database Search Wizard lets you specify the owner.

The table below describes the options and functionality on the first panel of the Database Search Wizard:

Option	Description
Select the owner(s) whose objects you would like to search	Lets you expand the nodes, select the target owner, and then click the right arrow button to include the target owner.

- 1 Click Next.

Rapid SQL opens the next panel of the wizard.

## DATABASE SEARCH WIZARD - PANEL 2

The second panel of the Database Search Wizard lets you specify the search criteria.

The table below describes the options and functionality on the second panel of the Database Search Wizard:

Option	Description
Search Database For	Lets you enter the search string.  Strings can also be searched for using DB2-standard wildcards. And for multiple string searches, separate each string with a vertical bar (for example, DEPARTMENTS wage_cap status). When searching for strings that already contain vertical bars, enclose each string in double quotation marks.
Match Case	Select Yes to make the search case sensitive.  NOTE: IBM DB2 LUW for OS/390 and Microsoft SQL Server 7.0 or before searches are always case insensitive.
Search DDL of these Objects	In the grid, select the target object check boxes.  NOTE: Event Monitors are available for IBM DB2 LUW for Linux, Unix, and Windows only.

- 1 Click Execute to start the operation.

Rapid SQL displays a progress dialog box while the search runs. When the search completes, Rapid SQL opens the [Database Search Window](#).

## DATABASE SEARCH RESULTS

Rapid SQL displays Database Search operation results in a Database Search Window, listing all of the objects containing the search string in the left pane. You can browse instances of the search string by selecting different objects in the tree. The DDL of the objects displays in the right pane and the search string is highlighted.

The table below describes the buttons on the Database Search Window toolbar:

Button	Description
Search	Opens the first panel of the <a href="#">Database Search Wizard</a> .
Criteria	Opens the Search Criteria dialog box.
Open	Opens the editor for the target object.
Extract	Lets you extract the target object.
Print	Lets you print the target object SQL.

## SCRIPT EXECUTION FACILITY

Rapid SQL's Script Execution Facility is a stand-alone utility that establishes multiple threads and database connections letting you simultaneously execute SQL statements against multiple Oracle, Sybase Adaptive Server, Microsoft SQL Server, and IBM DB2 LUW for Linux, Unix, and Windows datasources. After completing a scheduled job, Rapid SQL generates a results report that lists errors, verifies script execution, and details the output of the job. The Script Execution Facility also works in conjunction with Rapid SQL's scheduling facilities, letting you schedule script execution jobs. When used in conjunction with a scheduler, Rapid SQL can automatically send the results report to any e-mail or network recipients. The Script Execution Facility is a tabbed dialog box where you set the parameters and options for the script execution. In the Script Execution Facility dialog box you can:

- Type or paste the target SQL script.
- Specify the datasources against which to execute the script.
- Specify the output mode for the results report.
- Open one of the Rapid SQL scheduling programs to schedule the script execution.
- Specify execution notification e-mail and Net Send addresses.

For more information, see [Completing the Script/File Execution Facility](#).

## FILE EXECUTION FACILITY

Rapid SQL's File Execution Facility is a stand-alone utility that establishes multiple threads and database connections that lets you execute parallel queries and ANSI SQL files against multiple, cross-platform datasources. The Script Execution Facility also works in conjunction with Rapid SQL's scheduling facilities, letting you schedule script execution jobs. After completing a scheduled job, Rapid SQL generates a results report that lists errors, verifies execution, and details the output of the job. When used in conjunction with a scheduler, Rapid SQL can automatically send the results report to any e-mail or network recipients. The File Execution Facility is a tabbed dialog box where you set the parameters and options for the file execution. In the File Execution Facility dialog box you can:

- Specify the target files or ANSI SQL scripts.
- Specify the datasources against which to execute the files.
- Specify the output mode for the results report.
- Open one of the Rapid SQL scheduling programs to schedule the file execution.
- Specify execution notification e-mail and Net Send addresses.

For more information, see [Completing the Script/File Execution Facility](#).

## COMPLETING THE SCRIPT/FILE EXECUTION FACILITY

Rapid SQL lets you run parallel queries against multiple datasources with the File Execution Facility.

- 1 On the *Tools* menu, click *Script Execution Facility* or *File Execution Facility*.

OR

On the *Tools* toolbar, click *Script Execution Facility* or *File Execution Facility*.

Rapid SQL opens the Script or File Execution Facility dialog box.

- 2 Complete the [Script](#) Tab ([Script Execution Facility](#))
- 3 Complete the [Files](#) Tab ([File Execution Facility](#))
- 4 Complete the [Target](#) Tab.
- 5 Complete the [Output](#) Tab.
- 6 Complete the [Notify](#) Tab.

For more information, see:

[File Execution Facility](#)

[Script Execution Facility](#)

## SCRIPT EXECUTION FACILITY - SCRIPT TAB

The table below describes the options and functionality on the Script Tab of the File/Script Execution Facility:

Option	Description
Script box	Lets you type or paste a script.

For more information, see:

[File Execution Facility](#)

[Script Execution Facility](#)

[Completing the Script/File Execution Facility](#)

## FILE EXECUTION FACILITY - FILES TAB

The table below describes the options and functionality on the Files Tab of the File Execution Facility:

Option	Description
Show Full File Paths	Select to display the full path. Deselect to display only the file name.
File Name	Displays the file names.
Add	Click to open the Select Files dialog box.
Remove	Click to remove the selected file.
View	Opens the View File dialog box.
Up	Click to move the selected file up in the list.
Down	Click to move the selected file down in the list.

For more information, see:

[File Execution Facility](#)

[Script Execution Facility](#)

[Completing the Script/File Execution Facility](#)

## FILE/SCRIPT EXECUTION FACILITY - TARGET TAB

The table below describes the options and functionality on the Target Tab of the File/Script Execution Facility:

Option	Description
Select the Target Datasource(s) to Execute the Script Against	Only Connected Datasources - Displays only datasources that are currently connected in the Datasource grid. All DBMS Types - Displays all DBMS types in the Datasource grid.
Datasource grid	Displays the target datasource(s) to execute the script/file against. Select a datasource name. If the datasource has multiple databases, type in a database in the Database box.

For more information, see:

[File Execution Facility](#)

[Script Execution Facility](#)

[Completing the Script/File Execution Facility](#)

## FILE/SCRIPT EXECUTION FACILITY - OUTPUT TAB

The table below describes the options and functionality on the Output Tab of the File/Script Execution Facility:

Option	Description
Graphical Output	If selected, specifies a graphical output.
File Output	<p>If selected, specifies a file output.</p> <p>Directory - Type or browse to enter the full path and directory name in which you want to place the output file.</p> <p>File Type - Specifies a file type.</p> <p>Include column titles when saving - If selected, lets you save column titles.</p> <p>Open files with registered applications - If selected, opens files with registered applications.</p>

For more information, see:

[File Execution Facility](#)

[Script Execution Facility](#)

[Completing the Script/File Execution Facility](#)

## FILE/SCRIPT EXECUTION FACILITY - NOTIFY TAB

The table below describes the options and functionality on the Notify Tab of the File/Script Execution Facility:

Option	Description
Job Description	Lets you enter a job description. This description will be the subject of the notification E-mail.
E-mail address	Lets you enter E-mail addresses. Separate each E-mail address with a semicolon (;).
Net Send User Names	Lets you enter net send user names. Separate each name with a semicolon (;).

For more information, see:

[File Execution Facility](#)

[Script Execution Facility](#)

[Completing the Script/File Execution Facility](#)

## SCHEDULING

The Rapid SQL scheduling programs and utilities let database administrators schedule and execute jobs on local datasources 24-hours-a-day, 7-days-a-week.

Rapid SQL offers the following programs and utilities to let you schedule routine tasks and jobs:

- [Microsoft Task Scheduler](#)
- [ETSQLX Command Line Utility](#)

## MICROSOFT TASK SCHEDULER

Rapid SQL lets you use the Microsoft Task Scheduler to schedule jobs. The Microsoft Task Scheduler is included with various Microsoft applications. If you do not have this program on your system, the first time you attempt to schedule a job, Rapid SQL provides you with a link to the Microsoft Web site where you can download the Microsoft Task Scheduler at no cost.

**NOTE:** Rapid SQL's ETSQLX command line utility runs a scheduled job even if Rapid SQL is not running. For more information, see [ETSQLX Command Line Utility](#).

### To open the Microsoft Job Scheduler from within Rapid SQL

- 1 On the **Tools** menu, click **Scheduler**.
- 2 Use the associated online Help for assistance with using this Microsoft utility.

## ETSQLX COMMAND LINE UTILITY

Rapid SQL's ETSQLX command line utility, is a multi threaded, cross-platform, SQL scripting engine. You can use ETSQLX in conjunction with the Microsoft Task Scheduler to schedule and automate routine jobs. ETSQLX creates batch files (with the extension.cfg) containing commands to execute automated and scheduled jobs. ETSQLX creates a directory, CFG, in which it stores the.cfg files. You can run.cfg files directly from the command line.

**NOTE:** ETSQLX supports.csv,.tab,.htm, and.html formats for result reports attachments.

## VISUAL DIFFERENCE

Rapid SQL lets you compare two files or database objects. Using the Visual Difference dialog box, you can easily synchronize and analyze database objects or files across multiple database platforms. The files are displayed side by side in the Visual Difference dialog box. The Visual Difference Utility highlights any differences between two files. Viewing differences between objects and files helps you negotiate between the different phases of development as well as providing a visual aid to rapidly changing and evolving production environments.

**NOTE:** Because contents of the Visual Difference dialog box are read-only, you will not be able to modify your files or objects directly from this dialog box.

The Visual Difference dialog box is composed of two panes; the left pane displays your source object or file and the right pane shows your target object or file. The Visual Difference dialog box also contains its own toolbar which lets you:

- Search
- Navigate differences
- Set options
- Print

### Opening the Visual Difference Dialog Box

1 On the *Tools* menu, click *Visual Diff*.

OR

On the *Tools* toolbar, click Visual Diff.

Rapid SQL opens the Visual Difference dialog box.

For more information, see:

[Comparing Files](#)

[Comparing Database Objects](#)

[Navigating in the Visual Difference Dialog Box](#)

[Printing a Pane of the Visual Difference Dialog Box](#)

[Searching in the Visual Difference Dialog Box](#)

[Setting Options in the Visual Difference Dialog Box](#)

## COMPARING FILES

You can compare two files side-by-side in the Visual Difference dialog box. The file you want to compare is called the Source. The file you want to compare the first file to is the Target.



## Comparing Items

- 1 On the *Tools* menu, click *Visual Diff*.

OR

On the *Tools* toolbar, click Visual Diff.

Rapid SQL opens the Visual Difference dialog box.

- 2 On the *Visual Difference* toolbar, click the *Source* icon or click the Down arrow next to the *Source* icon and then click *File*.

Rapid SQL opens the Select the 1st File to Compare dialog box.

- 3 Click the file that you want to be the Source file.

- 4 On the *Visual Difference* toolbar, click the *Target* icon or click the Down arrow next to the *Target* icon and then click *File*.

Rapid SQL opens the Select the 2nd File to Compare dialog box.

**NOTE:** The Visual Difference Utility highlights all differences between the two files.

For more information, see [Visual Difference Utility](#).

## COMPARING DATABASE OBJECTS

The schema of database objects is automatically extracted so you can view the underlying differences between object and perform a side-by-side comparison in the Visual Difference Dialog.

### Comparing Database Objects

- 1 On the *Tools* menu, click *Visual Diff*.

OR

On the *Tools* toolbar, click Visual Diff.

Rapid SQL opens the Visual Difference dialog box.

- 2 On the *Visual Difference* toolbar, click the Down arrow next to the *Source* icon and then click *Database Object*.

Rapid SQL opens the Select the 1st Database Object to Compare dialog box.

- 3 Click the datasource and then click *OK* to connect.

- 4 Navigate through the datasource tree and double-click the database object that you want to be the *Source*.

- 5 On the *Visual Difference* toolbar, click the Down arrow next to the *Target* icon and then click *Database Object*.

Rapid SQL opens the Select the 2nd Database Object to Compare dialog box.

**NOTE:** The Visual Difference Utility highlights all differences between the two database objects.

For more information, see [Visual Difference Utility](#).

## NAVIGATING IN THE VISUAL DIFFERENCE DIALOG BOX

You can navigate through the Visual Difference dialog box using the up and down arrow buttons. You can move back and forth between highlighted differences in your compared files or database objects.

### Going To the Next Difference

- 1 From the *Visual Difference* dialog box, click down arrow to go to the next difference.

### Going To the Previous Difference

- 1 From the *Visual Difference* dialog box, click up arrow to go to the next difference.

For more information, see [Visual Difference Utility](#).

## PRINTING A PANE OF THE VISUAL DIFFERENCE DIALOG BOX

You can print each pane of the Visual Difference dialog box.

### Printing a Pane of the Visual Difference Dialog Box

- 1 Position your cursor inside the pane you want to print.
- 2 Click the *Print* icon on the *Visual Difference* toolbar.

Rapid SQL opens the Print Setup dialog box.

- 3 Click OK to print the pane.

**NOTE:** You can only print one pane of the Visual Difference dialog box at a time.

For more information, see [Visual Difference Utility](#).

## SEARCHING IN THE VISUAL DIFFERENCE DIALOG BOX

The Visual Difference dialog box lets you search for text in your files or database objects.

## Searching for Text

- 1 Place your cursor inside the pane you want to search.
- 2 Click the *Find* icon on the *Visual Difference* toolbar.  
Rapid SQL opens the Find dialog box.
- 3 In the *Find What* box, enter the search string.
- 4 To match whole words only, select the *Match whole word only* check box.
- 5 To make the search case sensitive, select the *Match case* check box.
- 6 Click *Find Next* to find the next occurrence of your search string. You can also click the *Find Next* icon on the *Visual Difference* toolbar to search for the next occurrence at a later time.

For more information, see [Visual Difference Utility](#).

## SETTING OPTIONS IN THE VISUAL DIFFERENCE DIALOG BOX

The Visual Difference dialog box lets you set display and comparison options to help you customize the dialog box to view differences in a comprehensive manner.

### Setting Options

- 1 Click the *Options* icon on the *Visual Difference* toolbar.

Rapid SQL opens the Visual Diff Options dialog box.

The following table describes the check box options:

Option	Description	Default
Display Line Numbers	Indicates that line numbers should appear in the Visual Diff dialog box.	Off
Display Hidden Characters	Indicates that hidden characters (nonprintable) should be displayed.	Off
Ignore White Space	Indicates that White Space (such as spaces, carriage returns, line feeds, and tabs) should be ignored. If this option is set on, text will be considered equivalent regardless of white space, otherwise the text will be shown as being different.	On
Ignore Hidden Characters	Indicates that hidden characters (nonprintable) should be excluded.	Off
Ignore Case	Indicates that case should not be a differentiating factor.	On

- 2 Click OK.

Rapid SQL accepts the options.

For more information, see [Visual Difference Utility](#).

## QUERY BUILDER

Query Builder is a database productivity tool that lets you construct, structure, and manipulate up to five different types of queries simultaneously. It includes a separate graphical interface that opens within your current workspace. You can run Query Builder against all Embarcadero Technologies supported database platforms.

Query Builder displays the interconnections of your queries as you work. The powerful visual components of Query Builder let you see your query grow and change to fit your needs. Query Builder eases the task of drawing data from tables by automatically creating correct SQL code as you build a statement. You can use Query Builder to create and execute SELECT statements for tables and views. You can also test queries, and easily adjust your information, before you save. Query Builder does not rely on knowledge of the underlying SQL code.

You can save and reopen queries in Query Builder. Query Builder automatically checks for changes in your tables or columns between the time you save the query and the time you reopen it.

The table below describes the types of queries available in Query Builder:

Query Type	Description
<a href="#">SELECT</a>	Create, manipulate and execute SELECT Statements for tables and views.
<a href="#">INSERT</a>	Create and manipulate INSERT Statements for tables.
<a href="#">UPDATE</a>	Create and manipulate UPDATE Statements for tables.
<a href="#">DELETE</a>	Create and manipulate DELETE Statements for tables.
<a href="#">CREATE VIEW</a>	Create and manipulate CREATE VIEW Statements for tables and views.

**NOTE:** You can execute SELECT statements directly from Query Builder. INSERT, UPDATE, DELETE, and CREATE VIEW statements must be moved to an ISQL Editor for execution.

For more information, see:

[Query Builder Design](#)

[Using Query Builder](#)

## QUERY BUILDER DESIGN

Query Builder lets you build DML statements using an intuitive, graphical interface. It offers you a powerful and flexible way to quickly create complex statements without sacrificing time manipulating SQL code. Query Builder lets you add tables or columns, create joins, and change statements within the graphic display without leaving Rapid SQL. It also lets you have multiple sessions working at the same time.

Query Builder includes many different features to assist you in building and manipulating your query:

- [Query Builder Statement Properties](#)
- [Workspace Windows](#)
- [Query Builder Explorer](#)
- [Tool Bar](#)
- [SQL Diagram Pane](#)
- [SQL Statement Pane](#)

For more information, see [Build Query](#).

## WORKSPACE WINDOWS

The Workspace Windows provide a comprehensive view of your data structure and query. The table below describes the Workspace Windows:

Pane	Description
<a href="#">Query Builder Explorer Window</a>	Includes two tabs that display selected object details: Tables/Views DML
<a href="#">SQL Diagram Pane</a>	Displays tables or views included in the current query.
<a href="#">SQL Statement Pane</a>	Displays the SQL code, and when appropriate, a Results Tab.

For more information, see [Build Query](#).

## QUERY BUILDER EXPLORER WINDOW

The Query Builder Explorer is a separate tree that exposes all the tables and views in your target database. It also displays your current query structure. The Query Builder Explorer includes two tabs that display information about the selected objects:

- [Tables/Views](#)
- [DML](#)

## Tables/Views Tab

The Tables/View Tab displays information about the selected tables or views. You can use the drop-down lists to change your table or view, and when appropriate, the owner. The table below describes each drop-down list on the Tables/Views Tab:

List	Description
First	Displays all databases for a target Microsoft SQL Server or Sybase ASE.
Second	Displays all valid owners.

**NOTE:** To change your current database, select the new database in the Explorer, and then open another Query Builder session. Query Builder prompts you to [save](#) the current session prior to opening a new session.

## DML Tab

The DML Tab displays all the basic elements of a query statement in the SQL Statement Tree. You can access any element of the current statement display and perform SQL editing from the SQL Statement Tree.

For more information, see:

[Creating a Clause Using the SQL Statement Tree](#)

[Build Query](#)

## SQL DIAGRAM PANE

The SQL Diagram Pane displays tables, views, and joins included in the current query. You can manipulate elements of your query, using the mouse functionality, in the SQL Diagram Pane. From the SQL Diagram Pane you can:

- [Add and Remove tables and views.](#)
- [Create and delete joins.](#)
- [Add and Subtract columns.](#)

All changes in the SQL diagram reflect in correct SQL code in the SQL Statement Pane.

For more information, see [Build Query](#).

## SQL STATEMENT PANE

The SQL Statement Pane displays the current query SQL code. When you run a query, Query Builder displays results to your query in the SQL Statement Pane. The SQL Statement Pane is divided into two tabs:

- [SQL Tab](#)

- [Results Tab](#)

## SQL Tab

The SQL Tab displays the query in progress. It displays each element of your query as you build it, and updates as you do edits such as selecting or deselecting columns, adding clauses, and creating joins. Rapid SQL lets you open the current statement directly into an ISQL editor or copy it to the clipboard for later use.

## Results Tab

The Results Tab displays the results of your executed query in the Results grid. To edit data, use the [Data Editor](#) application from Query Builder. When you begin building a new query, the tab title changes to Old Results until you execute the new query.

For more information, see [Build Query](#).

## QUERY BUILDER TOOL BAR

The Query Builder tool bar lets you access commonly used features.

The table below describes Query Builder tool bar functionality:

Name	Function
Copy	Copies the current SQL statement to the clipboard.
Statement Box	Displays the type of statement currently on display in the main workspace window.
Stop Execution	Stops an executing query.
Execute	Executes the current SELECT or CREATE VIEW statement. If the button is not available, the statement is not executable.
New	Adjusts to the target node in the Query Builder Explorer window.
Edit	Displays, on the DML Tab, the ORDER BY or GROUP BY dialog boxes when target node is selected.
Delete	Deletes the target object.
Auto Layout	Resets the main workspace to the auto layout mode.
Auto Join	Finds and joins, automatically, like items by name.
Statement Check	Checks query syntax.
Edit Data	Opens Data Editor.
Close	Closes the current query.

**NOTE:** Query Builder adjusts tool availability to match the current query functionality.

For more information, see [Build Query](#).

## TABLES AND VIEWS SHORTCUT MENUS

Query Builder includes a shortcut menu that lets you manipulate a table or view. The table below describes the table shortcut options:

Option	Description
Delete	Removes the table from the SQL Diagram Pane, and the SQL Statement.
Title Font	Specifies the table title font for this diagram.
Column Font	Specifies the column font for this diagram.
Background Color	Specifies the table background color for this diagram.
Select Star	Selects every column in the table.
Select None	Deselects every column in the table.
Bring to Front	Moves the table to the top layer of the diagram.
Properties	Opens the <a href="#">Table Properties dialog box</a> .

**NOTE:** Your selection applies to all selected tables and views.

For more information, see [Build Query](#).

## TABLES AND VIEWS KEYBOARD COMMANDS

Query Builder provides a number of keyboard shortcuts that let you quickly construct queries. The table below describes the keyboard commands:

Keyboard Command	Location	Description
ESCAPE	SQL Diagram Pane	Breaks off a join.
F5	Query Builder	Refreshes screen and runs Schema Change Detection. In a CREATE VIEW, this key adds the new view to the Table Tree Pane.
CTRL A	SQL Diagram Pane	Selects all tables and joins in the current diagram.
F1	Query builder and application	Obtains context sensitive Help.

For more information, see [Build Query](#).



## QUERY BUILDER DIALOG BOXES

Query Builder includes a number of dialog boxes to assist you in building and customizing your query.

Dialog Box	Description
<a href="#">Statement Properties</a>	Specifies general properties in an individual Query Builder session.
<a href="#">Table Properties</a>	Specifies column selection and alias names for a table or view.
<a href="#">Column Properties</a>	Specifies column functionality within SELECT and CREATE VIEW statements.

For more information, see [Build Query](#).

### STATEMENT PROPERTIES

The Statement Properties dialog box lets you customize properties in an individual Query Builder session. For example, you can set options to limit the number of rows returned in a query to save execution time, or turn off the auto join function to manually control all joins for an individual query. The table below describes the options and functionality of the Statement Properties dialog box.

Interface Element	Option	Description	Default
Code Generation	Generate Use Database statement	Adds a line of SQL code indicating which database or instance is used in the statement.	Selected
	Generate owner names	Adds a line of SQL code showing the table owner name as part of the query.	Selected
	Include Row Count limits	Includes the output row limit set in the Execution settings.	Selected
Execution	Max Row Count in Results Set	Sets row count limits to build and check a query without congesting server processes when a query executes.	1000 rows
General	Show Column Data types in Query Diagram	Lets Query Builder reveal the data type in each column for tables in the SQL Diagram Pane.	Not selected
	Confirm on Item delete	Lets Query Builder open a Confirm Delete dialog box when an item is deleted.  NOTE: Clearing this function can result in unexpected changes to your query diagram and statement.	Selected
	Auto populate views	Lets Query Builder automatically populate views.	Not selected
Auto Join	Require Indexes	Joins indexed columns automatically, and requires indexed columns for joins.	Selected
	Require same data type	Automatically joins columns with the same data type.	Selected
Syntax Checker	Automatic Syntax Check	Lets Query Builder check syntax every time an execute statement, refresh or copy statement begins.	Selected
	Run Automatically	Lets Query Builder automatically detect like names and data types and create joins for multiple tables.	Selected
Display	Columns Font	Lets you set the font, font style, size, and color of column fonts.	Available
	Title Font	Lets you set the font, font style, size, and color of table/view title fonts.	Available
	Table Color	Lets you set the background color of your tables in the SQL Diagram Pane.	Available

**NOTE:** If you set options while Query Builder is running, Rapid SQL displays a warning indicating that you are about to change options or properties.

For more information, see [Completing the Statement Properties Dialog Box](#).

**COMPLETING THE STATEMENT PROPERTIES DIALOG BOX**

To complete the Statement Properties dialog box, do the following:

- 1 On the Query Builder menu, click Statement Properties.

OR

In the SQL Diagram Pane, right-click, and then click Statement Properties.

- 2 Set options.

- 3 Click OK.

Query Builder saves the options.

For more information, see [Build Query](#).

**TABLE PROPERTIES**

The Tables Properties dialog box lets you set parameters for tables or views in your SQL Diagram. The table below describes the options and functionality on the Table Properties dialog box.

Option	Description
Table Alias	Creates an alias name for your table.
Show Datatypes	Shows or hides the datatype for every column in the target table.
Displayed Columns	Displays columns visible in the SQL Diagram.
Hidden Columns	Displays columns hidden in the SQL Diagram.
Hide All	Moves all non selected columns in the table to the Hidden Columns window.
Display All	Moves all columns in the table to the Displayed Columns window.
Right Arrow	Moves a target file from Displayed Columns to Hidden Columns.
Left Arrow	Moves a target file from Hidden Columns to Displayed Columns.

For more information, see [Completing the Table Properties Dialog Box](#).

**COMPLETING THE TABLE PROPERTIES DIALOG BOX**

To complete the Table Properties dialog box, do the following:

- 1 Double click the target table or view title bar.

OR

Right-click target table or view, and then click Properties.

- 2 If you only want to hide or display columns in your table, click the arrow button on the table title bar.
- 3 You can also edit view properties from the Table Properties dialog box.
- 4 Click OK.

Query Builder saves the changes.

For more information, see [Build Query](#).

## COLUMN PROPERTIES

The Column Properties dialog box lets you set properties for individual columns in your SELECT or CREATE VIEW statements. You can set aggregate functions and create an alias for an individual column.

The Column Properties dialog box is not available for INSERT, UPDATE or DELETE statements.

The table below describes the options and functions Columns Properties dialog box:

Interface Element	Description
Tables/Views	Displays all tables and views in the SQL Diagram Pane.
Aggregate	Specifies aggregate options for the target column. AVG - An average is taken for a column with an int or numeric datatype. COUNT - Returns the number of rows which contain data for the target column. MAX - Returns the highest number in a row in the column. MIN - Returns the lowest number in a row in the column. SUM - Returns the sum of the target column in all rows which contain data. This function is only operable on int or numeric datatypes.
Alias	Displays the alias name for the target column. Lets you type the name of the alias. NOTE: Query Builder displays the results of an aggregate column without a column name unless you create an alias for that column.
Available Columns	Displays all available columns in the target table or view.
Selected Columns	Displays all selected columns in the target table or view. To create an aggregate function or alias for a different column, select target column, select an aggregate function, and then type the name of the alias.
Select All	Moves all columns in the Available Columns box to the Selected Columns box.
Clear All	Moves all columns in the Selected Columns box to the Available Columns box.

Interface Element	Description
Right Arrow	Moves target column in the Available Columns box to the Selected Columns box.
Left Arrow	Moves target column in the Selected Columns box to the Available Columns box.
Select List Statement	Displays the current query.

### Completing the Column Properties Dialog Box

To complete the Column Properties dialog box, do the following:

- 1 On the SQL Statement Tree, double-click target column.
- 2 Select options.
- 3 Click OK.

For more information, see [Build Query](#).

### JOIN PROPERTIES

Query Builder lets you edit joins with the Join editor. You can edit join parameters in a SELECT, UPDATE, DELETE, and CREATE VIEW Statement.

The table below describes the options and functionality on the Join dialog box.

Option	Description
From Table Column	The primary column in the join.
To Table Column	The secondary column in the join.
Select the join relation operator	Click the target join operator. If it is not equals, the operator displays on the join in the SQL Diagram Pane.
Join Type: Inner	Click to make the join an inner join. Aggregates are only available for inner joins.
Join Type: Left Outer	Click to make the join a left outer join.
Join Type: Right Outer	Click to make the join a right outer join.

**NOTE:** For IBM DB2 LUW for Linux, Unix, and Windows servers, there is an additional join object in the SQL Statement Tree. The Join On node displays join relations between columns in IBM DB2 LUW for Linux, Unix, and Windows tables and views.

## Completing the Join Dialog Box

To complete the Join dialog box, do the following:

- 1 In the *SQL Diagram Pane*, right-click the target join, and then click Properties.

OR

In the *SQL Diagram Pane*, double-click the target join.

OR

On the *SQL Statement Tree*, expand the *Where* and *And* nodes, and then double-click the target join.

- 2 Select options.
- 3 Click OK.

For more information, see [Build Query](#).

## USING QUERY BUILDER

Query Builder provides a visual display of your queries as you construct them. You can run Query Builder against any registered datasource in Rapid SQL. Query Builder lets you build five separate types of queries simultaneously:

- [SELECT](#)
- [INSERT](#)
- [UPDATE](#)
- [DELETE](#)
- [CREATE VIEW](#)

You can execute a SELECT statement from Query Builder. To execute an INSERT, UPDATE, DELETE, and CREATE VIEW statement, copy them to an ISQL Editor. You can also copy the statements to the clipboard for later use in the ISQL Editor. Query Builder also lets you [save](#) a statement at any time so that you can open them later for editing or execution.

Rapid SQL lets you open Query Builder with multiple [tables](#) or [views](#) with the same or different owners. If you open tables or views with different owners, Query Builder displays "All Owners" in the Owner drop-down list. You can start multiple Query Builder sessions from Rapid SQL. You can use different tables and views for each query. You can also toggle back and forth among each of the queries.

You can save and reopen queries in Query Builder. Query Builder automatically checks for changes in your database or instance between the time you save the query and the time you reopen it with the [Schema Change detection](#) component.

Query Builder is integrated with [Data Editor](#) so you can edit data in real time and then continue to build your query with the new information embedded in the query.

## USING QUERY BUILDER

To use Query Builder, do the following:

- [Select an instance or database](#)
- [Select a statement.](#)
- [Select a table\(s\) or view\(s\).](#)
- [Select a column or columns.](#)

**NOTE:** You can start Query Builder directly from a table or view which automatically selects the instance or database which contains that table or view.

- 1 On the Tools menu, click Query Builder.

OR

On the Datasource Explorer, expand the Database or Schema node, click Tables, and then on the Command Menu, click Build Query.

OR

On the Tools tool bar, click Build Query.

OR

On the Explorer, expand Tables, right-click the target table(s), and then click Build Query.

OR

On the Explorer, expand Views, right-click the target view(s), and then click Build Query.

Rapid SQL opens Query Builder.

## SELECTING A DATABASE

To create an SQL statement, first select an instance or database.

**NOTE:** You can start Query Builder directly from a table or view which automatically selects the database which contains that table or view.

If you are working with Microsoft SQL Server or Sybase ASE, Query Builder provides two drop-down lists. The first drop-down list displays all available databases for the target server. The second drop-down list displays owners.

If you are working with Oracle or IBM DB2 LUW for Linux, Unix, and Windows the first drop-down list is unavailable.

**NOTE:** You can start Query Builder directly from a table or view which automatically selects the database which contains that table or view.

- 1 Start Query Builder.
- 2 In the database drop-down list, click the target instance or database.

- 3 In the owners drop-down list, select the appropriate owner.

Query Builder is ready for Statement selection.

- 4 To select different instances or databases while Query Builder is running, on the Tables/ Views Tab, in the database drop-down list, click the target instance or database.

Rapid SQL clears the current query and displays a warning prompt.

- 5 To save the current query, click Yes.

Rapid SQL opens the Save As dialog box.

- 6 To continue without saving, click No.

Rapid SQL clears the SQL Diagram Pane and SQL Statement Pane.

For more information on saving queries in Query Builder, see [Saving and Reopening Queries](#).

#### SELECTING A STATEMENT

Query Builder lets you build [SELECT](#), [INSERT](#), [UPDATE](#), [DELETE](#), and [CREATE VIEW](#) queries simultaneously.

To select a statement, do the following:

- 1 On the Query Builder tool bar, click the statement drop-down list, and then click the target statement type.

OR

In the *SQL Diagram Pane*, right-click, and then click the target statement type.

For more information, see [Using Query Builder](#).

#### SELECTING TABLES AND VIEWS

To build a query, open one or more tables or views in the [SQL Diagram Pane](#). You can use different tables or views for each type of query.

**TIP:** For multiple tables: Press SHIFT+click for adjacent tables or CTRL+click for nonadjacent tables. You can also drag the bounding line with your pointer to select multiple tables.



To select a Table or View, do the following:

- 1 In the *Tables/Views* Tab, drag the target table or view to the Diagram Pane.

OR

In the *Tables/Views* Tab, click target table or view and then, on the Query Builder tool bar, click Add.

OR

In the *Tables/Views* Tab, right-click target table or view, and then click Add.

Query Builder displays the target table(s) and view(s) in the SQL Diagram Pane.

For more information, see [Using Query Builder](#).

## SELECTING COLUMNS

You must select at least one column to build a query. Query Builder displays columns in each table in the SQL Diagram window. By default, Query Builder exposes every column in a table. You can select the columns you want to use for your query. Query Builder orders them, in your statement, in the select order.

Query Builder lets you select columns in the:

- [SQL Diagram Pane](#).
- [SQL Statement Tree](#).

You can select an individual column or all columns. Query Builder orders them, in your statement, in the select order. You can reorder columns after you set them in your diagram or statement.

For more information, see:

[Selecting Columns in the SQL Diagram Pane](#)

[Selecting Columns in the SQL Statement Tree](#)

[Selecting All Columns](#)

## SELECTING COLUMNS IN THE SQL DIAGRAM PANE

To select a column in the SQL Diagram Pane, do the following:

- 1 Select the check box to the left of the target column name.

For more information, see [Using Query Builder](#).

## SELECTING COLUMNS IN THE SQL STATEMENT TREE

Query Builder lets you select and set individual properties using the [Selected Column Properties Dialog Box](#).

For more information, see [Using Query Builder](#).

### SELECTING ALL COLUMNS

Query Builder uses columns in statements based on the order of selection. When you select all columns, Query Builder displays the columns in the order they appear in the table.

**NOTE:** Query Builder lets you select all columns in single or multiple tables.

To select all columns, do the following:

- 1 On the *Query Builder* menu, click Select Star.

OR

On the SQL Statement Tree, or in the SQL Diagram, right-click target table, or any of a group of selected tables, and then click Select Star.

For more information, see [Using Query Builder](#).

### BUILDING A QUERY

Query Builder lets you build five different queries, which you can run separately or simultaneously, depending on your needs. The table below describes these queries:

Query Type	Description
<a href="#">SELECT</a>	Lets you create, manipulate and execute SELECT Statements for tables and views.
<a href="#">INSERT</a>	Lets you create and manipulate INSERT Statements for tables.
<a href="#">UPDATE</a>	Lets you create and manipulate UPDATE Statements for tables.
<a href="#">DELETE</a>	Lets you create and manipulate DELETE Statements for tables.
<a href="#">CREATE VIEW</a>	Lets you create and manipulate CREATE VIEW Statements for tables and views.

To build a Query, do the following:

- [Select an instance or database.](#)
- [Select a Statement.](#)
- [Select your table\(s\) or view\(s\).](#)
- [Select your columns.](#)

**NOTE:** You can start Query Builder directly from a table or view which automatically selects the database which contains that table or view.

Query Builder lets you build queries that include both [tables](#) and [views](#) in the SQL Diagram Pane for [SELECT](#) and [CREATE VIEW](#) statements. For the INSERT, UPDATE, and DELETE statements, use one or the other object, but you cannot use both.

Once you make your selections, you can edit, restructure, and streamline your query. Query Builder offers many options for streamlining your queries.

For more information, see:

[Working with Tables and Views in the SQL Diagram Pane](#)

[Working with Columns in the SQL Diagram Pane](#)

[Joins](#)

[Creating a Clause using the SQL Statement Tree](#)

[Moving Tables and Columns in the SQL Statement Tree](#)

[Subqueries](#)

[Aliases](#)

## BUILDING A SELECT STATEMENT

Query Builder lets you construct and execute simple-to-complex SELECT statements using data from any table or view. You can also create and edit [joins](#) for SELECT statements. Query Builder can check your query and warn you if there are syntax errors with the [Syntax Checker](#).

To build a SELECT statement, do the following:

- 1 On the Tools menu, click Query Builder.

OR

On the Datasource Explorer, expand the Database or Schema node, click Tables, and then on the Command menu, click Build Query.

OR

On the Tools tool bar, click Query Builder.

OR

On the Explorer, expand Tables, right-click the target table(s), and then click Build Query.

OR

On the Explorer, expand Views, right-click the target view(s), and then click Build Query.

Rapid SQL opens Query Builder.

- 2 In the statement drop-down list, click SELECT.
- 3 In the *Table Tree* Pane, select target table(s) or view(s) and move them to the *SQL Diagram Pane*.
- 4 In the target table or view, click target column(s), or click *Select Star* to select every column.
- 5 To check syntax, click Check.
- 6 To copy the statement, click Copy.
- 7 To execute the statement, click Execute.

### Copying a SELECT Statement from the SQL Statement Pane

To copy any part of a statement from the SQL Statement Pane, do the following:

- 1 Open Query Builder, then begin a new SELECT statement.  
OR  
Open an existing SELECT statement.
- 2 In the SQL Statement Pane, select all, or the target portion of the statement.
- 3 On the Query Builder tool bar, click Copy.  
OR  
In the SQL Statement Pane, right-click, and then click Copy.  
  
Query Builder makes the target statement portion available on the clipboard.

For more information, see [Building a Query](#).

### BUILDING AN INSERT STATEMENT

Query Builder lets you construct and execute simple-to-complex INSERT statements using data from any table. To execute an INSERT statement, copy it to an ISQL Editor. You can also copy the statement to the clipboard for later use in the ISQL Editor. Query Builder also lets you [save](#) your statement at any time so that you can open it later for editing or execution.

### Building an INSERT Statement

To build an INSERT Statement, do the following:

- 1 On the Tools menu, click Query Builder.  
OR  
On the Datasource Explorer, expand the Database or Schema node, click Tables, and then on the Command menu, click Build Query.  
OR  
On the Tools tool bar, click Query Builder.  
OR  
On the Explorer, expand Tables, right-click the target table(s), and then click Build Query.  
OR  
On the Explorer, expand Views, right-click the target view(s), and then click Build Query.  
  
Rapid SQL opens Query Builder.
- 2 In the statement drop-down list, click INSERT.
- 3 In the *Table Tree* Pane, select target table, and move it to the *SQL Diagram* Pane.
- 4 In the target table, click target column(s).

## Copying an INSERT Statement from the SQL Statement Pane

To copy any part of a statement from the SQL Statement Pane, do the following:

- 1 Open Query Builder, then begin a new INSERT statement.  
OR  
Open an existing INSERT statement.
- 2 In the SQL Statement Pane, select all, or the target portion of the statement.
- 3 On the Query Builder tool bar, click Copy.  
OR  
In the SQL Statement Pane, right-click, and then click Copy.  
Query Builder makes the target statement portion available on the clipboard.

For more information, see [Building a Query](#).

## BUILDING AN UPDATE STATEMENT

Query Builder lets you construct and execute simple-to-complex UPDATE statement using data from any table. To execute an UPDATE statement, copy it to an ISQL Editor. You can also copy the statement to the clipboard for later use in the ISQL Editor. Query Builder also lets you [save](#) your statement at any time so that you can open it later for editing or execution.

## Building an UPDATE Statement

To build an UPDATE statement, do the following:

- 1 On the Tools menu, click Query Builder.  
OR  
On the Datasource Explorer, expand the Database or Schema node, click Tables, and then on the Command menu, click Build Query.  
OR  
On the Tools tool bar, click Query Builder.  
OR  
On the Explorer, expand Tables, right-click the target table(s), and then click Build Query.  
OR  
On the Explorer, expand Views, right-click the target view(s), and then click Build Query.  
Rapid SQL opens Query Builder.
- 2 In the statement drop-down list, click *UPDATE*.
- 3 In the *Table Tree* Pane, select target table and move it to the SQL Diagram Pane.
- 4 In the target table, click target column(s).

### Copying an UPDATE Statement from the SQL Statement Pane

To copy any part of a statement from the SQL Statement Pane, do the following:

- 1 Open Query Builder, then begin a new UPDATE statement.  
OR  
Open an existing UPDATE statement.
- 2 In the SQL Statement Pane, select all, or the target portion of the statement.
- 3 On the Query Builder tool bar, click Copy.  
OR  
In the SQL Statement Pane, right-click, and then click Copy.  
  
Query Builder makes the target statement portion available on the clipboard.

For more information, see [Building a Query](#).

### BUILDING A DELETE STATEMENT

Query Builder lets you construct DELETE statements using data from any table. Query Builder displays a Confirmation Option Message box when you create a DELETE statement. You can set the [Statement Properties](#) dialog box to display or hide this message when creating a DELETE statement.

To execute a DELETE statement, copy it to an ISQL Editor. You can also copy the statement to the clipboard for later use in the ISQL Editor. Query Builder also lets you [save](#) your statement at any time so that you can open it later for editing or execution.

### Building a DELETE Statement

To build a DELETE statement, do the following:

- 1 On the Tools menu, click Query Builder.  
OR  
On the Datasource Explorer, expand the Database or Schema node, click Tables, and then on the Command menu, click Build Query.  
OR  
On the Tools tool bar, click Query Builder.  
OR  
On the Explorer, expand Tables, right-click the target table(s), and then click Build Query.  
OR  
On the Explorer, expand Views, right-click the target view(s), and then click Build Query.  
Rapid SQL opens Query Builder.
- 2 In the statement drop-down list, click *DELETE*.

- 3 In the *Table Tree* Pane, select target table, and move it to the *SQL Diagram* Pane.

### **Copying a DELETE Statement from the SQL Statement Pane**

To copy any part of a statement from the SQL Statement Pane, do the following:

- 1 Open Query Builder, then begin a new DELETE statement.  
OR  
Open an existing DELETE statement.
- 2 In the SQL Statement Pane, select all, or the target portion of the statement.
- 3 On the Query Builder tool bar, click Copy.  
OR  
In the SQL Statement Pane, right-click, and then click Copy.  
Query Builder makes the target statement portion available on the clipboard.

For more information, see [Building a Query](#).

### **BUILDING A CREATE VIEW STATEMENT**

Query Builder lets you construct and execute simple-to-complex CREATE VIEW statements using data from any table or view. You can also copy the statement to the clipboard for later use in the ISQL Editor. Query Builder also lets you [save](#) your statement at any time so that you can open it later for editing or execution.

To build a CREATE VIEW statement, do the following:

- 1 On the Tools menu, click Query Builder.  
OR  
On the Datasource Explorer, expand the Database or Schema node, click Tables, and on the Command menu, click Build Query.  
OR  
On the Tools tool bar, click Query Builder.  
OR  
On the Explorer, expand Tables, right-click the target table(s), and then click Build Query.  
OR  
On the Explorer, expand Views, right-click the target view(s), and then click Build Query.  
Rapid SQL opens Query Builder.
- 2 In the statement drop-down list, click CREATE VIEW.
- 3 In the *Table Tree* Pane, select target table or view and move it to the *SQL Diagram* Pane.

**NOTE:** Query Builder supports multiple tables and views in a CREATE VIEW statement.

- 4 In the target table or view, click the target column(s).
- 5 To check syntax, click Check.
- 6 To copy the statement, click Copy.
- 7 To execute the CREATE VIEW Statement, click the SQL Statement Pane, and then press any key.

Query Builder opens the Edit SQL dialog box.

- 8 Click OK.

**CAUTION:** If you have used this method previously, and you selected the Please do not show me this dialog again check box, on the Edit SQL dialog box, Query Builder does not display the Edit SQL dialog box. It pastes your statement directly to the ISQL Editor.

Rapid SQL opens the ISQL Editor.

- 9 In the *ISQL Editor*, on the line, CREATE VIEW NAME AS, replace the word NAME with a name for your view.
- 10 On the tool bar, click Execute.

Rapid SQL executes the CREATE VIEW query.

- 11 To close the Editor, click Close.

Rapid SQL opens the ISQL Editor save message.

- 12 Click No.

Rapid SQL returns to Query Builder.

- 13 To add the view to the table tree, on the *Query Builder* menu, click Refresh.

OR

Press F5.

Query Builder adds the view to the Table Tree Pane.

### **Copying a CREATE VIEW Statement from the SQL Statement Pane**

To copy any part of a statement from the SQL Statement Pane, do the following:

- 1 Open Query Builder, then begin a new CREATE VIEW statement.

OR

Open an existing CREATE VIEW statement.

- 2 In the SQL Statement Pane, select all, or the target portion of the statement.



- 3 On the Query Builder tool bar, click Copy.

OR

In the SQL Statement Pane, right-click, and then click Copy.

Query Builder makes the target statement portion available on the clipboard.

For more information, see [Building a Query](#).

## WORKING WITH TABLES AND VIEWS IN THE SQL DIAGRAM PANE

Query Builder lets you organize your tables and views in the SQL Diagram Pane. You can also customize appearance, change visual aspects, and adjust layout while continuing to manufacture a query. You can resize or customize a selected table and view, or move them to the front or back of the diagram. The key symbol indicates a column that is indexed or participates in a primary key.

- [Selecting and Deselecting Tables and Views](#)
- [Moving Tables and Views](#)
- [Moving Additional Tables and Views to the SQL Diagram Pane](#)
- [Deleting Tables and Views](#)

Query Builder can automatically dictate a layout in the SQL Diagram Pane using the [Auto Layout](#) button.

### SELECTING AND DESELECTING TABLES AND VIEWS

You can select tables and views in the SQL Diagram Pane. You can make changes to more than one table or view simultaneously by selecting multiple tables or views.

To select and deselect Tables and Views, do the following:

- 1 To select a table, click the table title bar.
- 2 To select more than one table, drag the pointer to enclose all target tables with the bounding line.

Query Builder selects all target tables; none have handles.

- 3 To select all tables, in the SQL Diagram, right-click, and then click *Select All*.
- 4 Click the SQL Diagram workspace to deselect all tables.

For more information, see [Working with Tables and Views](#).

### MOVING TABLES AND VIEWS

Query Builder lets you move tables and views in the SQL Diagram Pane. It also moves selections and joins with the tables and views.

To move Tables and Views, do the following:

- 1 To move a table or view, drag the title bar to the target location.

**NOTE:** If you select more than one table or view, Query Builder moves all selected tables and views and any joins with the pointer.

For more information, see [Working with Tables and Views](#).

#### MOVING ADDITIONAL TABLES AND VIEWS TO THE SQL DIAGRAM PANE

Query Builder sets tables and views in your statement in the order that you move them to the SQL Diagram Pane. Tables and views moved into the Diagram Pane appear first in your statement, including all joins connecting that table. To change the order of tables, move them back into the Table Tree and re-select them in the order in which you would like to join them.

#### Moving Additional Tables or Views

To move additional tables or views, do the following:

- 1 Click the target table or view and drag it to the Diagram Pane.

For multiple tables or views: Use SHIFT+click for adjacent tables or views or use CTRL+click for nonadjacent tables and views.

OR

Click the target table or view, and then on the Query Builder tool bar, click Add.

OR

Right-click the target table or view, and then click *Add*.

For multiple tables or views: Use SHIFT+click for adjacent tables or views or use CTRL+click for non-adjacent tables and views.

**NOTE:** Moving a table or view to the SQL Diagram Pane is not available while a query is executing.

For more information, see [Working with Tables and Views](#).

## DELETING A TABLE OR VIEW

To delete tables from the SQL Diagram Pane, do the following:

- 1 Right-click the target table or view, and then click Delete.

OR

In the SQL Diagram, click target table or view, and then on the Query Builder tool bar, click Delete.

OR

In the SQL Diagram, right-click the target table or view, and then click Delete.

Query Builder deletes the table from the SQL Diagram, SQL Statement, and SQL Statement Tree.

For more information, see [Working with Tables and Views](#).

## WORKING WITH COLUMNS IN THE SQL DIAGRAM PANE

You can customize queries by selecting and deselecting columns in the SQL Diagram Pane. You can customize columns using the [Selected Column Properties dialog box](#).

### SELECTING AND DESELECTING COLUMNS

You can select and deselect columns in the SQL Diagram. Query Builder lets you select and deselect individual columns or all columns. Your results reflect the order of selection. You can change the order of columns after you set them in your diagram or statement.

**TIP:** You can also select, re-order and deselect columns in the [SQL Statement Tree](#).

### Selecting Individual Columns

To select individual columns, do the following:

- 1 To select a column, in the SQL Diagram, select the check box to the left of the target column name.

### Deselecting Individual Columns

To deselect individual columns, do the following:

- 1 To deselect a column, in the SQL Diagram, select the check box to the left of the target column name.

**NOTE:** When you clear the columns, Query Builder deletes the columns and any sub clauses from the SQL Statement Pane and SQL Statement Tree.

## Selecting All Columns

To select all columns, do the following:

- 1 On the *Query Builder* menu, click *Select Star*.

OR

On the SQL Statement Tree, or in the SQL Diagram, right-click target table, or any of a group of selected tables, and then click *Select Star*.

**NOTE:** Query Builder uses columns in statements based on the order of selection. When you select all columns, Query Builder displays the columns as they appear in the table.

## Deselecting All Columns

To deselect all columns, do the following:

- 1 On the *Query Builder* menu, click *Select None*.

OR

On the SQL Statement Tree, or in the SQL Diagram, right-click target table, or any of a group of selected tables, and then click *Select None*.

Query Builder adds or removes selected columns from the SQL Statement Tree and the SQL Statement Pane.

## SELECTING ALL OR DISTINCT COLUMNS

Selecting ALL or DISTINCT columns is a way to filter data in your query. Selecting ALL columns means all rows displays results in the grid regardless of duplication in non-primary key columns. The DISTINCT column function is a query process that limits duplicate data in non-primary key columns to rows with the first iteration of any identical data. For example, if there are two identical addresses for different last names, and the column with a primary key does not participate in the query, only the row with the first instance of the address displays in the results of the query.

To select ALL or DISTINCT columns, do the following:

- 1 In the Statement Tree pane, right-click the ALL or DISTINCT node, click *Properties*, and then select the ALL or DISTINCT check box.

OR

In the Statement Tree pane, double click the ALL or DISTINCT node. Query Builder toggles to the opposite function.

**NOTE:** You can change between ALL or DISTINCT at any time prior to executing or copying a query.

## JOINS

Joins let you distill the information in your database to a usable form. Query Builder lets you create, manipulate, and edit work with joins without requiring knowledge of the underlying SQL code. Query Builder lets you create any type of join for SELECT and CREATE VIEW Statements. You can create self joins for UPDATE or DELETE Statements. You cannot create joins for INSERT Statements.

Query Builder includes four types of joins. The table below describes joins and their availability in Query Builder:

Join	Statement Availability	Description
<a href="#">Inner Join</a>	SELECT, CREATE VIEW, DELETE, UPDATE	Returns data from the joined tables that match the query's join criteria and set a relation between tables or views. Inner joins return results where the join condition is true.
<a href="#">Left Outer Join</a>	SELECT, CREATE VIEW	Returns all data from the primary table and data from the joined tables that match the query's join criteria and set a join relation operator from a column in a primary table or view to a column in a secondary table or view.
<a href="#">Right Outer Join</a>	SELECT, CREATE VIEW	Returns all data from the primary table and data from the joined tables that match the query's join criteria and set a join relation operator from a column in a secondary table or view to a column in a primary table or view.
<a href="#">Self Join</a>	SELECT, CREATE VIEW	Set a relation between columns in the same table.

In the Query Builder SQL Diagram Pane, you can create, edit, and delete joins. You can edit joins in the [Join dialog box](#).

Joins are the way you can filter data in relational databases. Query Builder lets you change the types of [joins](#) between tables, views and columns. It is important that you have some knowledge of the data in your tables, and the datatypes for each column. This information helps you frame a better query, and filter your data for maximum effect.

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

### INNER JOINS

Inner joins are the most common types of joins for SELECT statements. An inner join returns information from two tables where the relation between two target columns is true for both columns.

The join operand determines the relation results, for example, if the join operand is equals, then identical data, in two columns, is the only result. If the join operand is not equals, Query Builder only returns data that is different between two columns.

For example, if you have an inner join matching territory numbers between the table `dbo.Managers` and `dbo.Clients`, running the query returns all Managers and Clients with matching territory numbers:

Query Builder displays the following results from this query with an inner join

**NOTE:** Query Builder displays results of columns in the order of selection. You can reorder columns by deselecting and selecting in the SQL Diagram Pane, the Selected Columns Properties dialog box, or the SQL Statement Tree.

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

### LEFT OUTER JOINS

Left outer joins bring back a different data set than [inner joins](#). Left outer joins retrieve all the data in columns selected from the primary table, and only matching data from the joined or secondary table.

For example, in the same pair of tables, a left inner join from `dbo.Managers` to `dbo.Clients`, where the columns `Current Territory` and `Territory` are joined, displays different results.

**NOTE:** There is one additional manager who does not have a client, but because a left outer join includes all data from selected columns in the primary table, the last entry in the illustration is displayed.

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

### RIGHT OUTER JOINS

Right outer joins return opposite results from a [left outer join](#). In a right outer join, you are asking for all the information in the secondary table's column, and the join operator's matching information from the primary table.

For example, in the same set of data we used in the left outer join example, a right outer join returns all clients from `dbo.Client`, and only managers who match territory numbers, in the joined column.

**NOTE:** The managers are the same as the first, inner join, but a right outer join returns the additional clients without matching managers.

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

### SELF JOINS

A self join is a join within a single table. Query Builder lets you return specific information from a single table using a self join.

For example, in our example table, there is a column for the number of clients and another column with the goal client total for a territory.

A self join can ascertain which managers are reaching their quota. Notice that the join relation operator in the example is greater than or equal to, which shows managers exceeding quota as well.

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

#### ADDING AND DELETING A JOIN IN THE SQL DIAGRAM PANE

Query Builder lets you add and delete joins. This method adds a [WHERE](#) clause in your query. You can join different [tables](#) and or [views](#) in a [SELECT](#) or CREATE VIEW statement.

#### Adding a Join

To add a Join, do the following:

- 1 In the SQL Diagram Pane, drag the target column to the second column.

Query Builder displays both a line joining the two columns in the SQL Diagram Pane and the corresponding SQL code in the SQL Statement Pane.

#### Removing a Join

Query Builder lets you remove joins from your query. Query Builder automatically deletes joins from the query in the SQL Statement Pane, when you remove them from the SQL Diagram Pane.

To remove a join, do the following:

- 1 Click the target join, and then on the *Query Builder* tool bar, click Delete.

OR

Right-click the target join, and then click Delete.

Query Builder deletes the Join.

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

#### EDITING JOINS

Query Builder lets you edit joins with the Join editor. You can edit join parameters in a SELECT, UPDATE, DELETE, and CREATE VIEW Statement.

The table below describes the options in the Join dialog box:

Option	Description
From Table Column	The primary column in the join.
To Table Column	The secondary column in the join.
Select the join relation operator	Click the target join operator. If it is not equals, the operator displays on the join in the SQL Diagram Pane.
Join Type: Inner	Click to make the join an inner join. Aggregates are only available for inner joins.

Option	Description
Join Type: Left Outer	Click to make the join a left outer join.
Join Type: Right Outer	Click to make the join a right outer join.

### Completing the Join Dialog Box

- 1 In the *SQL Diagram Pane*, right-click the target join, and then click Properties.

OR

In the *SQL Diagram Pane*, double-click the target join.

OR

On the *SQL Statement Tree*, expand the *Where* and *And* nodes, and then double-click the target join.

Query Builder opens the Join dialog box.

**NOTE:** For IBM DB2 LUW for Linux, Unix, and Windows servers, there is an additional join object in the SQL Statement Tree. The Join On node displays join relations between columns in IBM DB2 LUW for Linux, Unix, and Windows tables and views.

### Changing a Join Color

Query Builder lets you change the color at a join in the SQL Diagram Pane. Complex statements using many tables and multiple joins can be easier to view if joins have different colors.

To change the color of a join, do the following:

- 1 Right-click the target join, and then click Color.

Query Builder opens the Color dialog box.

- 2 In the Basic colors grid, click a target color

OR

Click Define Custom Colors, then create a custom color.

**NOTE:** Query Builder lets you save custom colors for the current color. Click Add to Custom Color to have the option of using that color for your queries.

- 3 Click OK.

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



## AUTO LAYOUT

The Auto Layout function displays [tables](#) and [views](#) in the [SQL Diagram Pane](#). It makes the best use of the available area in the SQL Diagram Pane by placing your tables and views in the most efficient manner. If the [automatic join](#) function is on, Query Builder displays all joins between columns in your diagram. Query Builder lets you run the automatic layout function any time you have tables or views in the SQL Diagram Pane.

### Using Auto Layout

To use Auto Layout, do the following:

- 1 On the Query Builder menu, click Auto Layout.
- OR
- On the Query Builder tool bar, click Auto Layout.
- OR
- In the SQL Diagram Pane, right-click, and then click Auto Layout.
- Query Builder organizes your tables in the SQL Diagram Pane.

## AUTO JOINS

Query Builder includes an automatic join function that displays joins between selected [tables](#) and [views](#) in the [SQL Diagram Pane](#). The Auto Join function seeks columns with the same name and data type. You can set global automatic join parameters in the Rapid SQL Options Editor. You can use the [Statement Properties Editor](#) to set local join parameters for the current Query Builder session without changing the global parameters.

### Using Auto Join

To use Auto Join, do the following:

- 1 On the Query Builder menu, click Auto Join.
- OR
- On the Query Builder tool bar, click Auto Join.
- OR
- In the SQL Diagram Pane, right-click, and then click Auto Join.
- Query Builder joins columns in the SQL Diagram Pane.

## CREATING A CLAUSE USING THE SQL STATEMENT TREE

Query Builder lets you build more detailed WHERE, ORDER BY, GROUP BY, and HAVING clauses using the SQL Statement Tree. Query Builder lets you add clauses to SELECT, UPDATE, DELETE, and CREATE VIEW statements.

**NOTE:** Query Builder does not support clauses for INSERT statements.

The table below describes these clauses:

Clause	Description
<a href="#">WHERE</a>	Limits rows in the query.
<a href="#">ORDER BY</a>	Orders the results of the query to a target column.
<a href="#">GROUP BY</a>	Groups target columns in the query.
<a href="#">HAVING</a>	Filters out groups of data.

## CREATING A WHERE CLAUSE

Query Builder lets you create a WHERE clause from the SQL Statement Tree which automatically displays in your query.

**NOTE:** Any additional WHERE clauses are displayed as [HAVING](#) clauses.

The table below describes the options and functionality on the Where dialog box.

Option	Description
Operand (Left)	Lets you click the target column for the first part of your WHERE clause.  NOTE: Query Builder lists every column in all tables in the SQL Diagram in the Operand lists.
Operator	Lets you select the target operator.
Operand (Right)	Lets you click the target column for the second part of your WHERE clause.  Query Builder automatically writes the query language in the Statement option box.

**NOTE:** Query Builder does not display clause phrases created from the SQL Statement Tree in the SQL Diagram Pane.

## Creating a WHERE Clause

To Create a WHERE clause, do the following:

- 1 Click the WHERE node, and then on the Query Builder tool bar, click New.

OR

Right-click the WHERE node, and then click New.

For more information, see [Creating a Clause using the SQL Statement Tree](#).

#### DELETING A WHERE CLAUSE

To delete a WHERE clause, do the following:

- 1 Expand the AND node, and then on the Query Builder tool bar, click Delete.

OR

Expand the AND node, right-click target column and then click Delete.

Query Builder deletes the target clause and removes it from the SQL Statement Pane.

For more information, see [Creating a Clause using the SQL Statement Tree](#).

#### CREATING AN AND CLAUSE IN A WHERE CLAUSE

Query Builder lets you add an AND clause from the SQL Statement Tree which automatically displays in your query.

The table below describes the options and functionality on the Where dialog box.

Option	Description
Operand (Left)	Lets you click the target column for the first part of your WHERE clause.
Operator	Lets you select the target operator.
Operand (Right)	Lets you click the target column for the second part of your WHERE clause. Query Builder automatically writes the query language in the Statement option box.
New Button	Click to clear your selections but remain in the Where dialog box. Query Builder adds another AND clause to your query.

To open the Where dialog box, do the following:

- 1 Click the AND node, and then on the Query Builder tool bar, click New.

OR

Expand the WHERE node, right-click the AND node, and then click New.

For more information, see [Creating a Clause using the SQL Statement Tree](#).

## DELETING AN AND CLAUSE

To delete an AND clause, do the following:

- 1 Expand the AND node, click target column, and then on the Query Builder tool bar, click Delete.

OR

Expand the AND node, click target column, and then on the keyboard press DELETE.

OR

Expand the AND node, right-click the target column, and then click Delete.

Query Builder deletes the target clause and removes it from the SQL Statement Pane.

For more information, see [Creating a Clause using the SQL Statement Tree](#).

## INSERTING AN AND OR OR CLAUSE

Query Builder lets you insert an AND or an OR WHERE clause from the SQL Statement Tree which automatically displays in your query. Query Builder lets you insert AND or OR clauses at any appropriate point in the SQL Statement Tree.

The table below describes the options and functionality on the Where dialog box.

Option	Description
Operand (Left)	Lets you click the target column for the first part of your WHERE clause.
Operator	Lets you select the target operator.
Operand (Right)	Lets you click the target column for the second part of your WHERE clause. Query Builder automatically writes the query language in the Statement option box.
New Button	Click to clear your selections but remain in the Where dialog box. Query Builder adds another AND clause to your query.

To insert an AND or OR Clause, do the following:

- 1 On the SQL Statement Tree, expand the WHERE node, right-click the target AND node, then click Insert, and then click And or Or.

For more information, see [Creating a Clause using the SQL Statement Tree](#).

## DELETING AN OR CLAUSE

To delete an OR clause, do the following:

- 1 Expand the OR node, and then on the Query Builder tool bar, click Delete.

OR

Expand the OR node, right-click the target column and then click Delete.

Query Builder deletes the target clause and removes it from the SQL Statement Pane.

For more information, see [Creating a Clause using the SQL Statement Tree](#).

## CREATING AN ORDER BY CLAUSE

Query Builder lets you create an ORDER BY clause from the SQL Statement Tree which automatically displays in your query.

The table below describes the Order By Columns dialog box.

Option	Description
Available Columns	Select target column(s) and click the right arrow. Query Builder moves target column from the Available Columns list to the Order By Columns list.  NOTE: Query Builder sorts query results based on the order that columns are placed in the ORDER BY clause.
Order	Lets you select the target sort order. ASC - Ascending DESC - Descending Query Builder displays the SQL language in the Order By Statement box.

To open the Order By Columns dialog box, do the following:

- 1 On the SQL Statement Tree, click the ORDER BY node, and then on the Query Builder tool bar, click Properties.

OR

On the SQL Statement Tree, right-click the ORDER BY node, and then click Properties.

For more information, see [Creating a Clause using the SQL Statement Tree](#).

## CHANGING THE SORT ORDER IN AN ORDER BY CLAUSE

To quickly change the sort order of a column in a query, do the following:

- 1 On the SQL Statement Tree, expand the ORDER BY node, and then double-click the target column.

OR

On the SQL Statement Tree, expand the ORDER BY node, then right-click the target column, and then click Properties.

Query Builder opens the Order dialog box.

- 2 Click the target sort order, and then click OK.

Query Builder appends the Order By clause for target column with the appropriate sort order in the SQL Statement Pane.

For more information, see [Creating a Clause using the SQL Statement Tree](#).

## DELETING AN ORDER BY CLAUSE

To delete an ORDER BY clause, do the following:

- 1 Expand the ORDER BY node, and then on the Query Builder tool bar, click Delete.

OR

Expand the ORDER BY node, right-click the target column, and then click Delete.

Query Builder deletes the target clause and removes it from the SQL Statement Pane.

For more information, see [Creating a Clause using the SQL Statement Tree](#).

## CREATING A GROUP BY CLAUSE

The table below describes the options and functionality on the Group By Columns dialog box.

Option	Description
Selected Columns	Select target column(s) and click the right arrow. Or click the Select All button. Query Builder moves target column from the Selected Columns list to the Group By Columns list.  NOTE: Query Builder sorts query results based on the order that columns are placed in the ORDER BY clause.
Clear All Button	Click to move target column from the Group By Columns list to the Selected Columns list.  Query Builder displays the SQL language in the Group By Statement window.

## Creating a GROUP BY Clause

To create a GROUP BY clause from the SQL Statement Tree which automatically displays in your query, do the following:

- 1 On the SQL Statement Tree, double-click the GROUP BY node.

OR

On the SQL Statement Tree, right-click the GROUP BY node, and then click New.

Query Builder adds all the selected columns in your table(s) to the GROUP BY node in the SQL Statement Tree, and to the appropriate location in the SQL Statement Pane.

- 2 On the GROUP BY node, double-click any column.

OR

On the GROUP BY node, click any column, then on the Query Builder menu, click New.

OR

On the GROUP BY node, right-click any column, then click Properties.

For more information, see [Creating a Clause using the SQL Statement Tree](#).

## DELETING A GROUP BY CLAUSE

To delete a GROUP BY clause, do the following:

- 1 On the SQL Statement Tree expand the GROUP BY node, and then on the Query Builder tool bar, click Delete.

OR

On the SQL Statement Tree Expand the GROUP BY node, right-click the target column, and then click Delete.

Query Builder deletes the target clause and removes it from the SQL Statement Pane.

For more information, see [Creating a Clause using the SQL Statement Tree](#).

## CREATING A HAVING CLAUSE

A HAVING clause is a type of WHERE clause. It filters additional information from your tables. Query Builder lets you create a HAVING clause from the SQL Statement Tree which automatically displays in your query. Query Builder lists every column in all tables in the SQL Diagram in the Operand lists. Query Builder displays the datatype of a column in the operand boxes.

The table below describes the options and functionality on the Having dialog box.

Option	Description
Operand (Left)	Lets you click the target column for the first part of your HAVING clause.
Operator	Lets you select the target operator.
Operand (Right)	Lets you click the target column for the second part of your HAVING clause.  Query Builder automatically writes the query language in the Statement option box.
New Button	Click to clear your selections but remain in the Having dialog box.  Query Builder adds another AND clause to your query.

**NOTE:** Query Builder does not display clause phrases created from the SQL Statement Tree in the SQL Diagram Pane.

To create a HAVING clause, do the following:

- 1 On the SQL Statement Tree, expand the *HAVING* node, and then expand the *And* node. If there is not a join listed on the *And* node, double-click *And*. If there is a join listed, use the shortcut option below.

OR

On the SQL Statement Tree, right-click the *HAVING* node, and then click New.

For more information, see [Creating a Clause using the SQL Statement Tree](#).

#### DELETING A HAVING CLAUSE

To delete a HAVING clause, do the following:

- 1 On the SQL Statement Tree expand the HAVING node, and then on the Query Builder tool bar, click Delete.

OR

On the SQL Statement Tree expand the HAVING node, right-click the target column, and then click Delete.

Query Builder deletes the target clause and removes it from the SQL Statement Pane.

For more information, see [Creating a Clause using the SQL Statement Tree](#).



### CHANGING TABLES AND COLUMNS LOCATION IN THE SQL STATEMENT TREE

Query Builder lets you move tables and columns on the SQL Statement Tree by dragging them to new locations. You can move columns from the AND and OR nodes to an AND or OR node on the WHERE and HAVING clause nodes. Query Builder changes the query in the SQL Statement Pane to match each move. Query Builder moves tables or columns you are dragging below target table or column.

To move a table or column in the SQL Statement Tree, do the following:

- 1 Expand target node, then drag the target table or column to a new location.

Query Builder makes the appropriate change in the query in the SQL Statement Pane.

**NOTE:** Query Builder lets you select multiple tables or columns.

- 2 To move a table or column to the bottom of a node, drag it to the target node.

Query Builder displays the target table or column at the bottom of target node.

For more information, see [Creating a Clause using the SQL Statement Tree](#).

### SUBQUERIES

Query Builder lets you build subqueries for SELECT and CREATE VIEW statements in the [WHERE](#) or [HAVING](#) clause. The table below describes the options available for a subquery in Query Builder:

Operand	Location	Description
EXISTS	Left operand	Specifies data that exists in a column.
NOT EXISTS	Left operand	Specifies data that does not exist in a column.
ANY	Right operand	Specifies data satisfying the operator parameters.
ALL	Right operand	Specifies data satisfying the operator parameters.
SELECT	Right operand	Specifies data satisfying the operator parameters.

The table below describes the options and functionality on the Where or Having dialog boxes.

Option	Description
Operand (Left)	Lets you click the target column for the first part of your clause.
Operator	Lets you select the target operator.
Operand (Right)	Lets you click the target column for the second part of your clause. Query Builder displays the working subquery in the Statement window.
Subquery	Paste or type the SUBQUERY statement.

To use the [WHERE](#) and [HAVING](#) dialog boxes to create subqueries, do the following:

- 1 On the SQL Statement Tree, expand the *Where* or *Having* node, and then expand the *And* node. If there is not a join listed on the *And* node, double-click *And*. If there is a join listed, use the shortcut option below.

OR

On the SQL Statement Tree, right-click the *Where* or *Having* node, and then click *New*.

For more information, see [Creating a Clause using the SQL Statement Tree](#).

## SYNTAX CHECKER

The Syntax Checker scans SQL statements for errors. You can check your syntax at any time while you are fashioning a query, or a Procedure or Function. Query Builder can automatically run a syntax check to validate your query when you are executing or copying a statement.

**NOTE:** Query Builder lets you continue with your query even if there are errors detected in the syntax.

### Using the Syntax Checker

The table below describes the possible syntax errors the Query Builder Syntax Checker tool displays, in order:

Error	Description
Does the query contain duplicate aliases?	Query Builder returns an error message when it detects duplicate aliases.
If the query has a <a href="#">HAVING</a> clause, is there a <a href="#">GROUP BY</a> clause?	Query Builder returns an error message when it detects a HAVING clause without a GROUP BY clause.
If there are aggregates, or a <a href="#">GROUP BY</a> clause, are all columns in one or the other?	Query Builder returns an error message when it detects an aggregate, or a GROUP BY clause without all columns in one or the other.
Are there <a href="#">joins</a> against non-indexed columns, or columns not participating in a primary key?	Query Builder returns a warning when it detects a join against a non-indexed column, or a column not participating in a primary key.
Are there joins between different datatypes?	Query Builder returns a warning when it detects a join between different datatypes.
Are there cross-products in the query?	Query Builder returns a warning when it detects a cross-product in the query.

## SAVING AND REOPENING QUERIES

You can save and reopen queries in Query Builder. Saving a query saves the SQL Diagram, SQL Statement, and Query Builder Explorer view. Query Builder automatically checks for changes in your database or instance between the time you save the query and the time you reopen it. Query Builder prompts you to save any time you try to close Query Builder, or any time you attempt quitting Rapid SQL.

Query Builder runs Schema Change detection any time you set a query to execute, refresh the data, or open a saved query.

Rapid SQL lets you open multiple saved queries simultaneously.

### Saving Queries

To save a query using standard Save and Save As functions, do the following:

- 1 On the File menu, click Save or Save As.

OR

On the Main tool bar, click Save As.

Rapid SQL opens the Save As dialog box.

- 2 In the File name box, type the name of the query.

**NOTE:** By default, the product appends the.qbl extension to Query Builder files. If there is more than one Query Builder session in progress when you save, the file is further appended with an integer, for example.qbl2.

**TIP:** Rapid SQL lets you save data in text (\*.txt) and XML (\*.xml) file formats.

- 3 Click OK.

Rapid SQL saves the file and closes the Save As dialog box.

### Reopening Queries

You can open a query using standard Open functions. Query Builder displays the Query Builder diagram, statement and Query Builder Explorer Pane and it checks the instance or database for schema changes.

The Query Builder Schema Change Detection component checks for:

- Renamed or dropped tables referenced in the query. Renamed tables that have been renamed are considered dropped.
- Renamed or dropped columns referenced in the query. Renamed columns are considered dropped and inserted.
- [Columns](#) added or reordered in tables referenced in the query.

If Query Builder detects a change, it opens the Schema Change Detected dialog box. The dialog box displays details of changes to your schema.

Query Builder opens an ISQL Editor with the last saved versions of the SQL statement.

## USING DATA EDITOR WITH QUERY BUILDER

SELECT statements. Rapid SQL lets you open multiple Data Editor sessions so that you can continue to change your data until you find the best match for your query.

**CAUTION:** Data Editor is a real-time editor. Changes in your data using Data Editor are permanent.

### Opening the Data Editor from Query Builder

To open the Data Editor from Query Builder, do the following:

- 1 On the Tools menu, click Query Builder.

OR

On the Datasource Explorer, expand the Database or Schema node, click Tables, and then on the Command menu, click Build Query.

OR

On the Tools tool bar, click Query Builder.

OR

On the Explorer, expand Tables, right-click the target table(s), and then click Build Query.

OR

On the Explorer, expand Views, right-click the target view(s), and then click Build Query.

Rapid SQL opens Query Builder.

- 2 Select a [database](#) or instance.

- 3 Select a [table](#).

- 4 Select a [column](#), or columns.

- 5 On the Query Builder menu, click Edit Data.

OR

On the Query Builder tool bar, click Edit Data.

Rapid SQL opens Data Editor.

For more information on using the Data Editor, see [Using Data Editor](#).

## DATA EDITOR

The Edit Data function opens the Data Editor. You can use the Data Editor to edit your tables in real-time. The Data Editor supports all editable datatypes and is an alternative way to add, edit, or delete data from your tables.

**NOTE:** You can use Data Editor within [Query Builder](#) to edit data in tables while you create SELECT statements. You can open multiple Data Editor sessions so that you can continue to change your data until you find the best match query.

The Data Editor includes a [Data Editor Filter](#) that lets you select the columns in your table that you want to edit. You must select at least one column to use the Data Editor. The Data Editor Filter is not available for the Query Builder.

For more information, see:

[Data Editor Design](#)

[Using Data Editor](#)

## DATA EDITOR DESIGN

The Data Editor includes the following components:

- [Edit Window](#)
- [ISQL Window](#)
- [Tool Bar](#)
- [Data Editor Filter](#)
- [Date/Time Format Builder](#)
- [Using Data Editor](#)

## DATA EDITOR EDIT WINDOW

Data Editor displays all the information in the target table in the Data Editor Edit Window. You can edit data directly in this window.

For more information, see:

[ISQL Window](#)

[Tool Bar](#)

[Data Editor Filter](#)

[Date/Time Format Builder](#)

[Using Data Editor](#)

## DATA EDITOR ISQL WINDOW

The Data Editor ISQL Window displays the active SQL statement, which uses the data from the target table.

When appropriate, Data Editor displays a History Tab. The History Tab displays all SQL Statements created in the current session. If there is an error, Data Editor displays an Error Tab. The Error Tab details any errors in data entry encountered during execution.

For more information, see:

[Edit Window](#)

[Tool Bar](#)

[Data Editor Filter](#)

[Date/Time Format Builder](#)

[Using Data Editor](#)

## DATA EDITOR TOOL BAR

The Data Editor tool bar lets you access commonly used features.

The table below describes the function of each Data Editor tool.

Description	Function
Stop Button.	Stops loading data to the Data Editor. Data Editor displays rows up to the stopping point.
List of options for the target table.	Displays the editing mode for the target table.
Execute SQL button	Executes the current SQL statement for the target table.
Insert Record button	Inserts new record for the target table. New records display at the end of the table. For related information, see <a href="#">Default Value Handling</a> .
Save Current Row button	Saves data in the current selected row. Data Editor prompts to save when you attempt to leave a row in Live mode. For related information, see <a href="#">Default Value Handling</a> .
Remove Data button	Removes data in target row. Data Editor displays an optional prompt.
Clear SQL Text button	Clears SQL text from the SQL Statement Pane.
Undo button	Undoes the most recent operation.
Redo button	Redoes the most recent operation.
First Record button	Moves to the first record in the target table.
Last Record button	Moves to the final record in the target table.
Filter Data button	Filters table using the target cell as the filter parameter.
Refresh button	Reloads data for target table

Description	Function
Calendar button	Sets correct format for target date/time cell. Enables the Calendar window.
Date/Time Format Builder button	Opens the <a href="#">Date/Time Format Builder dialog box</a> .
Date/Time Format Undo button.	Undoes the last date/time format display.
Date/Time Format Redo button.	Redoes the last date/time format display.
Close button	Closes and exits Data Editor.

For more information, see:

[Edit Window](#)

[ISQL Window](#)

[Data Editor Filter](#)

[Date/Time Format Builder](#)

[Using Data Editor](#)

#### DATA EDITOR FILTER

The Data Editor Filter displays the columns of a target table and the corresponding SELECT SQL Statement. You can select columns from the filter for selective data editing.

For more information, see:

[Edit Window](#)

[ISQL Window](#)

[Tool Bar](#)

[Date/Time Format Builder](#)

[Using Data Editor](#)

#### NOTES ON XML TYPES AND UNICODE DISPLAY IN THE DATA EDITOR

When working with data in the Data editor, keep the following in mind:

- XML data types are supported for IBM DB2 for Windows, Unix, and Linux, Microsoft SQL Server, and Oracle. In the Data editor, XML data types are displayed and entered as LOB content.

- Support for display of Unicode characters is provided as follows:
  - IBM DB2 for Windows, Unix, and Linux V8 and V9: **character**, **clob**, **varchar**, and **longvarchar** types
  - SQL Server 2000: for **nchar**, **nvarchar**, and **ntext** types
  - SQL Server 2005: **nchar**, **nvarchar**, **ntext**, and **nvarchar(max)** types.
  - Oracle 8i, 9i, and 10g: **NCHAR**, **NVARCHAR2** and **NCLOB** for non-Unicode UTF8 Character Set Instances and **NCHAR**, **NVARCHAR2**, **CHAR**, **VARCHAR2**, **LONG**, **NCLOB** and **CLOB** for Unicode UTF8 Character Set Instances
  - Sybase 12.5 and 15.2: **UNICHAR**, **UNIVARCHAR** and **UNITEXT** for non-Unicode UTF8 Character Set Instances and **UNICHAR**, **UNIVARCHAR**, **UNITEXT**, **NCHAR**, **NVARCHAR**, **CHAR**, **VARCHAR** and **TEXT** for Unicode UTF8 Character Set Instances

## USING DATA EDITOR

Data Editor lets you edit data in your tables with any editable datatype without leaving the parent application. Data Editor lets you use your new data immediately.

**CAUTION:** Data Editor is a real-time editor. Changes in your data using Data Editor are permanent.

The table below describes the functions and options of the Data Editor:

Option	Description
Live	Edits data one row at a time. You must execute when you leave the row.
Batch	Edits data in multiple rows before executing.

**NOTE:** You can also use the Data Editor to edit [date and time functions](#) in a table.

**NOTE:** If you make an incorrect edit in a cell, Data Editor displays the error on the Error Tab of the ISQL Editor. Data Editor does not incorporate this error(s) in data into the table. Data Editor saves any changes in data prior to the error message.

**CAUTION:** Data Editor is a real-time editor. Changes in your data using Data Editor are permanent.

For more information, see:

[Edit Window](#)

[ISQL Window](#)

[Tool Bar](#)



[Data Editor Filter](#)

[Date/Time Format Builder](#)

[Default Value Handling](#)

## EDITING DATE AND TIME FUNCTIONS

The Data Editor lets you edit date and time functions in a table. Data Editor uses a calendar tool to guarantee accurate input for date and time data. You can also change the display of date and time using the [Date/Time Format Builder](#).

For more information, see:

[Calendar Button](#)

[Date/Time Format Builder](#)

## DATE/TIME FORMAT BUILDER

The Date/Time Format Builder lets you customize your date/time display. The [Data Editor](#) uses this format to display your dates and times. You control how the Data Editor displays the dates and time by using predefined formats, or by customizing a format to fit your needs.

The Data Editor uses the default date/time format of your Operating System. If you do not make any global changes in Rapid SQL, the Date/Time Format Builder displays dates and times using the default formats of your operating system. If you make changes to dates and times in the Data Editor, Rapid SQL commits the changes in the format used by the database.

**NOTE:** The changes you make using the Date/Time Format Builder do not affect the way your database stores dates and times.

## EDITING THE DATE/TIME DISPLAY

You can edit the date/time display on a global, table, or column level. The table below describes the different ways you can edit your date/time format display:

Option	Description	Access
<a href="#">Global</a>	Lets you make global changes to the Data Editor date display.	Options Editor
<a href="#">Grid</a>	Lets you make changes to the date display of the entire Data Editor grid for that session only.	Data Editor grid
<a href="#">Column</a>	Lets you make changes to the date display of a single column in the Data Editor for that session only.	Data Editor column

**NOTE:** Date/Time formats changed on a table or column level are valid for that session only.

### Editing Date/Time Globally

You can use the Options Editor to make global changes to your date/time display in the [Data Editor](#). When you change the date/time format, using the Options Editor, the Data Editor displays all dates and times in the global manner. To change the date/time display for a particular session, see [Editing Grid Date/Time Format](#) or [Editing Column Date/Time Format](#).

To edit the date and time globally, do the following:

- 1 On the *File* menu, click *Options*.

OR

On the *Main* tool bar, click Options.

Rapid SQL opens the Options Editor.

- 2 On the Options Editor, click the Data Editor Tab.

- 3 On the Data Editor Tab, click ...

Rapid SQL opens the Date/Time Format Builder dialog box.

- 4 On the Date/Time Format Builder dialog box, click the Date/Time Format list, and then click the target predefined date/time format.

- 5 To customize the date/time format to your specifications, click Customize.

Rapid SQL opens the Date/Time Format Builder dialog box.

- 6 On the Date/Time Format Builder dialog box, select the appropriate Date/Time Format Options:

Option	Description
Date/Time Format	Displays the predefined Date/Time format.
Day Format	Lets you choose the day display.
Separator	Lets you choose the display separator between the day, month, and year.
Month Format	Lets you choose the month display.
Year Format	Lets you choose the year display.
Date Order	Lets you choose the date order display.
Hour Format	Lets you choose the hour display.
Minute	Lets you choose the minute display.
Sec Format	Lets you choose the second display.
AM/PM	Lets you choose the AM/PM display.
Date/Time Order	Lets you choose the date/time order display.
Format Display	Displays the current format.
Sample	Displays a sample of the current format.

- 7 When you have finished selecting the Date/Time format options, click OK.

Rapid SQL accepts the date/time format changes and closes the Date/Time Format Builder dialog box.

- 8 On the Options Editor, select the appropriate Default Date/Time Format options:

Option	Description
Use Calendar Control as default	If selected, Rapid SQL uses the <a href="#">Calendar Control</a> window.
Two-digit year system setting warning	If selected, Rapid SQL sends a warning when you use a two-digit year system setting.

- 9 Click OK.

Rapid SQL accepts the Default Date/Time Format changes and closes the Options Editor.

**NOTE:** To use a different format for a particular session, change the date/time at the session level.

### Editing Grid Date/Time

You can change the date/time display for a particular session when working in the [Data Editor](#). The Data Editor does not maintain the format changes once you close your session. To make this display permanent, use the [Editing Global Date/Time Format](#).

To edit the grid date and time, do the following:

- 1 On the Datasource Explorer, select the target table.
- 2 Right-click the table, and then click Edit Data.  
Rapid SQL opens the [Data Editor](#).
- 3 On the Data Editor tool bar, click Date/Time Format Builder.  
Rapid SQL opens the Date/Time Format Builder.
- 4 On the Date/Time Format Builder, click the Date/Time Format list, and then click the target predefined date/time format.
- 5 To customize the date/time format to your specifications, click Customize.  
Rapid SQL opens the Date/Time Format Builder dialog box.
- 6 On the Date/Time Format Builder dialog box, select the appropriate Date/Time Format Options:

Option	Description
Date/Time Format	Displays the predefined Date/Time format.
Day Format	Lets you choose the day display.
Separator	Lets you choose the display separator between the day, month, and year.

Option	Description
Month Format	Lets you choose the month display.
Year Format	Lets you choose the year display.
Date Order	Lets you choose the date order display.
Hour Format	Lets you choose the hour display.
Minute	Lets you choose the minute display.
Sec Format	Lets you choose the second display.
AM/PM	Lets you choose the AM/PM display.
Date/Time Order	Lets you choose the date/time order display.
Format Display	Displays the current format.
Sample	Displays a sample of the current format.

- When you have finished selecting the Date/Time format options, click OK.

Rapid SQL accepts the date/time format changes and closes the Date/Time Format Builder dialog box.

- To undo changes, on the Data Editor tool bar, click Undo Change.
- To redo changes, on the Data Editor tool bar, click Redo Change.

**NOTE:** Date/Time formats changed on a table level are valid for that session only.

### Editing Column Date/Time

You can change the date/time display for a particular column when working in the [Data Editor](#). The Data Editor does not maintain the format changes once you close your session. To change the format for the entire grid, see [Editing Grid Date/Time Format](#). To make this display permanent, [Editing Global Date/Time Format](#).

To edit the column date and time, do the following:

- On the Datasource Explorer, select the target table.
- Right-click the table, and click Edit Data.

Rapid SQL opens the [Data Editor](#).

- On the Data Editor, click the column header to select the column.
- Right-click the column and click Format.

Rapid SQL opens the Date/Time Format Builder.

- On the Date/Time Format Builder dialog box, click the Date/Time Format list, and then click the target predefined date/time format.
- To customize the date/time format to your specifications, click Customize.

Rapid SQL opens the Date/Time Format Builder dialog box.

- 7 On the Date/Time Format Builder dialog box, select the appropriate Date/Time Format Options:

Option	Description
Date/Time Format	Displays the predefined Date/Time format.
Day Format	Lets you choose the day display.
Separator	Lets you choose the display separator between the day, month, and year.
Month Format	Lets you choose the month display.
Year Format	Lets you choose the year display.
Date Order	Lets you choose the date order display.
Hour Format	Lets you choose the hour display.
Minute	Lets you choose the minute display.
Sec Format	Lets you choose the second display.
AM/PM	Lets you choose the AM/PM display.
Date/Time Order	Lets you choose the date/time order display.
Format Display	Displays the current format.
Sample	Displays a sample of the current format.

- 8 When you have finished selecting the Date/Time format options, click OK.

Rapid SQL accepts the date/time format changes and closes the Date/Time Format Builder dialog box.

- To undo changes, on the Data Editor tool bar, click Undo Format.
- To redo changes, on the Data Editor tool bar, click Redo Format.

**NOTE:** Date/Time formats changed on a column level are valid for that session only.

## DEFAULT VALUE HANDLING

When adding new records, when the transaction committing a row occurs, if no value has been entered for a column defined as having a default value, the default value is used for that column and validation is not required. The default value is not visible in the Data Editor grid until you commit the record and use the Reload Data button to refresh the grid.

In the case of tables with all columns defined as having default values, for example, you can add multiple records by repeatedly clicking the Insert New Record and Save Current Row buttons.

## Page Setup

The table below describes the options and functionality on the Page Setup dialog box:

Option	Functionality
Margins	Lets you select the size of the left, right, top, and bottom margins.
Titles and Gridlines	Lets you select options.
Preview	Displays how the table will appear when printed.
Page Order	Lets you specify when to print columns and rows.
Center on Page	Lets you select how table floats on the page.

## CODE GENERATION FACILITY

The Rapid SQL Code Generation Facility offers a quick way to generate DML statements for tables and views. The Code Generation Facility provides an efficient way to establish and enforce coding standards within the application by generating code with standardized formatting. The Code Generation Facility also lets you:

- Create packages and procedures for Oracle.
- Create procedures for IBM DB2 LUW, Microsoft SQL Server, and Sybase ASE.

**NOTE:** The Code Generation Facility can generate procedures for IBM DB2 LUW that are based on tables but not views.

The table below describes the options and functionality on the Embarcadero Code Generator:

Option	Description
Datasource List	To change the default datasource, click the list and then click a new datasource.  NOTE: If you are not connected to the target datasource, click the Connect button to establish a connection.  The current datasource, and database and datasource user name where applicable, are preselected when you open Code Generation Facility.
Database List	Lets you click the list and then click a target database.
Owner List	Lets you click the list and then click the new owner.
Tables Option Button	Click this option to see a list of the tables in the scroll box. In the scroll box, click the target table to be used as a code generation base.  NOTE: The list of objects corresponds to the selected datasource owner schema. If there are no tables in the target schema, the list will be empty.

Option	Description
Views Option Button	Click this option to see a list of the views in the scroll box. In the scroll box, click the target view to be used as a code generation base.  NOTE: The list of objects corresponds to the selected datasource owner schema. If there are no views in the target schema, the list will be empty.
Select 1 or More Where Clause	In the grid, select the check boxes that correspond to the target where clauses.  NOTE: Columns of primary keys are preselected.
Select 1 or More Output Columns	In the grid, select the check boxes that correspond to the target output columns.  NOTE: All existing columns are preselected.
Provide and Output File Name	In the box, enter the file name in the File box or click the Browse button to locate the output file.
Open	Select to open the file after the procedure runs.
Execute Immediately	Select to execute the file immediately.
Generate	To generate an SQL statement for tables, click the Object Type list and then click the object type.  NOTE: If you are generating an SQL statement for Views, you only have the Select option.
Grant Execute to	If you are a DBA or have DBA privileges, select the Users, Roles, and Groups check boxes to grant execute privileges.

### Using the Code Generation Facility

- 1 On the *Tools* menu, click *Code Generation Facility*.

OR

On the *Tools* toolbar, click *Code Generation Facility*.

Rapid SQL opens the Embarcadero Code Generator dialog box.

## IMPORT DATA

It is often necessary to import data into database tables from an external source. Developers commonly need to bring sample test data into a database to assist with Use Case scenarios. These scenarios may simulate particular data retrievals where segments of the data are preferred over performing a full data load. Business Analysts often acquire spreadsheets from outside sources. It is helpful to them to load this data into tables to perform more in-depth queries and pull meaningful data to make informed decisions.

To leverage the power of Microsoft Excel, database users may prefer to pull data from the database and load it into spreadsheet. Once using manipulations like Average, Aggregate, Row Count, or simple additions or deletions have been used on the data, the need remains to import this massaged data back into the database. Additionally, once the data is placed back in the database, recurring reporting and documentation are easily accomplished.

Rapid SQL's Import Data tool eliminates the time-consuming, manual process of working with data. The Import Data Wizard lets you pull data from a text file or a Microsoft Excel spreadsheet.

### Important Notes

None

### Completing the Import Data Wizard

- 1 On the *Tools* menu, click *Import Data*.

Rapid SQL opens the [first panel](#) of the Import Data Wizard.

- 2 On the first panel of the Wizard, specify the location of the file and enter the catalogue, schema (owner), and table into which the data will be imported.
- 3 Click Next.
- 4 On the second panel of the Wizard, select the character delimited type for the columns in the data file – Tab, Semi colon, Comma, Space, or Other (Custom – Tilde, Ampersand etc.)
- 5 You can select the option that the first row of the data contains field names.
- 6 Click Next.
- 7 On the third panel of the Wizard, you can preview and confirm the data format. You can also use the custom mapping to match columns in the table.
- 8 If you need to make any changes, click Back to scroll back to the appropriate panels of the wizard to make your corrections.
- 9 Click Finish.

Rapid SQL generates the Insert statements that can be saved as a \*.sql file for re-use across several datasources (versus simply loading the data directly into the database).

## IMPORT DATA WIZARD - PANEL 1

The table below describes the options and functionality on this panel of the Import Data Wizard:

Option	Description
Specify the file to be used in this data load operation.	Lets you type the file path or browse for the file.



Option	Description
Which table do you want to load data into?	Lets you select the catalog, schema, and the table to create the insert statements.

For more information, see [Import Data](#).

## IMPORT DATA WIZARD - PANEL 2 FOR TEXT FILES

The table below describes the options and functionality on this panel of the Import Data Wizard:

Option	Description
What character delimited the columns in the data file?	Lets you select the character delimited type for the columns in the data file: Tab, Semicolon, Comma, Space, or Other (Custom – Tilde, Ampersand etc.)
First Row Contains Field Names	Select for the first row of the data to contain field names.

For more information, see [Import Data](#).

## IMPORT DATA WIZARD - PANEL 3 FOR TEXT FILES

This panel of the Import Data Wizard lets you assign column names to fields. Column names in red indicate an invalid mandatory column names. Blue columns are already used.

For more information, see [Import Data](#).

## IMPORT DATA WIZARD - PANEL 2 FOR EXCEL FILES

The table below describes the options and functionality on this panel of the Import Data Wizard:

Option	Description
The spreadsheet you have selected contains more than one worksheet. Which worksheet contains the data you wish to import?	Lets you select the worksheet to import.
Start Cell	Lets you type the first cell of data to import.
End Cell	Lets you type the last cell of data to import.
First Row Contains Field Names	Select to assign column names to fields in the grid.

For more information, see [Import Data](#).

## EMBARCADERO PRODUCTS

The Tools menu lists all installed Embarcadero Technologies products. This lets you toggle to or start another Embarcadero product.

To open the Performance Center web client, do the following:

## CODE WORKBENCH

The Code Workbench is a multi-tabbed dialog that lets you enable and configure the following, time-saving ISQL editor features:

- **Code templates** - common, multi-line blocks of code you can add to a script with a couple of keystrokes and modify to the purpose of the script. For information on using this feature, see [Using Code Templates](#).
- **Auto replace** - commonly used SQL syntax or commands that can be added to a script using a shortcut, replacement expression. For information on using this feature, see [Using Auto Replace Expressions](#).

The Code Workbench lets you modify, add, and delete Code Template and Auto Replace resources. In addition, you can save your Code Workbench settings for the purpose of user-sharing or customization. The following topics provide detailed instructions and information on setting up these features:

- [Enabling the Code Templates and Auto Replace Features](#)
- [Maintaining Code Template Definitions](#)
- [Maintaining Auto Replace Expressions](#)
- [Loading and Saving Custom Code Workbench Settings](#)

## ENABLING THE CODE TEMPLATES AND AUTO REPLACE FEATURES

In order to use the Code Templates or Auto Replace features, they must be enabled.

### To enable the Code Templates or Auto Replace features

- 1 From the **Tools** menu, select **Code Workbench**. The **Code Workbench** dialog opens.
- 2 Use the following table as a guide to using the controls on the **Settings** panel:

Group	Settings and descriptions	
<b>Auto Replace</b>	<b>Enable Auto Replace</b>	Select this checkbox to use this feature in the ISQL editor.
	<b>Replace substrings</b>	If unselected, only blank-delimited occurrences of auto replace expressions are replaced when an activation event occurs. If selected, all occurrences of auto replace expressions are replaced.  For example, consider an Auto Replace expression of <b>BEG</b> that is to be replaced with <b>Begin</b> . With <b>Replace substrings</b> enabled, <b>BEGBEG</b> would be replaced with <b>BeginBegin</b> when an activation event occurs. With <b>Replace substrings</b> disabled, no replacement would occur.
<b>Code Templates</b>	<b>Enable Code Templates</b>	Select this check box to use this feature in the ISQL editor.

- 3 Either click **OK** to dismiss the dialog or use the other tabs to configure the code templates and auto replace features.

For information on configuring the Code Template and Auto Replace features, see [Maintaining Code Template Definitions](#) and [Maintaining Auto Replace Expressions](#).

See also:

- [Using Code Templates](#)
- [Using Auto Replace Expressions](#)

## MAINTAINING CODE TEMPLATE DEFINITIONS

Rapid SQL loads a default set of code templates automatically on startup. The Code Workbench lets you add templates to the currently loaded templates, edit any existing templates, or delete templates.

### To modify the currently loaded code templates

- 1 From the **Tools** menu, select **Code Workbench**.
- 2 Select the **Code Templates** tab.

- 3 Use the following table as a guide to performing the template maintenance tasks offered on this tab:

Task	Steps
To add a new template	Click <b>Add</b> to open the <b>Edit Code Template expression</b> dialog and provide the following information:
	<p><b>Shortcut</b> - used to identify the template when the Code Assist menu is open.</p> <p><b>Description</b> - displayed in the <b>Code Templates</b> tab listing for this template.</p> <p><b>Platform</b> - Select a specific DBMS platform or select ALL. The template will be available in ISQL editor windows with connections to the specified platform or all platforms, accordingly.</p> <p><b>Template</b> - Type the code block that will be added at the cursor when this template is selected.</p>
To edit a template	Select a template from the list, click <b>Edit</b> to open the <b>Edit Code Template expression</b> dialog, and edit the template values as per the descriptions above.
To delete a template	Select a template from the list, click <b>Delete</b> and when prompted to confirm, click <b>Yes</b> .

**NOTE:** You can also save the currently loaded code templates and other Code Workbench settings to an XML file for subsequent loading. This lets you share settings among users and tailor Code Template and Auto Replace settings to particular environments or projects. For details, see [Loading and Saving Custom Code Workbench Settings](#).

## MAINTAINING AUTO REPLACE EXPRESSIONS

Rapid SQL is packaged with a set of default Auto Replace definitions, stored in the registry and loaded automatically on startup. The Code Workbench lets you add new definitions, edit any existing definitions, or delete definitions from the currently loaded set.

### To modify the currently loaded expressions

- 1 From the **Tools** menu, select **Code Workbench**.
- 2 Select the **Auto Replace** tab.

- 3 Use the following table as a guide to performing the template maintenance tasks offered on this tab:

Task	Steps
To add a new expression	<p>Click <b>Add</b> to open the <b>Edit Auto Replace Expression</b> dialog and provide the following information:</p> <p><b>Expression</b> - the shortcut that will be replaced.</p> <p><b>Activation</b> - Type one or more individual keys that are to trigger replacement in the ISQL editor. Valid activation keystrokes are SPACE, TAB, and RETURN. As you type, the following representations are shown in the <b>Activation</b> box: \s represents SPACE, \t represents TAB, and \n represents RETURN (newline).</p> <p><b>Replace With</b> - Type the command or SQL syntax that will be replaced when one of the activation events occurs.</p>
To edit an expression	Select an expression from the list, click <b>Edit</b> to open the <b>Edit Auto Replace Expression</b> dialog, and edit the expression values as per the descriptions above.
To create an expression based on an existing expression	Select an expression from the list, click <b>Clone</b> to open the <b>Edit Auto Replace Expression</b> dialog, prepopulated with values of the existing expression. Edit the expression values as per the descriptions above.
To delete an expression	Select an expression from the list, click <b>Delete</b> and when prompted to confirm, click <b>Yes</b> .

**NOTE:** You can also save the currently loaded Auto Replace expressions and other Code Workbench settings to an XML file and subsequently load those settings. This lets you share settings among users and tailor Auto Replace and Code Template settings to particular environments or projects. For details, see [Loading and Saving Custom Code Workbench Settings](#).

## LOADING AND SAVING CUSTOM CODE WORKBENCH SETTINGS

At each startup, Rapid SQL loads a default set of code templates and auto replace definitions stored in the registry. The enabled/disabled status of the two features is retained on shutdown.

After changing the enabled/disabled status of the Code Templates or Auto Replace features, or modifying the associated templates or replacement expressions, you can save the status and currently loaded resources to an XML file. The ability to save, load, and restore settings means you can:

- Share Code Workbench settings among users
- Create custom Code Workbench settings tailored to DBMS platforms or coding projects

## To make use of Code Workbench file options

- 1 From the **Tools** menu, select **Code Workbench**.
- 2 Use the following table as a guide to performing the Code Workbench file option tasks:

Task	Steps
To save the current Code Workbench settings to an XML file:	Click <b>Export Settings</b> and use the <b>Save As</b> dialog to select a location and provide a name for the exported file.
To replace the currently loaded Code Workbench settings with a set of saved settings:	Click <b>Import Settings</b> and use the <b>Open</b> dialog to locate and select the settings file to be loaded.
To restore the default Code Workbench settings from the registry:	Click <b>Restore Settings</b> .

See also:

- [Maintaining Code Template Definitions](#)
- [Maintaining Auto Replace Expressions](#)

## CODE ANALYST

The Code Analyst is a tool to identify time-consuming lines of code. Code Analyst lets you:

- Perform detailed response time analysis on the execution of [Procedures](#) and [Foreign Keys](#).
- Benchmark the execution of one or more procedures or functions to determine exactly what code objects and lines of code are taking the longest to run.
- Save response time metrics and perform intelligent compares against current execution times so you can determine deviations from previous acceptable response times.

**TIP:** You can set Code Analyst options in the [Code Analyst Options](#).

**NOTE:** Availability of this feature depends on your Rapid SQL licensing. For more information, see [Licensing](#).

### Important Notes

- For DB2, before profiling with Code Analyst, [Compile](#) all procedures with the debugging option selected.
- For Oracle, when using the Oracle Debugger, [Compile](#) all procedures with the debugging option selected before profiling with Code Analyst.

### Common Tasks

[Creating a Code Analyst Session](#)

[Identifying and Fixing Bottlenecks Using Code Analyst](#)

[Comparing Code Analyst Sessions](#)

[Cloning a Code Analyst Session](#)

[Deleting a Code Analyst Session](#)

[Stopping a Code Analyst Session Execution](#)

[Executing a Code Analyst Session](#)

[Scheduling a Code Analyst Session](#)

[Unscheduling a Code Analyst Session](#)

[Refreshing a Code Analyst Session](#)

[Saving Results in Code Analyst](#)

[Printing Results in Code Analyst](#)

[Viewing Run Details in Code Analyst](#)

[Viewing Unit Summary Information in Code Analyst](#)

[Viewing Unit Details in Code Analyst](#)

[Setting View Options for the Unit Detail Tab in Code Analyst](#)

[Extracting SQL Text in Code Analyst](#)

[Executing SQL in Code Analyst](#)

## CODE ANALYST DBMS NOTES

Code Analyst is available for:

- Microsoft SQL Server 7 or later
- Oracle 7 or later
- IBM DB2 LUW 8
- Sybase ASE 12.0.0.3 or later

Rapid SQL utilizes debugger technology to capture the data for each line of executed code. For Oracle, you can use the debugger or using Oracle's supplied DBMS\_Profiler package.

**TIP:** For Oracle, you can specify to use the debugger or the DBMS\_Profiler package on the [Code Analyst Options](#).

The Code Analyst will step through each line of code, stopping to record data for those lines of code onto which a breakpoint can be issued. Some debuggers cannot capture time metrics for all lines of a stored procedure or function.

Procedures and functions that contain looping constructs will require more time to run. The additional amount of time needed to run is proportional to the number of iterations in the loop.

For more information, see:

[IBM DB2 LUW 8i Data Captured by Code Analyst](#)

[Microsoft SQL Server Data Captured by Code Analyst](#)

[Oracle Data Captured by Code Analyst](#)

[Sybase ASE Data Captured by Code Analyst](#)

### IBM DB2 LUW 8i DATA CAPTURED BY CODE ANALYST

Code Analyst utilizes the IBM Debugger when capturing time data.

The debugger is verified to run on IBM DB2 LUW version 7.2 and up. There is a known issue running version 7.2 with Fixpack 9.

DB2 has documented limitations on lines of code can be profiled.



The following are SQL statements that are NOT valid break point lines:

```
BEGIN
BEGIN
BEGIN NOT ATOMIC
BEGIN ATOMIC
CLOSE CURSOR
DECLARE cursor WITH RETURN FOR <sql statement>
DECLARE , var without default
DECLARE CONDITION (CONDITION) FOR SQLSTATE (VALUE) "..."
DECLARE CONTINUE HANDLER
DECLARE CURSOR
DECLARE EXIT HANDLER
DECLARE RESULT_SET_LOCATOR [VARYING]
DECLARE SQLSTATE
DECLARE SQLCODE (unless there is a default)
DECLARE UNDO HANDLER (unless they are entered)
DO
ELSE
END
END CASE
END IF
END FOR
END REPEAT
END WHILE
ITERATE
LEAVE
LOOP
OPEN CURSOR
REPEAT (as a keyword alone)
RESIGNAL
```

SIGNAL

THEN

labels, e.g. P1:

**NOTE:** Code containing these statements will not have times associated with them.

For more information, see [Code Analyst DBMS Notes](#).

## MICROSOFT SQL SERVER DATA CAPTURED BY CODE ANALYST

In order to execute a Code Analyst session against a Microsoft SQL Server database, the SQL Server debugger must be installed and functioning properly. Please refer to [Embarcadero SQL Debugger for Microsoft SQL Server](#) for details concerning set up.

### Related Information

[Code Analyst DBMS Notes](#)

## ORACLE DATA CAPTURED BY CODE ANALYST

When using the PL/SQL Profiler, Oracle has documented an issue regarding extremely large times being returned by the profiler. The times are sometimes hundred times larger than the actual run time of the stored procedure or function. Oracle documents that this is a vendor/os problem rather than an Oracle problem, because the RDTSC instruction is reporting wrong time stamp counter. They indicate that they have seen this problem on some INTEL Pentium processors.

### Related Information

[Code Analyst DBMS Notes](#)

## SYBASE ASE DATA CAPTURED BY CODE ANALYST

Sybase has documented a problem with their debugger API. The problem involves reporting the wrong line number through the debugger. Because of this bug, Code Analyst may report back data for blank lines or lines that contain comments. Sybase has fixed this problem release 12.5.2 of the database. All procedures affected must be dropped and recreated in order to correct the problem.

### Related Information

[Code Analyst DBMS Notes](#)

# CODE ANALYST REQUIREMENTS

## Debuggers

The Code Analyst uses fully configured Embarcadero SQL Debuggers to profile, and therefore availability of this feature depends on your Rapid SQL licensing. For more information, see the following topics:

- [Licensing](#)
- [Rapid SQL Add-On Tools](#)

## Using Code Analyst with the Oracle Profiler

Oracle users have the option of either using the Oracle Debugger or the Oracle Profiler with Code Analyst to capture statistics. To use Code Analyst with the profiler option, Oracle's profiler package must be installed. The install is user specific, so it must be installed by each user wishing to use Code Analyst. To install the package, users can invoke the [Rapid SQL PL/SQL Profiler](#).

**TIP:** You can set profiler options in the [Code Analyst Options](#). You can specify that Code Analyst display the actual run time on the database, and does not include the time it takes to get to the server.

## Privileges

For Oracle, SYS privileges are required to install the Code Analyst tables. If you do not have SYS privileges, ask your Server Administrator to log into the Oracle datasource as SYSDBA, and then open Code Analyst to install the tables.

During install, the following privileges are set for the Code Analyst tables.

- DB2 – Permissions are granted to the Public Group
- Oracle – Permissions are granted to the Public group.
- Microsoft SQL Server – Permissions are granted to the Public role.
- Sybase – Permissions are granted to the Public group.

All users can use the Code Analyst but each user will only see their own run ids. Users need to belong to the public group.

**TIP:** You can check this/modify privileges in the Users Editor.

# INSTALLING CODE ANALYST

To install Code Analyst, do the following:

- 1 On the Tools menu, select *Code Analyst*.
- 2 In Select the database you would like to install the tables on, select a database.

- 3 For IBM DB2 LUW for Open Systems, in select the tablespace you would like to install the tables on, select a tablespace.
- 4 For IBM DB2 LUW for Open Systems, in select the schema you would like to install the tables on, select a schema. The default is EMBTCA schema.
- 5 In Select the filegroup you would like to install the tables to, select a filegroup.

Code Analyst installs the following repository tables in the repository:

- EMBT\_CODE\_ANA\_RUNS - Holds all the code analyst sessions created by users.
- EMBT\_CODE\_ANA\_UNITS - Holds all the objects to be run for all.
- EMBT\_CODE\_ANA\_PARAMS - Contains all the parameters for the objects that were run.
- EMBT\_CODE\_ANA\_DATA - Contains the run data and is used to populate all the charts and statistics.
- EMBT\_CODE\_ANA\_VERSION - Contains the version number of code analyst.

Code Analyst opens to the Run Summary tab.

- 6 Create a session using the [Creating a Code Analyst Session](#).

## UNINSTALLING CODE ANALYST

The Uninstall functionality lets you uninstall Code Analyst from the server.

**NOTE:** To uninstall a repository table, you need create table and grant privileges. Generally, you need sysadmin privileges.

- 1 On the Tools menu, select *Code Analyst*.
- 2 Select a session or object, and then select Uninstall.

Code Analyst removes the repository tables in the repository:

- EMBT\_CODE\_ANA\_RUNS - Holds all the code analyst sessions created by users.
- EMBT\_CODE\_ANA\_UNITS - Holds all the objects to be run for all.
- EMBT\_CODE\_ANA\_PARAMS - Contains all the parameters for the objects that were run.
- EMBT\_CODE\_ANA\_DATA - Contains the run data and is used to populate all the charts and statistics.
- EMBT\_CODE\_ANA\_VERSION - Contains the version number of code analyst.

## CODE ANALYST PRODUCT DESIGN

Code Analyst performs detailed response time analysis. Code Analyst steps through each line of code and profiles those lines of code that the debugger or profiler can capture time metrics for.

**NOTE:** Some debuggers do not capture time metrics for all lines of a procedure or function. For more information, see [Code Analyst DBMS Notes](#).

After capturing the time metrics, Code Analyst displays the data in an easy-to-read format on the tabs.

The Code Analyst is comprised of the following tabs:

Tab	Option	Description
Run Summary	Session	Lets you select the run session(s).
	Session	Displays the name of the session as created by the user.
	Run ID	Displays the run ID for the run(s). This number is system generated.
	Run Date	Displays the time and date of the session.
	Total Profile Time (ms)	Displays the total time taken for the profiled code to execute. This time is limited to the lines of code that are profiled. Overhead is not included in this calculation.
	Total Analysis Time	Displays the total time taken for the session to complete, including all overhead time needed to analyze the procedure or function.
	Scheduler	Displays the scheduler used to schedule the session. This information displays until the scheduled job has been run.
Run Detail	Session	Lets you select the object execution session.
	Run	Lets you select the object execution.
	Unit Type	Lets you select the object type for the object execution.
	Unit Owner	Lets you select the object owner for the object execution.
	Unit Database	Lets you select the object database for the object execution.
	Time Unit	Lets you specify the time unit for the Unit Execution graph.
	Unit Owner	Displays the owner of the procedure or function.
	Unit Name	Displays the name of the procedure or function.
	Unit Type	Displays the types of captured objects, including Anonymous Block, Function, Package Body, and Procedure for Oracle databases. Also displays SQL Statement and Procedure for the other platforms.
	Unit Database	Displays the database on which the object is stored.
	Total Profiled Time	Displays the total time taken for the profiled code to execute.
	% of Profiled Time	Displays the percentage of the Total Profiled Time for the run that this unit accounts for.
Comparison	Base Run	Lets you select the earlier object execution.

Tab	Option	Description
	New Run	Lets you select the later object execution.
	Unit Owner	Displays the owner of the procedure or function.
	Unit Name	Displays the name of the procedure or function.
	Unit Type	Displays the types of captured objects, including Anonymous Block, Function, Package Body, and Procedure for Oracle databases. Also displays SQL Statement and Procedure for the other platforms.
	Unit Database	Displays the object database for the object execution.
	Time Diff	Displays the time difference in milliseconds between the base run and the new run.
	New Profiled Time	Displays the profiled time of the new run.
	Base Profiled Time	Displays the profiled time of the base run.
Unit Summary	Unit Owner	Lets you select the object owner for the session(s).
	Unit Name	Lets you select the object name for the session(s).
	Unit Database	Lets you select the object database for the session(s).
	Number of Top Runs	Lets you specify the number of top object executions to display in the Top 5 Runs graph and select the unit of time for the Unit Time graph.
	Session	Displays the name of the session as created by the user.
	Run ID	Displays the unique id for the Run. This number is system generated.
	Run Date	Displays the time and date of the session.
	Total Analysis Time	Displays the total time taken for the session to complete, including all overhead time needed to analyze the procedure or function.
	Total Profiled Time	Displays the total time taken for the profiled code to execute.
	Unit Profiled Time	Displays the unit time for the session(s).
	% of Profiled Time	Displays the percentage of the Total Profiled Time for the run that this unit accounts for.
	% of Run Time	Displays the percentage of object execution time for the session(s).
Unit Detail	Session	Lets you select the session.
	Run	Lets you select the object execution.
	Unit Name	Lets you select the object name for the session.
	Number of Top Lines	Lets you specify the number of top lines in the Top 5 Lines Execution Time graph and select the total time units.
	Percentage Calculation	Lets you specify total object execution time or object run time.
	Calls	Displays the number of times the line was executed.
	Total Time	Displays the total time the line was executed.
	% of Total Profiled	Displays the percentage of the Total Profiled Time that this line of code was responsible for.

Tab	Option	Description
	Avg Time	Displays the average profiled time for this line.
	Min Time	Displays the minimum recorded time for execution of this line.
	Max Time	Displays the maximum recorded time for execution of this line.
	Dependency	Displays the UNIT_NUMBER of the dependency object that was called by that line. Lets you right-click and quickly go to that UNIT_NUMBER to see its Unit detail information.
	Source	Displays the object's SQL source code.

## Common Tasks

[Creating a Code Analyst Session](#)

[Identifying and Fixing Bottlenecks Using Code Analyst](#)

[Comparing Code Analyst Sessions](#)

[Cloning a Code Analyst Session](#)

[Deleting a Code Analyst Session](#)

[Stopping a Code Analyst Session Execution](#)

[Executing a Code Analyst Session](#)

[Scheduling a Code Analyst Session](#)

[Unscheduling a Code Analyst Session](#)

[Refreshing a Code Analyst Session](#)

[Saving Results in Code Analyst](#)

[Printing Results in Code Analyst](#)

[Viewing Run Details in Code Analyst](#)

[Viewing Unit Summary Information in Code Analyst](#)

[Viewing Unit Details in Code Analyst](#)

[Setting View Options for the Unit Detail Tab in Code Analyst](#)

[Extracting SQL Text in Code Analyst](#)

[Executing SQL in Code Analyst](#)

## CODE ANALYST TUTORIAL

The following tutorial guides you through the process of using the Code Analyst.

## Creating a Code Analyst Session

- 1 On the Tools menu, select Code Analyst.

Initially, Code Analyst installs the repository tables. For more information, see [Installing Code Analyst](#). Then Rapid SQL opens the Code Analyst to the Run Summary tab.

- 2 On the Code Analyst Tools toolbar, click the Create New Collection button.

Rapid SQL opens the first panel of the Code Analyst Wizard.

- 3 Select the individual object or group of objects to analyze. In this example, an individual stored procedure (CREATE\_ADMISSION2) is selected.

**TIP:** Code Analyst does not let you select objects that do not have stored procedures.

- 4 Click Next.

If the object(s) selected to be analyzed requires parameters, the second panel of the wizard prompts you to enter the parameters.

- 5 Double-click the object to set the parameters.


- 6 For IBM DB2 LUW for Open Systems and Oracle, the Compile button opens the Confirm Compile dialog box that lets you compile the objects to ensure that the Code Analyst can capture the time metrics.

- 7 Click Finish.

Code Analyst displays a message that the Code Analyst will run longer than the actual code. Then Code Analyst analyzes the objects, using the Embarcadero SQL Debugger to profile and then opens the Run Detail tab.

**TIP:** You can select the "Please do not show me this dialog again" option in the dialog box or set the option on the [Code Analyst Options](#).



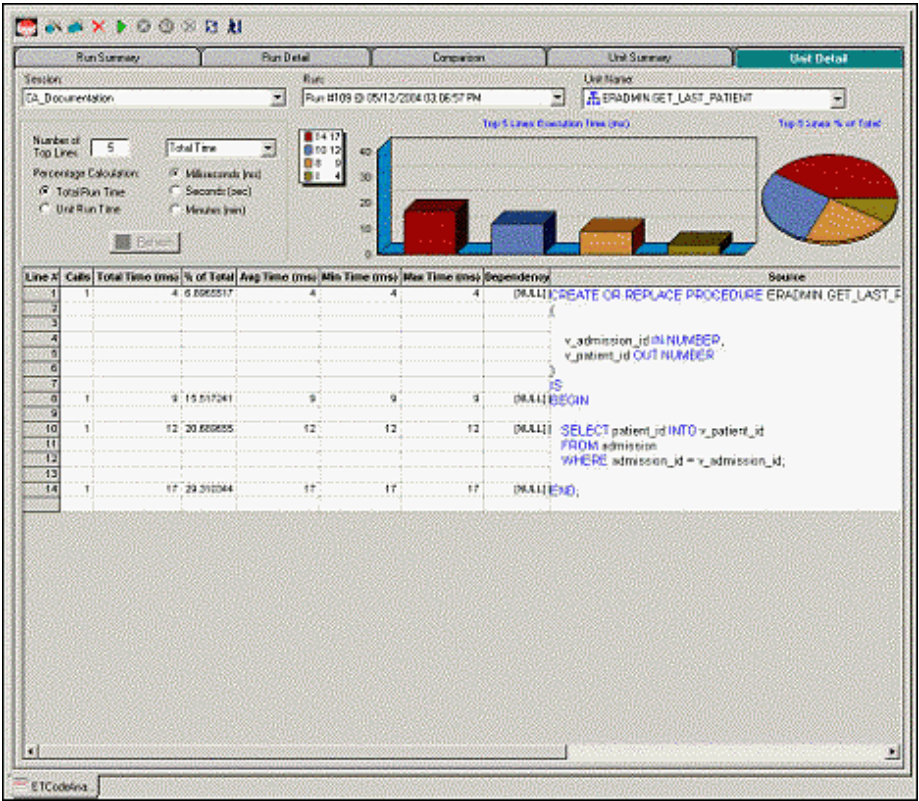


## Identifying and Fixing Bottlenecks Using Code Analyst

The Run Detail tab displays the total time for the objects being analyzed. The tab information may be enough to identify the potential bottleneck.

- 1 To view more detailed information, double-click the Unit Name.

Code Analyst opens to the Unit Detail tab.

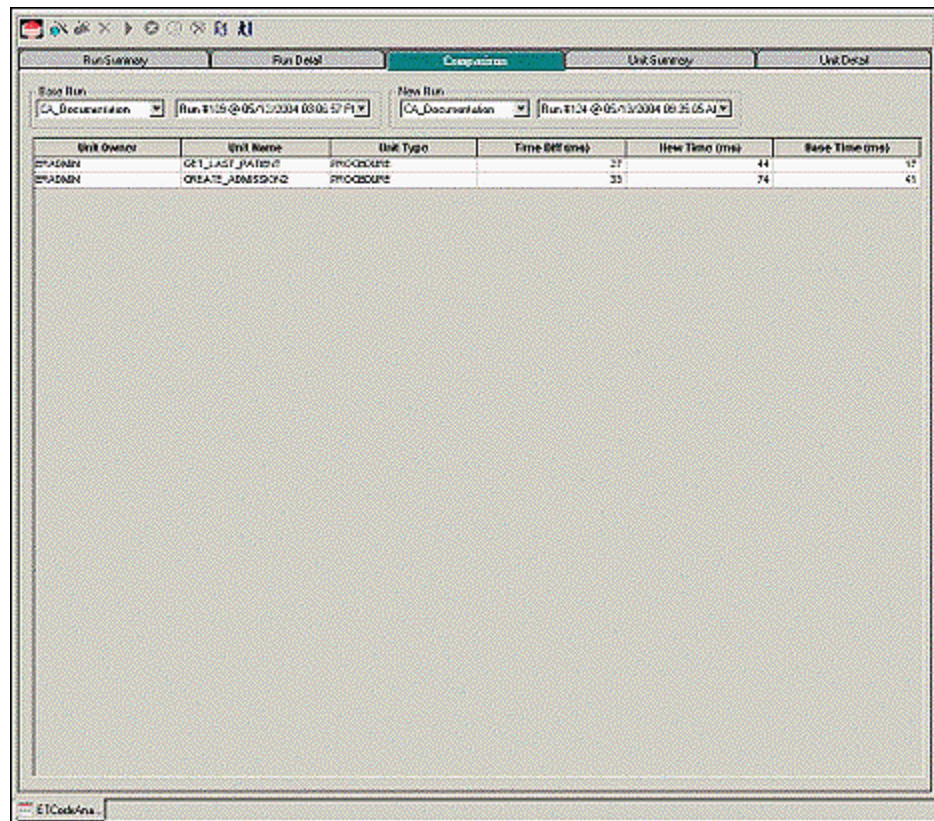


The Unit Detail Tab displays the object code and other information related to the individual lines of code. You can determine which line of code is taking too long and why. The Unit Detail Tab is where you troubleshoot, and then resolve the problem in the [Modifying objects using editors](#).

- 2 In Rapid SQL, open the object editor, and then modify the code on the Definition tab.
- 3 Click Alter.
- 4 In Code Analyst, on the Unit Detail Tab, click Execute.

## Comparing Code Analyst Sessions

- 1 Click the Comparison tab.



Code Analyst has a Comparison facility to allow quick compares of two object executions, showing the base time and the new time, as well as the time differences. The Comparison tab lets you compare which of the two procedures or functions ran faster.

- 2 Examine the Time Diff which indicates improvement to code.
- 3 If necessary, continue to modify the code on the Definition tab of the object editor, and then press Alter.
- 4 In Code Analyst, on the Unit Detail Tab, click Execute.
- 5 Examine the Time Diff until the bottleneck is solved.

## Common Tasks

[Creating a Code Analyst Session](#)

[Identifying and Fixing Bottlenecks Using Code Analyst](#)

[Comparing Code Analyst Sessions](#)

[Cloning a Code Analyst Session](#)



[Deleting a Code Analyst Session](#)[Stopping a Code Analyst Session Execution](#)[Executing a Code Analyst Session](#)[Scheduling a Code Analyst Session](#)[Unscheduling a Code Analyst Session](#)[Refreshing a Code Analyst Session](#)[Saving Results in Code Analyst](#)[Printing Results in Code Analyst](#)[Viewing Run Details in Code Analyst](#)[Viewing Unit Summary Information in Code Analyst](#)[Viewing Unit Details in Code Analyst](#)[Setting View Options for the Unit Detail Tab in Code Analyst](#)[Extracting SQL Text in Code Analyst](#)[Executing SQL in Code Analyst](#)

## USING THE CODE ANALYST

When working with database code stored in database objects, it is sometimes difficult to pinpoint bottlenecks within the code. When situations like this arise, Code Analyst can assist in identifying the trouble spots. Code Analyst can be used to analyze one object or a group of objects. You select one or multiple objects, execute them, view the results, and save those results for later viewing or comparing.

### Common Tasks

[Creating a Code Analyst Session](#)[Identifying and Fixing Bottlenecks Using Code Analyst](#)[Comparing Code Analyst Sessions](#)[Cloning a Code Analyst Session](#)[Deleting a Code Analyst Session](#)[Stopping a Code Analyst Session Execution](#)[Executing a Code Analyst Session](#)[Scheduling a Code Analyst Session](#)[Unscheduling a Code Analyst Session](#)

[Refreshing a Code Analyst Session](#)

[Saving Results in Code Analyst](#)

[Printing Results in Code Analyst](#)

[Viewing Run Details in Code Analyst](#)

[Viewing Unit Summary Information in Code Analyst](#)

[Viewing Unit Details in Code Analyst](#)

[Setting View Options for the Unit Detail Tab in Code Analyst](#)

[Extracting SQL Text in Code Analyst](#)

[Executing SQL in Code Analyst](#)

## CREATING A CODE ANALYST SESSION

The Code Analyst Wizard creates a new Code Analysis session that creates data for the Code Analyst tabs:

- 1 On the Tools menu, select Code Analyst.
- 2 On the Code Analyst Tools toolbar, click Create New Collection.

Rapid SQL opens the first panel of the Code Analyst Wizard.

- 3 Select the individual object or group of objects to analyze.
- 4 Click Next.

If the object(s) selected to be analyzed requires parameters, the second panel of the wizard prompts you to enter the parameters.

- 5 Double-click the object to set the parameters.
- 6 For IBM DB2 LUW for Open Systems and Oracle, the Compile button opens the Confirm Compile dialog box that lets you compile the objects to ensure that the Code Analyst can capture the time metrics.
- 7 Click Finish.

Code Analyst displays a message that the Code Analyst will run longer than the actual code. Then Code Analyst analyzes the objects, using the Embarcadero SQL Debugger to profile and then opens the Run Detail tab.

**TIP:** You can also select the "Please do not show me this dialog again" option in the dialog box or set the option on the [Code Analyst Options](#).

For more information, see [Code Analyst Wizard](#).

## CODE ANALYST WIZARD

The table below describes the options and functionality of the Code Analyst Wizard:

Option	Description
Session Name	Lets you type a name.
Select by Owner	Code Analyst Wizard queries the database to get the list of procedures and functions and lets you select objects to retrieve Codes for. Select to display available objects by owner, and then select the database(s).
Select by Object	Code Analyst Wizard queries the database to get the list of procedures and functions and lets you select objects to retrieve Codes for. Select to display available objects by object, and then select the object(s).
Object Name	Double-click each object to specify the input parameters. Specify which object executes first by clicking the Up and Down buttons.
Compile	IBM DB2 LUW FOR OPEN SYSTEMS AND ORACLE ONLY: Opens the Confirm Compile dialog box that lets you compile the objects to ensure that the Code Analyst can capture the time metrics.
Schedule	Opens the Select Scheduler dialog box or opens scheduling application.
Finish	Code Analyst analyzes the code.

For more information, see: [Creating a Code Analyst Session](#).

## IDENTIFYING AND FIXING BOTTLENECKS USING CODE ANALYST

The Unit Detail Tab displays the object code and other information related to the individual lines of code. You can identify time-consuming lines of code in the Unit Detail Tab. The Unit Detail Tab is where you troubleshoot, and then resolve the problem in the [Modifying objects using editors](#).

- 1 On the Tools menu, select Code Analyst.
- 2 Click the Unit Detail Tab.  
Percent of Run Time displays the percentage of object execution time for the session(s).
- 3 Identify an object that contains time-consuming code.
- 4 In Rapid SQL, open the object editor, and then modify the code on the Definition tab.
- 5 Click Alter.
- 6 In Code Analyst, on the Unit Detail Tab, click Execute.
- 7 Click the Comparison tab.

The Comparison Tab lets you compare times of the objects in two different object executions to determine which run was more efficient. The Comparison Tab displays the base time and the new time, as well as the time differences. The Comparison tab lets you compare which of the two procedures or functions ran faster.

- 8 Examine the Time Diff which indicates improvement to code.
- 9 If necessary, continue to modify the code on the Definition tab of the object editor, and then press Alter.
- 10 Create new Code Analyst sessions and examine the Time Diff until the bottleneck is solved.

For more information, see:

[Code Analyst Product Design](#)

[Using the Code Analyst](#)

## COMPARING CODE ANALYST SESSIONS

The Comparison Tab lets you compare times of the objects in two different object executions to determine which run was more efficient. The Comparison Tab displays the base time and the new time, as well as the time differences. The Comparison tab lets you compare which of the two procedures or functions ran faster.

- 1 On the Tools menu, select Code Analyst.
- 2 On the Run Summary tab, right-click the sessions, and then select Compare.
- 3 Examine the Time Diff which indicates improvement to code.

For more information, see:

[Code Analyst Product Design](#)

[Using the Code Analyst](#)

## CLONING A CODE ANALYST SESSION

The Clone Collection functionality lets you clone an existing Code Analyst session using the Code Analyst Wizard. Clone lets you reset the parameters or the order of the objects in the session without creating a new session.

- 1 On the Tools menu, select Code Analyst.
- 2 On the Run Summary tab, select the session to clone.
- 3 On the Code Analyst Tools toolbar, click Clone Collection.

Rapid SQL opens the [first panel of the Code Analyst Wizard](#).

For more information, see:

[Code Analyst Product Design](#)

[Using the Code Analyst](#)

## DELETING A CODE ANALYST SESSION

The Delete Collection functionality lets you delete the selected Code Analyst session.

- 1 On the Tools menu, select Code Analyst.
- 2 On the Run Summary tab, select the session to delete.
- 3 On the Code Analyst Tools toolbar, click Delete Collection.

Code Analyst deletes the session.

For more information, see:

[Code Analyst Product Design](#)

[Using the Code Analyst](#)

## STOPPING A CODE ANALYST SESSION EXECUTION

The Stop Execution kills the execution of the selected collection.

- 1 On the Tools menu, select Code Analyst.
- 2 On the Run Summary tab, select the session to kill.
- 3 On the Code Analyst Tools toolbar, click Stop Execution.

Code Analyst kills the execution.

For more information, see:

[Code Analyst Product Design](#)

[Using the Code Analyst](#)

## EXECUTING A CODE ANALYST SESSION

The Execute Collection functionality extracts the SQL text and then executes the code.

- 1 On the Tools menu, select Code Analyst.
- 2 On the Run Summary tab, select the session to execute.
- 3 On the Code Analyst Tools toolbar, click Execute Collection.

Code analyst extracts and executes the SQL.

For more information, see:



[Code Analyst Product Design](#)

[Using the Code Analyst](#)

## SCHEDULING A CODE ANALYST SESSION

The Schedule Session functionality lets you schedule the session for a future run.

- 1 On the Tools menu, select Code Analyst.
- 2 On the Run Summary tab, select the session to schedule.
- 3 On the Code Analyst Tools toolbar, click Schedule Session.

Code Analyst opens the default scheduler.

For more information, see:

[Scheduling](#)

[Code Analyst Product Design](#)

[Using the Code Analyst](#)

## UNSCHEDULING A CODE ANALYST SESSION

The Delete Session functionality lets you remove the session from a schedule.

- 1 On the Tools menu, select Code Analyst.
- 2 On the Run Summary tab, select the session to unschedule.
- 3 On the Code Analyst Tools toolbar, click Delete Session.

For more information, see:

[Scheduling](#)

[Code Analyst Product Design](#)

[Using the Code Analyst](#)

## REFRESHING A CODE ANALYST SESSION

The Refresh Data functionality refreshes the data.

- 1 On the Tools menu, select Code Analyst.
- 2 On the tab, on the Code Analyst Tools toolbar, click Refresh Data.

For more information, see:

[Code Analyst Product Design](#)

[Using the Code Analyst](#)

## SAVING RESULTS IN CODE ANALYST

The Save functionality lets you save results for later viewing or comparing.

- 1 On the Tools menu, select Code Analyst.
- 2 On the tab, right-click the session or unit, and then select Save.

Code Analyst opens the Save Results dialog box.

For more information, see:

[Code Analyst Product Design](#)

[Using the Code Analyst](#)

## PRINTING RESULTS IN CODE ANALYST

The Print functionality lets you print results for later viewing or comparing.

- 1 On the Tools menu, select Code Analyst.
- 2 On the tab, right-click the session or unit, and then select Print.

Code Analyst opens the Print Results dialog box.

For more information, see:

[Code Analyst Product Design](#)

[Using the Code Analyst](#)

## VIEWING RUN DETAILS IN CODE ANALYST

The Run Detail tab displays the total time for the objects being analyzed. The tab information may be enough to identify the potential bottleneck.

To open the Run Details Tab in Code Analyst, do the following:

- 1 On the Tools menu, select Code Analyst.
- 2 On the Run Summary tab, right-click the session, and then select Run Detail.

OR

- 3 On the Unit Summary tab, right-click the session, and then select Run Detail.

For more information, see:

[Code Analyst Product Design](#)

## [Using the Code Analyst](#)

### VIEWING UNIT SUMMARY INFORMATION IN CODE ANALYST

The Unit Summary Tab in Code Analyst displays the individual runs for a session.

To open the Unit Summary Tab in Code Analyst, do the following:

- 1 On the Tools menu, select Code Analyst.
- 2 On the Comparison tab, right-click the session, and then select Unit Summary.

For more information, see:

[Code Analyst Product Design](#)

[Using the Code Analyst](#)

### VIEWING UNIT DETAILS IN CODE ANALYST

The Unit Detail Tab displays the object code and other information related to the individual lines of code. You can identify time-consuming lines of code in the Unit Detail Tab. The Unit Detail Tab is where you troubleshoot, and then resolve the problem in the [Modifying objects using editors](#).

To open the Unit Details tab in Code Analyst, do the following:

- 1 On the Tools menu, select Code Analyst.
- 2 On the Unit Summary tab, right-click the session, and then select Unit Detail.

For more information, see:

[Code Analyst Product Design](#)

[Using the Code Analyst](#)

### SETTING VIEW OPTIONS FOR THE UNIT DETAIL TAB IN CODE ANALYST

The table below describes the options on the shortcut menu for the Unit Details Tab in Code Analyst:

Option	Description
Dependency Details	Displays dependency details.
Show Only Hit Lines	Displays only those lines with time metrics.
Show Only Missed Lines	Displays only those lines without time metrics.
Show All Lines	Resets the view to show all lines.

Option	Description
Advanced View	Displays the default view.
Normal View	Displays a limited number of data columns.

## EXTRACTING SQL TEXT IN CODE ANALYST

The Extract SQL Text functionality extracts SQL text to an ISQL window.

- 1 On the Tools menu, select Code Analyst.
- 2 On the Unit Detail tab, right-click the session, and then select Extract SQL Text.

Code analyst extracts the SQL text to an ISQL window.

For more information, see:

[Extract](#)

[Code Analyst Product Design](#)

[Using the Code Analyst](#)

## EXECUTING SQL IN CODE ANALYST

The Execute SQL functionality extracts SQL text to an ISQL window, and then executes the code.

- 1 On the Tools menu, select Code Analyst.
- 2 On the Unit Detail tab, right-click the session, and then select Execute SQL.

Code analyst extracts the SQL text to an ISQL window and executes the code.

For more information, see:

[Execute](#)

[Code Analyst Product Design](#)

[Using the Code Analyst](#)

# RAPID SQL ADD-ON TOOLS

Rapid SQL includes the following add-on tools:

[Embarcadero SQL Debugger for IBM DB2 LUW](#)

[Embarcadero SQL Debugger for Microsoft SQL Server](#)

[Embarcadero SQL Debugger for Oracle](#)

[Embarcadero SQL Debugger for Sybase ASE](#)

[Rapid SQL PL/SQL Profiler](#)

## EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows lets you locate and fix bugs in procedures and triggers for IBM DB2 LUW for Linux, Unix, and Windows version 7.2 or later. Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows lets you debug triggers by debugging the procedures that call them.

**NOTE:** Availability of this feature depends on your Rapid SQL licensing. For more information, see [Licensing](#).

**TIP:** For Rapid SQL, [Code Analyst](#) is a tool to identify time-consuming lines of code. Code Analyst lets you perform detailed response time analysis on the execution of Procedures and Functions.

The table below describes the sections of this chapter:

Section	Description
<a href="#">Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Features</a>	This section describes how Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows lets you identify problems within your code.
<a href="#">Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Interface</a>	This section describes the Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows graphical interface that includes an editor window and four debug view windows.
<a href="#">Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Functionality</a>	This section describes the functionality on the Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows.
<a href="#">Using the Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows</a>	This section describes how to run a debug session.

## EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS FEATURES

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows lets you identify problems within your code. The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows lets you:

- Interactively step through the flow of script execution.
- Examine the value of variables.
- Solve logical problems with your script design.

**NOTE:** The Debugger is available on the Rapid SQL main menu, the Procedures window, the DDL Editor and ISQL windows.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows offers fundamental debugging features and options to fine tune debugging. The table below describes these features:

Debugging Feature	Description
<a href="#">Step Into</a>	Lets you execute each instruction step-by-step and step inside a stored object.
<a href="#">Step Out</a>	Lets you stop stepping through the current object and execute the remainder of the script. This option is only active when the pointer indicates a child dependent instruction.
<a href="#">Step Over</a>	Lets you execute the current instruction without stepping into any child dependents.
<a href="#">Breakpoints</a>	Lets you specify positions in a program where the debugger stops execution.

To set specific Debugger values on Rapid SQL's Options Editor, see [Debugger Options](#).

## EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS REQUIREMENTS

Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows requires proper configuration of the server and client.

For more information, see:

[Prerequisites for Adding and Compiling Procedures](#)

[Configuring the IBM DB2 LUW for Linux, Unix, and Windows Server for Procedures](#)

[Prerequisites for Debugging Procedures](#)

## PREREQUISITES FOR ADDING AND COMPILING STORED PROCEDURES

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows requires the following products and components.

### Client

- IBM DB2 LUW for Linux, Unix, and Windows 7.2 or later
- DB2 Application Development Client
- DB2 Administration Client
- Communications Protocols
- Stored Procedure Builder
- Applications Development Interfaces
- System Bind Files
- DB2 Connect Server Support
- Documentation
- Base DB2 LUW for Windows/Unix Support
- Administration and Configuration Tools

### Server

- IBM DB2 LUW for Linux, Unix, and Windows 7.2 or later
- DB2 Enterprise Edition
- Communications Protocols
- Stored Procedure Builder
- Applications Development Interfaces
- System Bind Files
- DB2 Connect Server Support
- Documentation
- Base DB2 LUW for Windows/Unix Support
- Administration and Configuration Tools

- Microsoft Visual Studio, Visual C++

**NOTE:** The server must have a local directory structure and file C:\program files\sqllib\function\routine\sr\_cpath.bat. This file is installed with IBM DB2 LUW 7.2 and includes the C compiler options needed to compile the procedure on the server. If it is not found, install the IBM DB2 LUW 7.2 Administration and Configuration Tools option on the server.

## CONFIGURING THE IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS SERVER FOR PROCEDURES

Rapid SQL lets you create procedures on the targeted server using Rapid SQL.

To create or run any procedure, set up the configuration environment and enable the C compiler options on the server.

To configure your server, do the following:

- 1 Open a DB2 Command Window, and then type:

```
DB2set DB2_SQLROUTINE_COMPILER_PATH="C:\program
files\sqllib\function\routine\sr_cpath.bat"
```

DB2 sets the DB2\_SQLROUTINE\_COMPILER\_PATH DB2 registry variable to call the required initialization script for the C compiler on the server.

To enable the C compiler options on your server:

- 1 Open the file C:\program files\sqllib\function\routine\sr\_cpath.bat.
- 2 Remove the REM (remarks) prefix on the lines that match the version of Visual Studio that is installed on the server. VCV6 = version 6.0 and VCV5 = version 5.0.

**NOTE:** Only remove the REM prefix on the lines that apply to your installation of Visual Studio

- 3 Restart the DB2 services on the server.

## PREREQUISITES FOR DEBUGGING PROCEDURES

To enable debugging on the server, do the following:

- 1 Open a DB2 Command window and type:

```
Db2set DB2ROUTINE_DEBUG=ON
```

**NOTE:** Client must have a licensed or evaluation copy of the Embarcadero LUW SQL Debugger.

## EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS OPTIONS

You can specify debugger options from the Debug Tab of Rapid SQL's Options editor. The Debug Tab of the Options Editor lets you set the duration of your debug initialization and debug session, enable DBMS output, and refresh dependencies.



## Setting Debugger Options

To set debugger options, do the following:

- 1 On the *File* menu, click *Options*.

OR

On the *Main* toolbar, click Options.

Rapid SQL opens the Options Editor.

- 2 Specify debugger options. The table below describes the options available:

Option	Description	Default
Initialization Timeout (seconds)	Specifies the number of seconds Rapid SQL tries to initialize the debugger. If it cannot initialize the debugger in the specified time, a message displays in the Debug Output window.	60
Debug Session Timeout (seconds)	Specifies, in seconds, the length of your debug session.	7200
Enable DBMS Output	Toggles the print output. Enable this option if you use <code>dbms_output.put_line</code> calls in your procedures and you want these lines displayed.	Selected
Refresh Dependencies for each run	Refreshes dependencies for each run. This potentially time-consuming process is useful if the target procedure has rapidly varying dependencies that can require updating during the debugging process.	Cleared

- 3 Click Close.

Rapid SQL closes the Options Editor.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Features](#).

## EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS INTERFACE

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows includes an editor window and four debug view windows. When you open a *debug session*, Rapid SQL extracts the code for the object into a DDL Editor and opens four debug view windows at the bottom of the screen. The four debug view windows are optional, dockable windows designed to let you debug your script.

**TIP:** All Embarcadero debuggers display Performance Metrics that let you measure the execution time of each statement in the debug session.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows's includes five windows:

- 1 [DDL Editor window](#)
- 2 [Watch window](#)
- 3 [Variables window](#)
- 4 [Call Stack window](#)
- 5 [Dependency Tree window](#)

### Working with T-SQL Debugger Windows

Rapid SQL lets you resize, move, dock and float the following windows:

- [Watch window](#)
- [Variables window](#)
- [Call Stack window](#)
- [Dependency Tree window](#)

- 1 To resize the target window, click its frame and drag it.

Rapid SQL resizes the window.

- 2 To move and dock the target window, click its grab bar and drag it.

Rapid SQL moves the window to its new location and docks it with surrounding windows.

- 3 To float the target window, press Shift, then click its grab bar and drag it.

Rapid SQL frames the window in its own floating frame and moves the window to its new location.

### DDL EDITOR WINDOW FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

The DDL Editor displays your code in read-only format. When you start debugging, the SQL Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows extracts your code into a DDL Editor. The DDL Editor uses the default Rapid SQL syntax coloring.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Interface](#).

### WATCH WINDOW FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR

## LINUX, UNIX, AND WINDOWS

The Watch window displays the watch variables for the database object you are debugging. The Watch window also lets you specify variables you want to evaluate or modify while debugging your program.

For example, to check what happens when a variable (x) has a value of 100, you can double-click the variable in the DDL Editor, drag it into the Watch Window, and change the value to 100. When you execute the script, the Debugger uses the value x =100. This window is only visible when the T-SQL Debugger is active.

**NOTE:** Until you step at least once into a script, variables are not defined. Therefore, step at least once before dragging or typing a local variable in the Watch Window.

**NOTE:** When you exit a debug session and reenter it, the Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows retains any watch variables or breakpoints you have set.

### Opening and Closing the Watch Window

To open and close the Watch Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views sub-menu*, select or clear *Watch*.

OR

Press ALT+3.

### Setting a Watch Variable

To set a Watch Variable, do the following:

- 1 In the *DDL Editor*, double-click the target variable and drag it to the *Watch window*.
- 2 In the *Watch window*, change the value of the variable.
- 3 On the *DDL Editor*, click *Debug*.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows executes the script using the new variable.

### Removing a Watch Variable

To remove a Watch variable, do the following:

- 1 In the *Watch window*, click the target variable and press *DELETE*.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Interface](#).

## VARIABLES WINDOW FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR

## LINUX, UNIX, AND WINDOWS

The Variables window displays the local variables and their current values during script execution.

**NOTE:** You cannot edit the variables in the Variables window.

If the DDL Editor displays an external database object, and that object is a dependent of the object you are debugging, then the Variables Window automatically refreshes and displays the variables for that particular object. The Variables Window is only visible when the Debugger is active.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows also lets you monitor your variables while debugging.

### Opening and Closing the Variables Window

To open and close the Variables Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views sub-menu*, select or clear *Variable*.

OR

Press ALT+4.

### Monitoring Variables

To monitor the values of your variables while debugging, do the following:

- 1 In the SQL Editor, hold the pointer over the target variable.

Rapid SQL opens a ScreenTip displaying the current value of that variable.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Interface](#).

## CALL STACK WINDOW FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

The Call Stack window displays the stack of currently active calls. The Call Stack Window is only visible when the Debugger is active.

### Opening and Closing the Call Stack Window

To open and close the Call Stack Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views sub-menu*, select or clear *Call Stack*.

OR

Press ALT+5.

## Using the Call Stack Window

To display a line of code that references the call in the DDL Editor, do the following:

- 1 In the Call Stack window, double-click the target line.

In the DDL Editor, Rapid SQL displays a green arrow on the line of the referenced call.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Interface](#).

## DEPENDENCY TREE WINDOW FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

The Dependency Tree window displays any external database objects the script accesses. Rapid SQL displays these database objects in a hierarchical tree, with the child objects as database objects accessed by the parent objects. You can use this window to display the code for a dependent database object in the *DDL Editor* window. This window is only visible when the Debugger is active.

## Opening and Closing the Dependency Tree Window

To open and close the Dependency Tree Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views* sub-menu, select or clear *Dependencies*.

OR

Press ALT+6.

## Displaying Dependencies

To display the code for a dependent database object in the DDL Editor window, do the following:

- 1 In the Dependency Tree window, double-click the target object.

Rapid SQL displays the SQL of the target object in the DDL Editor window.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Interface](#).

## EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS FUNCTIONALITY

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows offers the following functionality:

- [Input Parameters](#)
- [Step Into](#)

- [Step Out](#)
- [Step Over](#)
- [Run to Cursor](#)
- [Insert or Remove a Breakpoint](#)
- [Toggle Breakpoint](#)
- [Go](#)
- [Stop](#)
- [Restart](#)
- [Break](#)
- [Close](#)

To use these functionalities, first [open a debugging session](#).

## INPUT PARAMETERS FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

Input parameters are set when you first create an object. If the object you want to debug requires input parameters, Rapid SQL opens a Procedure Execution dialog box and prompts you for the input parameters when you [open a debugging session](#).

The Procedure Execution dialog box also lets you:

- Save input parameters as \*.prm files to preserve specific input parameter configurations.
- Open \*.prm files to save the effort of reentering specific input parameters.
- Reset parameters to their default setting.

The table below describes the options and functionality on Procedure Execution dialog box:

The following table describes the options available in this dialog box:

Dialog box component	Description
Owner drop-down list	Displays the current procedure's owner
Procedure drop-down list	Displays the name of the current procedure.
Parameter window	Specify the required input parameters in this window. If input parameters are not required for the execution of the target procedure, a message appears in this window, stating that the procedure "has no input parameters. Press execute to run it."
Open button	Click to open an Open dialog box, from which you can open an existing *.prm file. The saved parameters immediately populate the dialog box upon opening.

Dialog box component	Description
Save button	Click to save the values of your input parameters as a *.prm file. You can reopen a saved *.prm file from this dialog box at any time.
Reset button	Click to reset the parameters in the Parameter window to their default values.
Execute button	Click to execute the procedure once you have entered values for all required parameters in the Parameter window.

Option	Description
Owner	Displays the current procedure's owner.
Procedure	Displays the name of the current procedure.
Parameter	Specify the required input parameters in this window. If input parameters are not required for the execution of the target procedure, a message displays in this window, stating that the procedure "has no input parameters. Press execute to run it."
Open	Click to open an existing *.prm file. The saved parameters immediately populate the dialog box upon opening.
Save	Click to save the values of your input parameters as a *.prm file. You can reopen a saved *.prm file from this dialog box at any time.
Reset	Click to reset the parameters in the Parameter window to their default values.
Continue	Click to execute the procedure once you have entered values for all required parameters in the Parameter window.

**NOTE:** You cannot debug a script that requires input parameters until you provide input parameters.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Functionality](#).

## STEP INTO FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

After you [open a debugging session](#), Step Into lets you execute the current instruction. If the current instruction makes a call to a stored SQL object, the Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows steps inside the nested child object.

To use the Step Into facility, do the following:

- 1 On the *Debug* menu, click *Step Into*.

OR

On the *DDL Editor* toolbar, click *Step Into*.

OR

In the DDL Editor window, right-click, and then click *Step Into*.

OR

Press F11.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows moves the arrow to execute the current instruction.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Functionality](#).

## STEP OUT FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

After you [open a debugging session](#), Step Out lets you execute the remainder of the dependent child object and resumes line-by-line, step-debugging in the parent object.

**NOTE:** Step Out is only active when the pointer indicates a child dependent instruction.

To use the Step Out facility, do the following:

- 1 On the *Debug* menu, click *Step Out*.

OR

On the *DDL Editor* toolbar, click *Step Out*.

OR

In the DDL Editor window, right-click, and then click *Step Out*.

OR

Press SHIFT+F11.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows stops stepping through the current object and executes the remainder of the script.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Functionality](#).

## STEP OVER FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX,



## UNIX, AND WINDOWS

After you [open a debugging session](#), Step Over lets you execute the current instruction without stepping into a nested child object if the instruction makes a call to a dependent object.

To use the Step Over, do the following:

- 1 On the *Debug* menu, click *Step Over*.

OR

On the *DDL Editor* toolbar, click *Step Over*.

OR

In the DDL Editor window, right-click, and then click *Step Over*.

OR

Press F10.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows executes the current instruction.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Functionality](#).

## RUN TO CURSOR FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

After you [open a debugging session](#), Run to Cursor lets you execute all instructions between the yellow arrow and the cursor.

To use the Run to Cursor facility, do the following:

- 1 *Scroll down from the yellow arrow to the target line.*
- 2 Click the target line.

Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows places the cursor on the target line.

- 3 On the *Debug* menu, click *Run to Cursor*.

OR

On the *DDL Editor* toolbar, click *Run to Cursor*.

OR

In the DDL Editor window, right-click, and then click *Run to Cursor*.

OR

Press CTRL+F10.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows executes all instructions between the pointer and the cursor.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Functionality](#).

## INSERT OR REMOVE BREAKPOINT FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

A breakpoint is a position in a program where a debugger stops execution. When you start debugging, Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows opens the script in a DDL Editor. A yellow arrow pointer indicates which line the Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows executes next.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows executes all lines of code between the yellow arrow and the first breakpoint. If no breakpoints are present, Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows debugs the entire script.

While debugging you can set one or more breakpoints in the currently executing object or in any object in the program call stack. You can [toggle](#), temporarily disable or enable breakpoints without having to add or remove breakpoints.

Rapid SQL's Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows displays each enabled breakpoint as a red dot in the left margin of the DDL Editor Window, and each disabled breakpoint as a red circle.

Rapid SQL stores all breakpoints you set so that when you debug the same script on separate occasions, you can reuse the same breakpoints. After you [open a debugging session](#), you can insert a breakpoint on the line where your cursor is located, and you can remove a breakpoint on the line where your cursor is located.

**NOTE:** Script execution stops at the first breakpoint.

To insert and remove breakpoints, do the following:

- 1 In the DDL Editor window, click the target line of SQL.
- 2 On the *Debug* menu, click *Breakpoint*.

OR

On the *DDL Editor* toolbar, click *Breakpoint*.

OR

In the *DDL Editor* window, right-click, and then click *Breakpoint*.

OR

Press F9.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows inserts a new breakpoint or removes an existing breakpoint on the target line of code.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Functionality](#).

## TOGGLE BREAKPOINT FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

After you [open a debugging session](#) and [insert a breakpoint](#), Toggle Breakpoint lets you enable or disable that breakpoint. Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows displays each enabled breakpoint as a red dot in the left margin of the DDL Editor Window, and each disabled breakpoint as a red circle.

You can toggle any breakpoint in the DDL Editor window. When you exit a debugging session and reenter it, the Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows retains any breakpoints you set.

To use the Toggle Breakpoint facility, do the following:

- 1 In the DDL Editor window, click the line of the target breakpoint.
- 2 On the *Debug* menu, click *Enable/Disable Breakpoint*.

OR

On the *DDL Editor* toolbar, click *Enable/Disable Breakpoint*.

OR

In the *DDL Editor* window, right-click, and then click *Enable/Disable Breakpoint*.

OR

Press CTRL+F9.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows toggles the breakpoint indicated by the pointer.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Functionality](#).

## GO FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

After you [open a debugging session](#), Go lets you execute all instructions stopping only when when it encounters a breakpoint or when the program is complete.

To use the Go facility, do the following:

- 1 On the *Debug* menu, click *Go*.

OR

On the *DDL Editor* toolbar, click *Go*.

OR

In the *DDL Editor* window, right-click, and then click *Go*.

OR

Press F5.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows executes all instructions.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Functionality](#).

## STOP FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

After you [open a debugging session](#), Stop lets you halt the script execution and terminate the session.

To use the Stop facility, do the following:

- 1 On the *Debug* menu, click *Stop Debugging*.

OR

On the *DDL Editor* toolbar, click *Stop Debugging*.

OR

In the *DDL Editor* window, right-click, and then click *Stop Debugging*.

OR

Press SHIFT+F5.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows stops the script execution and terminates the session.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Functionality](#).

## RESTART FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX,

## UNIX, AND WINDOWS

After you [open a debugging session](#), Restart lets you terminate the current debug session and open a new one. When the new session opens, Rapid SQL prompts you for new input parameters.

To use the Restart facility, do the following:

- 1 On the *Debug* menu, click *Restart*.

OR

On the *DDL Editor* toolbar, click Restart.

OR

In the *DDL Editor* window, right-click, and then click *Restart*.

OR

Press CTRL+SHIFT+F5.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows restarts the debug session.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Functionality](#).

## BREAK FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

After you [open a debugging session](#), Break lets you pause the debug session.

To use the Break facility, do the following:

- 1 On the *Debug* menu, click *Break*.

OR

On the *DDL Editor* toolbar, click Break.

OR

In the *DDL Editor* window, right-click, and then click *Break*.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows suspends the debug session.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Functionality](#).

## CLOSE FOR EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX,

## AND WINDOWS

After you [open a debugging session](#), Close lets you close the DDL Editor and the Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows.

To use the Close facility, do the following:

- 1 On the *DDL Editor* toolbar, click Close.

OR

In the upper right corner of the window, click Close.

OR

In the *DDL Editor* window, right-click, and then click *Close*.

The Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows closes the debug session.

For more information, see [Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows Functionality](#).

## USING THE EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

This section offers a general overview of how to use Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows's full range of debugging [functionality](#). After you [open a debugging session](#) for any procedure or trigger, you can [begin debugging](#).

### OPENING A DEBUGGING SESSION IN IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

When you open a debugging session, Rapid SQL opens the [five windows](#) of the Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows interface. If the target script requires input parameters, Rapid SQL opens a Procedure Execution dialog box and prompts you for the necessary input parameters before displaying the target code in the SQL Editor window. When Rapid SQL displays the target script in the SQL Editor window, you can [begin debugging](#).

**NOTE:** Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows only lets you debug the SQL script of procedures or triggers.

To debug triggers by debugging the procedures that call them, do the following:

- 1 On the Explorer Tab, click the node of the target trigger or procedure. The node expands and displays the Code object.
- 2 Double-click Code. The DDL Editor opens and displays the code of the target object.

- 3 On the Debug menu, click Start Debugging.

OR

On the SQL Editor toolbar, click Debug.

OR

In the *DDL Editor* window, right-click, and then click *Debug*.

OR

Press CTRL+F5.

- 4 On the toolbar, click Debug.

OR

In the *DDL Editor* window, right-click, and then click *Debug*.

If the script requests input parameters, Rapid SQL opens a Procedure Execution dialog box. If the script does not require input parameters, Rapid SQL displays the script in the DDL Editor window for you to [begin debugging](#).

**NOTE:** You cannot use the Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows until it has fully initialized.

- 5 In the Procedure Execution dialog box, specify the appropriate parameters, and then click *Continue*.

Rapid SQL displays the script in the DDL Editor window for you to [begin debugging](#).

## DEBUGGING AN SQL SCRIPT WITH EMBARCADERO SQL DEBUGGER FOR IBM DB2 LUW FOR LINUX, UNIX, AND WINDOWS

After you [open a debugging session](#) and enter any required input parameters, you can begin working with your script in the Embarcadero SQL Debugger for IBM DB2 LUW for Linux, Unix, and Windows.

### Debugging an SQL Script

To debug a SQL Script, do the following:

- 1 On the *Debug* menu, click one of the T-SQL Debugger options ([Step Into](#), [Step Over](#), and so forth) or click *Go*.

OR

On the *DDL Editor* toolbar, click one of the T-SQL Debugger options ([Step Into](#), [Step Over](#), and so on) or click *Go*.

**NOTE:** You can monitor the progress of your debug session in the Variables window.

- 2 On the *Debug* menu, click *Breakpoint*.

OR

On the *DDL Editor* toolbar, click *Breakpoint*.

OR

Press F9.

**NOTE:** When you set a breakpoint, the Call Stack window shows what was called before the breakpoint.

**NOTE:** You can use the [Run to Cursor](#) option to test the lines of code between a breakpoint and your cursor (indicated by the yellow arrow in the DDL Editor).

- 3 To check your variables:

- 1) In the **DDL Editor**, click a variable in your script and drag it to the **Watch** window.
- 2) In the **Watch** window, change the value of the watch variable, and then click **Go** to run your script and see the results of the new value.

- 4 To check a record in stored objects:

- 1) Drag the record to the **Watch** window.
- 2) In the **Watch** window, change the value of the record, then click **Go** to run your script and see the results of the new value.

- 5 To check the dependencies:

- 1) In the **Dependency Tree** window double-click the target dependent object to extract the code into a new **DDL Editor**.
- 2) **Step through** the script while monitoring the [Dependency Tree window](#).

- 6 When you finish debugging the script, click *Close*.

Rapid SQL closes the T-SQL Debugger DDL Editor.

**NOTE:** When you exit a debug session and reenter it, the Embarcadero SQL Debugger for MSSQL retains any watch variables or breakpoints you have set.

## EMBARCADERO SQL DEBUGGER FOR MICROSOFT SQL SERVER

Embarcadero SQL Debugger for Microsoft is a programming tool that helps you locate and fix bugs in Microsoft SQL Server procedures and triggers for Microsoft SQL Server version 7.0 or later.

**NOTE:** Availability of this feature depends on your Rapid SQL licensing. For more information, see [Licensing](#).

### Objects

You can use Embarcadero SQL Debugger for Microsoft to debug the following objects:



- Procedures
- Triggers

You can only debug triggers by debugging the procedures that call them.

**TIP:** The '[Code Analyst](#)' is a tool to identify time-consuming lines of code. Code Analyst lets you perform detailed response time analysis on the execution of Procedures and Functions.

The table below describes the sections of this chapter:

Section	Description
<a href="#">Embarcadero SQL Debugger for Microsoft Features</a>	This section describes how Embarcadero SQL Debugger for Microsoft helps you identify problems within your code.
<a href="#">Embarcadero SQL Debugger for Microsoft Interface</a>	This section describes the Embarcadero SQL Debugger for Microsoft graphical interface that includes an editor window and four debug view windows.
<a href="#">Embarcadero SQL Debugger for Microsoft Functionality</a>	This section describes the way in which Embarcadero SQL Debugger for Microsoft functions.
<a href="#">Using the Embarcadero SQL Debugger for Microsoft</a>	This section describes how to run a debug session.

## EMBARCADERO SQL DEBUGGER FOR MICROSOFT FEATURES

The Embarcadero SQL Debugger for Microsoft helps you identify problems within your code. The Embarcadero SQL Debugger for Microsoft lets you:

- Interactively step through the flow of script execution.
- Examine the value of variables.
- Solve logical problems with your script design.

The Embarcadero SQL Debugger for Microsoft offers fundamental debugging features and several options to help fine tune debugging, as listed in the table below:

Debugging Feature	Description
<a href="#">Step Into</a>	Lets you execute each instruction step-by-step and step inside a stored object.
<a href="#">Step Out</a>	Lets you stop stepping through the current object and execute the remainder of the script. This option is only active when the pointer indicates a child dependent instruction.
<a href="#">Step Over</a>	Lets you execute the current instruction without stepping into any child dependents.

Debugging Feature	Description
<a href="#">Breakpoints</a>	A position in a program where the debugger stops execution.

To set specific Debugger values on Rapid SQL's Options Editor, see [Debugger Options](#).

## EMBARCADERO SQL DEBUGGER FOR MICROSOFT REQUIREMENTS

To use the Embarcadero SQL Debugger for Microsoft, you must properly configure the server and client. For more information, see:

[Server Requirements](#)

[Client Requirements](#)

### SERVER REQUIREMENTS

To use the Embarcadero SQL Debugger for Microsoft you must be running Windows 2000 or Windows NT 4.0 or later, and your Microsoft SQL Server version must be 7.0 or later.

### Setting Up the Server

There are three parts to setting up the server:

- [Installing the Microsoft SQL Debugger Interface subcomponent.](#)
- [Configuring the Service.](#)
- [Configuring DCOM on the server](#)

### Enabling SQL Debugger for Microsoft on SQL Server SP3

SQL Debugging is disabled by default in SQL Server SP3 and greater. Please refer to [Microsoft Support](#) for information regarding enabling the SQL Debugger for Microsoft on SQL Server SP3.

### INSTALLING THE MICROSOFT SQL DEBUGGER INTERFACE SUBCOMPONENT

The Microsoft server must have the Development Tools, Debugger Interface subcomponent of Microsoft SQL Server 7.0. To determine if the Debugger Interface subcomponent is installed, locate the following files in the \Program Files\Common Files\Microsoft Shared\SQL Debugging directory:

- SQLDBREG.exe
- SQLDBG.dll

If these files are not in the \Program Files\Common Files\Microsoft Shared\SQL Debugging directory, you must install them before running the Embarcadero SQL Debugger for Microsoft.

## Installing the Microsoft SQL Debugger Interface on the Server

To install the Debugger Interface subcomponent on the server after the initial installation, do the following:

- 1 Start Microsoft Visual Studio, Enterprise Edition Setup.

OR

Start Microsoft SQL Server 7.0 Setup.

- 2 Select Custom Install.

Microsoft SQL Server opens the Select Components dialog box:

- 3 In the Components box, select the Development Tools check box.
- 4 In the Sub-components box, select the Debugger Interface check box.
- 5 Click Next.

Microsoft SQL Server proceeds through the Microsoft SQL Server wizard to install the components.

### CONFIGURING THE SERVICE

Configuring the service is an operating-system-dependent operation. See the instructions below for your server operating system.

- [Windows 2000](#)
- [Windows NT 4.0](#)

### Windows 2000

- 1 On the Windows taskbar, click the Start button, click Settings, and then click Control Panel.
- 2 Double-click Administrative Tools, and then click Services.  
Windows opens the Services explorer.
- 3 In the right pane of the Services explorer, right click MSSQLServer, and then click Properties.  
Windows opens the Net Logon Properties dialog box.
- 4 Click the Logon Tab.
- 5 Select the This Account option button.
- 6 In the This Account box, type (or browse to locate) the logon user account (including domain name, if necessary) of the person using the Embarcadero SQL Debugger for Microsoft.  
**NOTE:** This person must have admin permissions on the server.
- 7 In the Password and Confirm Password boxes, type the password.
- 8 Click Apply.

- 9 Click the General Tab.
- 10 Click Start.

Windows starts the server and applies the changes.

### Important Notes about Microsoft SQL Server 2000 Service Pack 3 (SP3)

By default, after you install Microsoft SQL Server 2000 Service Pack 3 (SP3), you cannot use the Embarcadero SQL Debugger for Microsoft. You may receive the following error message:

"Server: Msg 514, Level 16, State 1, Procedure sp\_sdebug, Line 1 [Microsoft][ODBC SQL Server Driver][SQL Server]Unable to communicate with debugger on [SQL Server Name] (Error = 0x80070005). Debugging disabled for connection 53."

Microsoft made this change for security reasons. To enable debugging, a member of the sysadmins server role, such as sa, must explicitly enable debugging by running the following code:

```
Exec sp_sdebug 'legacy_on'
```

You must repeat this procedure whenever you restart the server.

### Windows NT 4.0

- 1 On the Windows taskbar, click the Start button, select Settings and then click Control Panel.
- 2 Double-click Services.  
Windows opens the Services dialog box.
- 3 In the Service list, select MSSQLServer and then click Startup.
- 4 In the Log On As: box, select the This Account option button.
- 5 In the This Account box, type the logon user account (including domain name, if necessary) of the person using the Embarcadero SQL Debugger for Microsoft.

**NOTE:** This person must have admin permissions on the server.

- 6 In the Password and Confirm Password boxes, type the password.
- 7 Click OK.

Windows returns to the Services dialog box.

- 8 Click Start.

Windows starts the server and applies the changes.

**CONFIGURING DCOM ON THE SERVER**

To configure DCOM, do the following:

- 1 After the server restarts, on the Windows taskbar, click the Start button, and then click Run.
- 2 In the Open box, type dcomcnfg.exe.
- 3 Click OK.

Windows opens the Distributed COM Configuration Properties dialog box.

- 4 Click the Default Security Tab.
- 5 In the Default Access Permissions box, click Edit Default.

Windows opens the Registry Value Permissions dialog box.

- 6 Click Add.

Windows opens the Add Users and Groups dialog box.

- 7 In the Names box, select SYSTEM, and then click Add.
- 8 Click the Type of Access list and then click Allow Access.
- 9 To let any user use the Embarcadero SQL Debugger for Microsoft, you must grant them remote access on the server. To grant remote access, you must configure their DCOM permissions on the server. In the Names box, click the target users and then click Add.

**NOTE:** You can add individual users or groups.

- 10 Click the Type of Access list and then click Allow Access.
- 11 Click OK.
- 12 Restart the server to apply the changes.

**CLIENT REQUIREMENTS**

There are three categories of client requirements for the Embarcadero SQL Debugger for Microsoft:

- Operating System
- Microsoft SQL Server 7.0, Client Connectivity Component
- Microsoft SQL Server 7.0, Development Tools - Debugger Interface Subcomponent

**Operating System**

The client must be running one of the following operating systems:

- Microsoft Windows 95
- Microsoft Windows 98

- Microsoft Windows NT 4.0 or later

### Important Notes about Microsoft SQL Server 2000 Service Pack 3 (SP3)

By default, after you install Microsoft SQL Server 2000 Service Pack 3 (SP3), you cannot use the Embarcadero SQL Debugger for Microsoft. You may receive the following error message:

“Server: Msg 514, Level 16, State 1, Procedure sp\_sdidebug, Line 1 [Microsoft][ODBC SQL Server Driver][SQL Server]Unable to communicate with debugger on [SQL Server Name] (Error = 0x80070005). Debugging disabled for connection 53.”

Microsoft made this change for security reasons. To enable debugging, a member of the sysadmins server role, such as sa, must explicitly enable debugging by running the following code:

```
Exec sp_sdidebug 'legacy_on'
```

You must repeat this procedure whenever you restart the server.

### Client Connectivity

The client must have the Client Connectivity component of Microsoft SQL Server 7.0 or later.

### Microsoft Debugger Interface

The client must have the Development Tools, Debugger Interface subcomponent of Microsoft SQL Server 7.0 or later. To determine if the Debugger Interface subcomponent is installed, locate the following files in the \Program Files\Common Files\Microsoft Shared\SQL Debugging directory:

- SQLDBREG.exe
- SQLDBG.dll

If these files are not in the \Program Files\Common Files\Microsoft Shared\SQL Debugging directory, you must install them before running the Embarcadero SQL Debugger for Microsoft.

### Installing the Microsoft SQL Debugger Interface on the Client

To install the Debugger Interface subcomponent on the client:

- 1 Start the Microsoft SQL Server Setup program.
- 2 Select Custom Install.  
Microsoft SQL Server opens the Select Components dialog box.
- 3 In the Components box, select the Development Tools check box.
- 4 In the Sub-Components box, select the Debugger Interface check box.

- 5 Click Next.

Microsoft SQL Server proceeds through the Microsoft SQL Server Wizard to install the components.

## EMBARCADERO SQL DEBUGGER FOR MICROSOFT OPTIONS

You can specify T-SQL Debugger options from the Debug Tab of Rapid SQL's Options editor. The Debug Tab of the Options Editor lets you set the duration of your debug initialization and debug session, enable DBMS output, and refresh dependencies.

### Setting Debugger Options

To set debugger options, do the following:

- 1 On the *File* menu, click *Options*.

OR

On the *Main* toolbar, click Options.

Rapid SQL opens the Options Editor.

- 2 Specify debugger options. The table below describes the options available:

Option	Description	Default
Initialization Timeout (seconds)	Specifies the number of seconds Rapid SQL tries to initialize the debugger. If it cannot initialize the debugger in the specified time, a message displays in the Debug Output window.	60
Debug Session Timeout (seconds)	Specifies, in seconds, the length of your debug session.	7200
Enable DBMS Output	Toggles the print output. Enable this option if you use <code>dbms_output.put_line</code> calls in your procedures and you want these lines displayed.	Selected
Refresh Dependencies for each run	Refreshes dependencies for each run. This potentially time-consuming process is useful if the target procedure has rapidly varying dependencies that can require updating during the debugging process.	Cleared

- 3 Click Close.

Rapid SQL closes the Options Editor.

## EMBARCADERO SQL DEBUGGER FOR MICROSOFT INTERFACE

The Embarcadero SQL Debugger for Microsoft has a graphical interface that includes an editor window and four debug view windows. When you open a *debug session*, Rapid SQL extracts the code for the object into a DDL Editor and opens four debug view windows at the bottom of the screen. The four debug view windows are optional, dockable windows designed to help you debug your script.

**TIP:** All Embarcadero debuggers display Performance Metrics that let you measure the execution time of each statement in the debug session.

Embarcadero SQL Debugger for Microsoft's five windows are:

- 1 [DDL Editor window](#)
- 2 [Watch window](#)
- 3 [Variables window](#)
- 4 [Call Stack window](#)
- 5 [Dependency Tree window](#)

### Working with T-SQL Debugger Windows

Rapid SQL lets you resize, move, dock and float the following T-SQL Debugger windows:

- Watch
- Variables
- Call Stack
- Dependency Tree

To work with the above windows, do the following:

- 1 To resize the target window, click its frame and drag it  
Rapid SQL resizes the window.
- 2 To move and dock the target window, click its grab bar and drag it.  
Rapid SQL moves the window to its new location and docks it with surrounding windows.
- 3 To float the target window, press Shift, then click its grab bar and drag it.  
Rapid SQL frames the window in its own floating frame and moves the window to its new location.



## DDL EDITOR WINDOW

The Embarcadero SQL Debugger for Microsoft provides a DDL Editor that displays your code in read-only format. When you start debugging, the SQL Embarcadero SQL Debugger for Microsoft extracts your code into a DDL Editor. The DDL Editor uses the default Rapid SQL syntax coloring.

## WATCH WINDOW

The Rapid SQL Embarcadero SQL Debugger for Microsoft provides a watch window that displays the watch variables for the database object you are debugging and lets you specify variables you want to evaluate or modify while debugging your program. For example, to check what happens when a variable (x) has a value of 100, you can double-click the variable in the DDL Editor, drag it into the Watch Window, and change the value to 100. When you execute the script, the Debugger uses the value x =100. This window is only visible when the T-SQL Debugger is active.

**NOTE:** Until you step at least once into a script, variables are not defined. Therefore, you must step at least once before dragging or typing a local variable in the Watch Window.

**NOTE:** When you exit a debug session and reenter it, the Embarcadero SQL Debugger for Microsoft retains any watch variables or breakpoints you have set.

### Opening and Closing the Watch Window

To open and close the Watch Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views* sub-menu, select or clear *Watch*.

OR

Press ALT+3.

### Setting a Watch Variable

To set a Watch Variable, do the following:

- 1 In the *DDL Editor*, double-click the target variable and drag it to the *Watch window*.

**NOTE:** Microsoft SQL Server requires that local variables begin with @. You must drag the @ to the Watch Window.

- 2 In the *Watch window*, change the value of the variable.
- 3 On the *DDL Editor*, click *Go*.

The Embarcadero SQL Debugger for Microsoft executes the script using the new variable.

### Removing a Watch Variable

To remove a Watch Variable, do the following:

- 1 In the *Watch window*, click the target variable and press *DELETE*.

### VARIABLES WINDOW

The Embarcadero SQL Debugger for Microsoft provides a variables window that displays the local variables and their current values during script execution. You cannot edit the variables in the Variables window. If the DDL Editor displays an external database object, and that object is a dependent of the object you are debugging, then the Variables Window automatically refreshes and displays the variables for that particular object. The Variables Window is only visible when the Debugger is active.

The Embarcadero SQL Debugger for Microsoft also lets you monitor your variables while debugging.

### Opening and Closing the Variables Window

To open and close the Variables Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views sub-menu*, select or clear *Variable*.

OR

Press ALT+4.

### Monitoring Variables

To monitor the values of your variables while debugging, do the following:

- 1 In the SQL Editor, hold the pointer over the target variable.

Rapid SQL opens a ScreenTip displaying the current value of that variable.

### CALL STACK WINDOW

The Embarcadero SQL Debugger for Microsoft provides a call stack window that displays the stack of currently active calls. The Call Stack Window is only visible when the Debugger is active.

### Opening and Closing the Call Stack Window

To open and close the Call Stack Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views sub-menu*, select or clear *Call Stack*.

OR

Press ALT+5.

## Using the Call Stack Window

To display a line of code that references the call in the DDL Editor, do the following:

- 1 In the Call Stack window, double-click the target line.

In the DDL Editor, Rapid SQL displays a green arrow on the line of the referenced call.

## DEPENDENCY TREE WINDOW

The Embarcadero SQL Debugger for Microsoft provides a dependency tree window that displays any external database objects the script accesses. Rapid SQL displays these database objects in a hierarchical tree, with the child objects as database objects accessed by the parent objects. You can use this window to display the code for a dependent database object in the *DDL Editor* window. This window is only visible when the Debugger is active.

## Opening and Closing the Dependency Tree Window

To open and close the Dependency Tree Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views* sub-menu, select or clear *Dependencies*.

OR

Press ALT+6.

## Displaying Dependencies

To display the code for a dependent database object in the DDL Editor window, do the following:

- 1 In the Dependency Tree window, double-click the target object.

Rapid SQL displays the SQL of the target object in the DDL Editor window.

## EMBARCADERO SQL DEBUGGER FOR MICROSOFT FUNCTIONALITY

The Embarcadero SQL Debugger for Microsoft offers the following functionality:

- [Input Parameters](#)
- [Step Into](#)
- [Step Out](#)
- [Step Over](#)
- [Run to Cursor](#)
- [Insert or Remove a Breakpoint](#)
- [Toggle Breakpoint](#)

- [Go](#)
- [Stop](#)
- [Restart](#)
- [Break](#)
- [Close](#)

To make use of the above functionality, you must first [open a debugging session](#).

## INPUT PARAMETERS

Input parameters are set when you first create an object. If the object you want to debug requires input parameters, Rapid SQL opens a Procedure or Trigger Execution dialog box and prompts you for the input parameters when you [open a debugging session](#).

This dialog box also lets you:

- Save input parameters as \*.prm files to preserve specific input parameter configurations.
- Open \*.prm files to save the effort of reentering specific input parameters.
- Reset parameters to their default setting.

The table below describes the options and functionality on Procedure Execution dialog box:

The following table describes the options available in this dialog box:

Dialog box component	Description
Owner drop-down list	Displays the current procedure's owner
Procedure drop-down list	Displays the name of the current procedure.
Parameter window	Specify the required input parameters in this window. If input parameters are not required for the execution of the target procedure, a message appears in this window, stating that the procedure "has no input parameters. Press execute to run it."
Open button	Click to open an Open dialog box, from which you can open an existing *.prm file. The saved parameters immediately populate the dialog box upon opening.
Save button	Click to save the values of your input parameters as a *.prm file. You can reopen a saved *.prm file from this dialog box at any time.
Reset button	Click to reset the parameters in the Parameter window to their default values.
Execute button	Click to execute the procedure once you have entered values for all required parameters in the Parameter window.

Option	Description
Owner	Displays the current procedure's owner.
Procedure	Displays the name of the current procedure.
Parameter	Specify the required input parameters in this window. If input parameters are not required for the execution of the target procedure, a message displays in this window, stating that the procedure "has no input parameters. Press execute to run it."
Open	Click to open an existing *.prm file. The saved parameters immediately populate the dialog box upon opening.
Save	Click to save the values of your input parameters as a *.prm file. You can reopen a saved *.prm file from this dialog box at any time.
Reset	Click to reset the parameters in the Parameter window to their default values.
Continue	Click to execute the procedure once you have entered values for all required parameters in the Parameter window.

**NOTE:** You cannot debug a script that requires input parameters until you provide input parameters.

## STEP INTO

After you [open a debugging session](#), Step Into lets you execute the current instruction. If the current instruction makes a call to a stored SQL object, the Embarcadero SQL Debugger for Microsoft steps inside the nested child object.

### Step Into

To use the Step Into facility, do the following:

- 1 On the *Debug* menu, click *Step Into*.

OR

On the *DDL Editor* toolbar, click Step Into.

OR

In the DDL Editor window, right-click and then click *Step Into*.

OR

Press F11.

The Embarcadero SQL Debugger for Microsoft moves the arrow to execute the current instruction.

## STEP OUT

After you [open a debugging session](#), Step Out lets you execute the remainder of the dependent child object and resumes line-by-line, step-debugging in the parent object.

**NOTE:** Step Out is only active when the pointer indicates a child dependent instruction.

### Step Out

To use the Step Out facility, do the following:

- 1 On the *Debug* menu, click *Step Out*.

OR

On the *DDL Editor* toolbar, click Step Out.

OR

In the DDL Editor window, right-click and then click *Step Out*.

OR

Press SHIFT+F11.

The Embarcadero SQL Debugger for Microsoft stops stepping through the current object and executes the remainder of the script.

## STEP OVER

After you [open a debugging session](#), Step Over lets you execute the current instruction without stepping into a nested child object if the instruction makes a call to a dependent object.

### Step Over

To use the Step Over facility, do the following:

- 1 On the *Debug* menu, click *Step Over*.

OR

On the *DDL Editor* toolbar, click Step Over.

OR

In the DDL Editor window, right-click and then click *Step Over*.

OR

Press F10.

The Embarcadero SQL Debugger for Microsoft executes the current instruction.

## RUN TO CURSOR

After you [open a debugging session](#), Run to Cursor lets you execute all instructions between the yellow arrow and the cursor.

### Run to Cursor

To use the Run to Cursor facility, do the following:

- 1 Scroll down from the yellow arrow to the target line.
- 2 Click the target line.

Embarcadero SQL Debugger for Microsoft places the cursor on the target line.

- 3 On the *Debug* menu, click *Run to Cursor*.

OR

On the *DDL Editor* toolbar, click Run to Cursor.

OR

In the DDL Editor window, right-click and then click *Run to Cursor*.

OR

Press CTRL+F10.

The Embarcadero SQL Debugger for Microsoft executes all instructions between the pointer and the cursor.

## INSERT OR REMOVE BREAKPOINT

A breakpoint is a position in a program where a debugger stops execution. When you start debugging, Embarcadero SQL Debugger for Microsoft opens the script in a DDL Editor. A yellow arrow pointer indicates which line the Embarcadero SQL Debugger for Microsoft executes next. The Embarcadero SQL Debugger for Microsoft executes all lines of code between the yellow arrow and the first breakpoint. If no breakpoints are present, Embarcadero SQL Debugger for Microsoft debugs the entire script.

While debugging you can set one or more breakpoints in the currently executing object or in any object in the program call stack. Breakpoints can be [toggled](#), temporarily disabled, or enabled, without having to add or remove them. Rapid SQL Embarcadero SQL Debugger for Microsoft displays each enabled breakpoint as a red dot in the left margin of the DDL Editor Window, and each disabled breakpoint as a red circle.

Rapid SQL stores all breakpoints you set so that when you debug the same script on separate occasions, you can reuse the same breakpoints. After you [open a debugging session](#), Insert Break lets you insert a breakpoint on the line where your cursor is located, and Remove Break lets you remove a breakpoint on the line where your cursor is located.

**NOTE:** Script execution stops at the first breakpoint.

**Inserting or Removing a Breakpoint**

To insert or remove a breakpoint, do the following:

- 1 In the DDL Editor window, click the target line of SQL.
- 2 On the *Debug* menu, click *Breakpoint*.

OR

On the *DDL Editor* toolbar, click *Breakpoint*.

OR

In the *DDL Editor* window, right-click and then click *Breakpoint*.

OR

Press F9.

The Embarcadero SQL Debugger for Microsoft inserts a new breakpoint or removes an existing breakpoint on the target line of code.

**TOGGLE BREAKPOINT**

After you [open a debugging session](#) and [insert a breakpoint](#), Toggle Breakpoint lets you enable or disable that breakpoint. Embarcadero SQL Debugger for Microsoft displays each enabled breakpoint as a red dot in the left margin of the DDL Editor Window, and each disabled breakpoint as a red circle. You can toggle any breakpoint in the DDL Editor window. When you exit a debugging session and reenter it, the Embarcadero SQL Debugger for Microsoft retains any breakpoints you set.

**Toggling a Breakpoint**

To toggle a breakpoint, do the following:

- 1 In the DDL Editor window, click the line of the target breakpoint.
- 2 On the *Debug* menu, click *Enable/Disable Breakpoint*.

OR

On the *DDL Editor* toolbar, click *Enable/Disable Breakpoint*.

OR

In the *DDL Editor* window, right-click and then click *Enable/Disable Breakpoint*.

OR

Press CTRL+F9.

The Embarcadero SQL Debugger for Microsoft toggles the breakpoint indicated by the pointer.



## Go

After you [open a debugging session](#), Go lets you execute all instructions stopping only when it encounters a breakpoint or when the program is complete.

### Go

To use the Go facility, do the following:

- 1 On the *Debug* menu, click *Go*.  
OR  
On the *DDL Editor* toolbar, click *Go*.  
OR  
In the *DDL Editor* window, right-click and then click *Go*.  
OR  
Press F5.

The Embarcadero SQL Debugger for Microsoft executes all instructions.

## STOP

After you [open a debugging session](#), Stop lets you halt the script execution and terminate the session.

### Stop

To stop the debugger, do the following:

- 1 On the *Debug* menu, click *Stop Debugging*.  
OR  
On the *DDL Editor* toolbar, click *Stop Debugging*.  
OR  
In the *DDL Editor* window, right-click and then click *Stop Debugging*.  
OR  
Press SHIFT+F5.

The Embarcadero SQL Debugger for Microsoft stops the script execution and terminates the session.

## RESTART

After you [open a debugging session](#), Restart lets you terminate the current debug session and open a new one. When the new session opens, Rapid SQL prompts you for new input parameters.

### Restart

To restart the debugger, do the following:

- 1 On the *Debug* menu, click *Restart*.

OR

On the *DDL Editor* toolbar, click *Restart*.

OR

In the *DDL Editor* window, right-click and then click *Restart*.

OR

Press CTRL+SHIFT+F5.

The Embarcadero SQL Debugger for Microsoft restarts the debug session.

## BREAK

After you [open a debugging session](#), Break lets you pause the debug session.

### Break

To pause the debugger, do the following:

- 1 On the *Debug* menu, click *Break*.

OR

On the *DDL Editor* toolbar, click *Break*.

OR

In the *DDL Editor* window, right-click and then click *Break*.

The Embarcadero SQL Debugger for Microsoft suspends the debug session.

## CLOSE

After you [open a debugging session](#), Close lets you close the DDL Editor and the Embarcadero SQL Debugger for Microsoft.

**Close**

- 1 On the *DDL Editor* toolbar, click Close.

OR

In the upper right corner of the window, click Close.

OR

In the *DDL Editor* window, right-click and then click *Close*.

The Embarcadero SQL Debugger for Microsoft closes the debug session.

## USING THE EMBARCADERO SQL DEBUGGER FOR MICROSOFT

This section offers a general overview of how to use Embarcadero SQL Debugger for Microsoft's full range of debugging [functionality](#). After you [open a debugging session](#) for any procedure or trigger, you can [begin debugging](#).

For more detailed information, see [Debugging a Sample Script](#).

### OPENING A DEBUGGING SESSION

When you open a debugging session, Rapid SQL opens the [five windows](#) of the Embarcadero SQL Debugger for Microsoft interface. If the target script requires input parameters, Rapid SQL opens a Procedure Execution dialog box and prompts you for the necessary input parameters before displaying the target code in the SQL Editor window. When Rapid SQL displays the target script in the SQL Editor window, you can [begin debugging](#).

**NOTE:** Rapid SQL Embarcadero SQL Debugger for Microsoft only lets you debug the SQL script of procedures or triggers.

### Opening a Debugging Session

To debug a trigger or procedure, do the following:

- 1 On the *Explorer* Tab, click the node of the target procedure.  
Rapid SQL opens the node and displays two items: Code and Privileges.
- 2 Under the target object node, double-click Code.  
Rapid SQL opens a DDL Editor *displaying the code of the target object*.

- 3 On the *Debug* menu, click *Start Debugging*.

OR

On the SQL Editor toolbar, click *Debug*.

OR

In the *DDL Editor* window, right-click and then click *Debug*.

OR

Press CTRL+F5.

If the script requests input parameters, Rapid SQL opens a Procedure Execution dialog box. If the script does not require input parameters, Rapid SQL displays the script in the DDL Editor window for you to [begin debugging](#).

**NOTE:** You cannot use the Embarcadero SQL Debugger for Microsoft until it has fully initialized.

- 4 In the Procedure Execution dialog box, specify the appropriate parameters, and then click *Continue*.

Rapid SQL displays the script in the DDL Editor window for you to [begin debugging](#).

## DEBUGGING AN SQL SCRIPT

After you [open a debugging session](#) and enter any required input parameters, you can begin working with your script in the Embarcadero SQL Debugger for Microsoft.

### Debugging an SQL Script

To debug a SQL script, do the following:

- 1 On the *Debug* menu, click one of the Embarcadero SQL Debugger for MSSQL Server options ([Step Into](#), [Step Over](#), and so forth) or click *Go*.

OR

On the *DDL Editor* toolbar, click one of the Embarcadero SQL Debugger for MSSQL Server options ([Step Into](#), [Step Over](#), and so on) or click *Go*.

**NOTE:** You can monitor the progress of your debug session in the Variables window.

2 On the *Debug* menu, click *Breakpoint*.

OR

On the *DDL Editor* toolbar, click *Breakpoint*.

OR

Press F9.

**NOTE:** When you set a breakpoint, the Call Stack window shows what was called before the breakpoint.

**NOTE:** You can use the [Run to Cursor](#) option to test the lines of code between a breakpoint and your cursor (indicated by the yellow arrow in the DDL Editor).

3 To check your variables, do the following:

- 1) In the **DDL Editor**, click a variable in your script and drag it to the **Watch** window.
- 2) In the **Watch** window, change the value of the watch variable and then click **Go** to run your script and see the results of the new value.

4 To check a record in stored objects, do the following:

- 1) Drag the record to the **Watch** window.
- 2) In the **Watch** window, change the value of the record, then click **Go** to run your script and see the results of the new value.

5 To check the dependencies, do the following:

- 1) In the **Dependency Tree** window double-click the target dependent object to extract the code into a new **DDL Editor**.
- 2) **Step through** the script while monitoring the [Dependency Tree window](#).

6 When you finish debugging the script, click Close.

Rapid SQL closes the T-SQL Debugger DDL Editor.

**NOTE:** When you exit a debug session and reenter it, the Embarcadero SQL Debugger for MSSQL retains any watch variables or breakpoints you have set.

## DEBUGGING A SAMPLE SCRIPT

This walk-through demonstrates basic debugging functionality. During the course of this walk-through you debug two procedures using the Embarcadero SQL Debugger for Microsoft.

This section is divided into the following seven sections, each designed to familiarize you with basic debugging features and functionality:

- [Getting Started](#)
- [Testing a Procedure](#)
- [Starting the Debugging Session](#)

- [Breakpoints](#)
- [Step Into](#)
- [Step Out](#)
- [Correcting the Script](#)

## GETTING STARTED

This part of Debugging the Sample Script explains how to create the following two procedures to be used for debugging:

- check\_modulo
- calculate\_sum\_with\_overflow\_bug

**NOTE:** The procedure calculate\_sum\_with\_overflow\_bug intentionally includes a bug which prevents it from executing successfully. You use the Embarcadero SQL Debugger for Microsoft to identify this bug.

### Overview

The Getting Started section guides you through:

- [Creating procedure 1.](#)
- [Creating procedure 2.](#)
- [Confirming the creation of the procedures.](#)

### CREATING PROCEDURE 1

Procedure 1, check\_modulo, calculates the modulo of any two user-specified numbers. The user passes the numbers into the procedure as input parameters. The procedure returns the result as an output parameter. If the modulo equals zero, procedure execution returns the output "YES". If the modulo is not zero, procedure execution returns the output "NO". This procedure is nested in the second procedure, calculate\_sum\_with\_overflow\_bug.

To create this procedure, you must open Rapid SQL, connect to a MSSQL datasource, open a new SQL editor and, in the SQL editor, type or copy and paste the following code:

```
CREATE PROCEDURE username.check_modulo
@p_dividend_in INT,
@p_divisor_in INT,
@result VARCHAR(3)OUTPUT
AS
IF @p_dividend_in % @p_divisor_in = 0
    SELECT @result = 'YES'
ELSE
    SELECT @result = 'NO'
go
```

**NOTE:** For the purposes of this walk-through, this procedure was created under the user name Spence. Before executing the DDL above, substitute your user name for the word "username".

## Creating Procedure 1

To create Procedure 1, do the following:

- 1 Start Rapid SQL.
- 2 Connect to a Microsoft SQL Server 7.0 datasource.
- 3 On the Datasource menu, click the database node and then click the target database.

**NOTE:** For this walk-through, we recommend that you select a non-production database.

- 4 On the Main toolbar, click New.

OR

On the File menu, click New, and then click SQL.

OR

Press CTRL+N.

Rapid SQL opens an SQL Editor in the current workspace.

- 5 In the SQL Editor, type the DDL for procedure check\_modulo.

**NOTE:** You must substitute your user name once in the DDL for this procedure.

- 6 On the SQL Editor toolbar, click Execute.

Rapid SQL executes the script and creates Procedure 1, then opens the SQL Editor Results Tab with the results of the script execution. If you were not able to create the procedure, check the error messages to determine the problem.

## CREATING PROCEDURE 2

Procedure 2, calculate\_sum\_with\_overflow\_bug, requires two user-specified numbers as input parameters. Upon execution, the procedure calculates the sum of the all numbers divisible by five between the two user-specified numbers. This procedure calls sample procedure 1 (check\_modulo) to calculate the modulo of the user-specified numbers.

**NOTE:** The procedure calculate\_sum\_with\_overflow\_bug intentionally includes a bug which prevents it from executing successfully. You use the Embarcadero SQL Debugger for Microsoft to identify this bug.

**CAUTION:** When inputting parameters, you must enter the smaller number in the @p\_num1\_in int box.

To create this procedure, you must open Rapid SQL, connect to a MSSQL datasource, open a new SQL editor and, in the SQL editor, type or copy and paste the following code:

```
CREATE PROCEDURE username.calculate_sum_with_overflow_bug
@p_num1_in INT,
@p_num2_in INT,
@result TINYINT OUTPUT
/*INT-Integer (whole number) data from -2^31 (-2,147,483,648)
through 2^31 - 1 (2,147,483,647).
TINYINT-Integer data from 0 through 255.*/
AS
DECLARE @temp INT
DECLARE @temp_1 INT
DECLARE @v_divisor INT
DECLARE @v_condition VARCHAR(3)

SET @temp = @p_num1_in
SET @temp_1 = 0
SET @v_divisor = 5
SET @v_condition = 'NO'

WHILE 1=1
BEGIN

    SELECT @temp = @temp + 1 /*Increase temp starting from p_num1*/

    IF @temp = @p_num2_in /*Check if we reached p_num2*/
        /*If yes, leave the LOOP*/
        BREAK

    /*Call Procedure 2 to check if number is divisable by 5*/
    EXEC username.check_modulo @temp,@v_divisor,@result=@v_condition output

    IF @v_condition = 'YES'
        SELECT @temp_1 = @temp_1 + @temp

END /*WHILE LOOP*/

SELECT @result = @temp_1

RETURN
go
```

**NOTE:** For the purposes of this walk-through, this procedure was created under the user name Spence. Before executing the DDL above, substitute your user name for the word "username".

## Creating Procedure 2

To create Procedure 2, do the following:

- 1 Start Rapid SQL.
- 2 Connect to a Microsoft SQL Server 7.0 datasource.
- 3 On the Datasource menu, click the database node and then click the target database.

**NOTE:** For this walk-through, we recommend that you select a non-production database.



- 4 On the Main toolbar, click New.

OR

On the File menu, click New, and then click SQL.

OR

Press CTRL+N.

Rapid SQL opens an SQL Editor in the current workspace.

- 5 In the SQL Editor, type the DDL for procedure `calculate_sum_with_overflow_bug`.

**NOTE:** You must substitute your user name twice in the DDL for this procedure.

- 6 On the SQL Editor toolbar, click Execute.

Rapid SQL executes the script and creates Procedure 2, then opens the SQL Editor Results Tab with the results of the script execution. If you were not able to create the procedure, check the error messages to determine the problem.

#### CONFIRMING THE CREATION OF THE PROCEDURES

After you create [Procedure 1](#) and [Procedure 2](#), you can confirm their creation in Rapid SQL's Database Explorer.

#### Confirming the Creation of the Procedures

To confirm creation of the procedures, do the following:

- 1 On the Explorer Tab, click the Explorer list, and then click Organize By Owner.

The Explorer Tab refreshes with the new display configuration.

- 2 On the Explorer Tab, double-click the Databases node, and then double-click the target database node.

Rapid SQL displays the list of object owners.

- 3 Double-click your user name to display a list of your objects.
- 4 Double-click Procedures to display a list of procedures and confirm the creation of `check_modulo` and `calculate_sum_with_overflow_bug`.

You are now ready to begin [testing a procedure](#).

## TESTING A PROCEDURE

After you [confirm the creation of the procedures](#), you must execute the procedure `calculate_sum_with_overflow_bug` (which includes a bug) to view its error message. This procedure requires two integer input parameters: `@p_num1_in int` and `@p_num2_in int`. For all integers between these two integers, this procedure identifies those divisible by 5, and then returns their sum.

**CAUTION:** When inputting parameters, you must enter the smaller number in the `@p_num1_in int` box.

### Testing a Procedure

To test a procedure, do the following:

- 1 On the Explorer Tab, right-click `calculate_sum_with_overflow_bug`, and then click Execute.  
Rapid SQL opens the Procedure Execution window.
- 2 In the Value column of the `@p_num1_in` row, type 1.
- 3 In the Value column of the `@p_num2_in` row, type 11.
- 4 Click Execute.

Rapid SQL compiles the procedure and opens a Results Tab, displaying the sum 15. There are two numbers between 1 and 11 that are divisible by 5: 5, and 10. The sum of these two numbers is 15.

- 5 On the Explorer Tab, right-click `calculate_sum_with_overflow_bug`, and then click Execute.  
Rapid SQL again opens the Procedure Execution window.
- 6 In the Value column of the `@p_num1_in` row, type 100.
- 7 In the Value column of the `@p_num2_in` row, type 121.
- 8 On the Procedure Execution window toolbar, click Execute.

Rapid SQL returns an error stating "Arithmetic overflow occurred". You are now ready to [Start the Debugging Session](#).

## STARTING THE DEBUGGING SESSION

After you [test the procedure](#), you must open the procedure in Embarcadero SQL Debugger for Microsoft and enter input parameters before debugging. To start a session, do the following:

- 1 On the Explorer Tab, right-click the procedure, `calculate_sum_with_overflow_bug`, and then click Debug to start the debug session.  
  
Rapid SQL extracts the DDL for the procedure into a DDL Editor and opens the Procedure Execution dialog box.
- 2 In the Value column of the `@p_num1_in` row, type 100.

- 3 In the Value column of the @p\_num2\_in row, type 121.
- 4 Click Continue.

Rapid SQL closes the dialog box and opens the Embarcadero SQL Debugger interface, which includes the following five windows:

- [DDL Editor](#)
- [Watch Window](#)
- [Variables Window](#)
- [Call Stack Window](#)
- [Dependency Tree Window](#)

You are now ready to [insert breakpoints](#).

## BREAKPOINTS

After you [start the debugging session](#), you must insert a breakpoint into the code of the procedure calculate\_sum\_with\_overflow\_bug. Then you must run to the breakpoint. After you run to the breakpoint, Embarcadero SQL Debugger displays a yellow arrow on the red breakpoint icon and populates the Variables Window with values for the following variables:

Variable	Value
@temp	Current number
@p_num2_in	Second input parameter
@p_num1_in	First input parameter
@temp_1	Sum of the numbers, between the input parameters, divisible by 5
@result	Condition of the output parameter
@v_condition	Output parameter
@v_divisor	Divisor

## Breakpoints

To insert a breakpoint, do the following:

- 1 In the DDL Editor, scroll to and click the following line:

```
EXEC username.check_modulo @temp,@v_divisor,@result=@v_condition output
```

**NOTE:** This line is located near the end of the procedure's code.

- 2 On the *Debug* menu, click *Breakpoint*.

OR

On the *DDL Editor* toolbar, click *Breakpoint*.

OR

In the *DDL Editor* window, right-click and then click *Breakpoint*.

OR

Press F9.

Rapid SQL inserts a breakpoint (indicated by dot) next to the number of the target line.

- 3 On the *Debug* menu, click *Go*.

OR

On the *DDL Editor* toolbar, click *Go*.

OR

In the *DDL Editor* window, right-click and then click *Go*.

OR

Press F5.

Rapid SQL Embarcadero SQL Debugger for Microsoft displays the value of the variables before the breakpoint in the Variables Window.

You are now ready to [Step Into](#) the code.

## STEP INTO

After setting the [breakpoint](#), you must step into the dependent procedure, `check_modulo`. To step into the dependent procedure, do the following:

- 1 On the *Debug* menu, click *Step Into*.

OR

On the *DDL Editor* toolbar, click *Step Into*.

OR

In the *DDL Editor* window, right-click and then click *Step Into*.

OR

Press F11.

Rapid SQL extracts the DDL for the dependent, nested procedure into the DDL Editor.

## 2 Step Into again.

Rapid SQL executes the next part of the code and displays the values for the variables in the Variables Window.

The Call Stack Window displays calls to the procedures.

You are now ready to [Step Out](#) of the code.

## STEP OUT

After you [Step Into](#) the modulo\_check (nested procedure) code, you must step back out and return to the calculate\_sum\_with\_overflow\_bug (outside procedure) code. To step back out and return, do the following:

### 1 On the *Debug* menu, click *Step Out*.

OR

On the *DDL Editor* toolbar, click *Step Out*.

OR

In the *DDL Editor* window, right-click and then click *Step Out*.

OR

Press SHIFT+F11.

Rapid SQL opens the *DDL Editor* containing the code for calculate\_sum\_with\_overflow\_bug.

### 2 On the *Debug* menu, click *Go*.

OR

On the *DDL Editor* toolbar, click *Go*.

OR

In the *DDL Editor* window, right-click and then click *Go*.

OR

Press F5.

When the value of the variable, @temp is equal to the value of the variable, @p\_num2\_in, the WHILE LOOP is complete and the Embarcadero SQL Debugger for Microsoft continues to the next executable statement in the code.

### 3 While monitoring the value of the variables in the Variables Window, continue to click *Go* to cycle through the WHILE LOOP.

After executing the SELECT and RETURN statements, Rapid SQL closes the Debugger and opens a *DDL Editor* to the Results Tab.

Now you are ready to [correct the script](#).

## CORRECTING THE SCRIPT

When you finished [Stepping Out](#) of the nested code and encounter the error, you must do the following to fully fix the bug:

- Locate the source of the error
- Scroll to the line in the script displaying the error
- Analyze the code
- Correct the error
- Compile the corrected script

When you first executed the procedure, Rapid SQL displayed the error message "Arithmetic overflow error for data type tinyint, value = 450". According to *Microsoft SQL Server Books Online*: "This error occurs when an attempt is made to convert a float or real data type value into a data type that cannot store the result. This error prevents the operation from being completed."

The data type used in this procedure (TINYINT) stores values from 0 to 255. The sum of the four numbers between 100 and 121 that are divisible by 5 (105, 110, 115, and 120) is 450. But because the TINYINT variable @result can only accept a maximum value of 255, Rapid SQL returns the error message and the procedure fails.

### Correcting the Script

To correct the script, do the following:

- 1 On the Explorer Tab, right-click calculate\_sum\_with\_overflow\_bug, and then click Extract.  
Rapid SQL extracts the DDL for the procedure into a DDL Editor.
- 2 On the Edit toolbar, click Find.  
Rapid SQL opens the Find dialog box.
- 3 In the Find What box, type TINYINT.
- 4 Click Find Next.  
Rapid SQL selects the first occurrence of TINYINT.
- 5 Change the data type for @result from TINYINT to INT.
- 6 On the DDL Editor toolbar, click Execute to execute the modified script.  
Rapid SQL executes the script and opens the Results Tab.
- 7 On the Explorer Tab, right-click calculate\_sum\_with\_overflow\_bug, and then click Execute.  
Rapid SQL opens the Procedure Execution dialog box.
- 8 In the Value column of the @p\_num1\_in row, type 100.
- 9 In the Value column of the @p\_num2\_in row, type 121.

10 Click Execute.

Rapid SQL executes the procedure with the new data type and opens the Results Tab, returning the value 450. You successfully corrected the script and debugged the procedure.

## EMBARCADERO SQL DEBUGGER FOR ORACLE

Embarcadero SQL Debugger for Oracle is a programming tool that lets you debug functions, procedures and triggers for Oracle versions 7.3.3 or later.

### Objects

Using Embarcadero SQL Debugger for Oracle, you can debug the following objects:

- Functions
- Procedures
- Triggers

You can only debug triggers by debugging the functions or procedures that call them. You cannot debug packages, but you can debug the functions and procedures within packages.

**NOTE:** You cannot debug any objects contained in the [Exclusion List](#).

**NOTE:** Availability of this feature depends on your Rapid SQL licensing. For more information, see [Licensing](#).

**TIP:** The [Code Analyst](#) is a tool to identify time-consuming lines of code. Code Analyst lets you perform detailed response time analysis on the execution of Procedures and Functions.

The table below describes the sections of this chapter:

Section	Description
<a href="#">Features</a>	This section describes how the Embarcadero SQL Debugger for Oracle helps you identify problems within your code.
<a href="#">Interface</a>	This section describes the Embarcadero SQL Debugger for Oracle graphical interface that includes an editor window and four debug view windows.
<a href="#">Functionality</a>	This section describes the functions of the Embarcadero SQL Debugger for Oracle.
<a href="#">Using Embarcadero SQL Debugger for Oracle</a>	This section describes how to run a debug session.

## DEBUGGING FEATURES

Embarcadero SQL Debugger for Oracle is designed to help identify problems within your code. Embarcadero SQL Debugger for Oracle lets you:

- Interactively step through the flow of script execution.
- Examine the value of variables.
- Solve logical problems with your script design.

Embarcadero SQL Debugger for Oracle offers fundamental debugging features, an Oracle [Exclusion List](#) and several options to help fine tune debugging, as listed in the table below:

Debugging Feature	Description
<a href="#">Step Into</a>	Lets you execute each instruction step-by-step and step inside a stored object if the object is not on the <a href="#">Exclusion List</a> .
<a href="#">Step Out</a>	Lets you stop stepping through the current object and execute the remainder of the script. This option is only active when the pointer indicates a child-dependent instruction.
<a href="#">Step Over</a>	Lets you execute the current instruction without stepping into any child dependents.
<a href="#">Breakpoints</a>	A position in a program where the debugger stops execution.

To set specific Debugger values on the Options Editor, see [Debugger Options](#).

## EXCLUSION LIST

Upon installation, Rapid SQL sets up an Exclusion List on your computer which includes packages that the application cannot debug. The Exclusion List is located in the Rapid SQL directory, at the default installation location C:\Program Files\Embarcadero\Nov2001Shared\deborcex.etd. You can add or remove packages from this file by editing the Exclusion List.

### Editing the Exclusion List

To Edit the Exclusion List, do the following:

- 1 Open the Exclusion List, *deborcex.etd*, in a text editor, such as Notepad or WordPad.
- 2 To add a package, enter the name of the package at the end of the list. Use the following format: OWNER.OBJECT\_NAME.

**NOTE:** There must be a carriage return after each item on the list.



- 3 To remove a package from the Exclusion List, delete the package from the list.

**NOTE:** Embarcadero SQL Debugger for Oracle does debug a package procedure listed on the Exclusion List.

- 4 Save the changes to *deborcex.etd*.

## EMBARCADERO SQL DEBUGGER FOR ORACLE OPTIONS

You can specify PL/SQL Debugger options from the Debug Tab of Rapid SQL's Options editor. The Debug Tab of the Options Editor lets you set the duration of your debug initialization and debug session, enable DBMS output, and refresh dependencies.

### Setting Debugger Options

To set debugger options, do the following:

- 1 On the *File* menu, click *Options*.

OR

On the *Main* toolbar, click Options.

Rapid SQL opens the Options Editor.

- 2 On the Debug Tab, specify debugger options. The table below describes the options available:

Option	Description	Default
Initialization Timeout (seconds)	Specifies the number of seconds Rapid SQL tries to initialize the debugger. If it cannot initialize the debugger in the specified time, a message displays in the Debug Output window.	60
Debug Session Timeout (seconds)	Specifies, in seconds, the length of your debug session.	7200
Enable DBMS Output	Toggles the print output. Enable this option if you use <code>dbms_output.put_line</code> calls in your procedures and you want these lines displayed.	Selected
Refresh Dependencies for each run	Refreshes dependencies for each run. This potentially time-consuming process is useful if the target procedure has rapidly varying dependencies that can require updating during the debugging process.	Cleared

- 3 Click Close.

Rapid SQL closes the Options Editor.

## EMBARCADERO SQL DEBUGGER FOR ORACLE INTERFACE

Embarcadero SQL Debugger for Oracle has a graphical interface that includes an editor window and four debug view windows. When you open a debug session, Rapid SQL extracts the code for the object into an SQL Editor and opens four debug view windows at the bottom of the screen.

**TIP:** All Embarcadero debuggers display Performance Metrics that let you measure the execution time of each statement in the debug session.

The four debug view windows are optional, dockable, floatable windows designed to help debug your script. Embarcadero SQL Debugger for Oracle's five windows are:

- [SQL Editor window](#)
- [Watch window](#)
- [Variables window](#)
- [Call Stack window](#)
- [Dependency Tree window](#)

### Working with Embarcadero SQL Debugger Windows

Rapid SQL lets you resize, move, dock and float the following Debugger windows:

- Watch
- Variables
- Call Stack
- Dependency Tree

To work with the above windows, do the following:

- 1 To resize the target window, click its frame and drag it  
Rapid SQL resizes the window.
- 2 To move and dock the target window, click its grab bar and drag it.  
Rapid SQL moves the window to its new location and docks it with surrounding windows.
- 3 To float the target window, press Shift, then click its grab bar and drag it.  
Rapid SQL frames the window in its own floating frame and moves the window to its new location.

## SQL EDITOR WINDOW

Embarcadero SQL Debugger for Oracle provides an SQL Editor window that displays your code in Read-Only format. When you start debugging, Embarcadero SQL Debugger for Oracle extracts your code into an SQL Editor window, making it editable. The SQL Editor uses the default Rapid SQL syntax coloring.

**NOTE:** Rapid SQL displays LOB datatypes, and REF CURSOR variables, in the Results Tab.

## WATCH WINDOW

Embarcadero SQL Debugger for Oracle provides a Watch window that displays the watch variables for the database object you are debugging and lets you specify variables you want to evaluate or modify while debugging your program. For example, to check what happens when a variable (x) has a value of 100, you can double-click that variable in the SQL Editor, drag it into the Watch window, and change the value to 100. When you execute the script, the debugger uses the value x=100. This window is only visible when the PL/SQL Debugger is active.

**NOTE:** You can type a fully qualified record variable into the Watch window.

**NOTE:** When you exit a debug session and reenter it, the Embarcadero SQL Debugger for Oracle retains any watch variables or breakpoints you have set.

### Opening and Closing the Watch Window

To open and close the Watch Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views sub-menu*, select or clear *Watch*.
- OR
- Press ALT+3.

### Setting a Watch Variable

To set a Watch Variable, do the following:

- 1 In the *SQL Editor*, double-click the target variable and drag it to the *Watch window*.
- 2 In the *Watch window*, change the value of the variable.
- 3 On the *SQL Editor* toolbar, click *Go*.

Rapid SQL executes the script using the new value of the variable.

### Removing a Watch Variable

- 1 In the *Watch window*, click the target variable and press *DELETE*.

## VARIABLES WINDOW

Embarcadero SQL Debugger for Oracle provides a Variables window that displays the local variables and their current values during script execution. You cannot edit variables in the Variables window. This window is only visible when the Debugger is active. If the SQL Editor displays an external database object, and that object is a dependent of the object you are debugging, then the Variables window automatically refreshes and displays the variables for that particular object.

The Embarcadero SQL Debugger for Oracle also lets you monitor the value of your variables while debugging.

### Opening and Closing the Variables Window

To open and close the Variables Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views sub-menu*, select or clear *Variable*.  
OR  
Press ALT+4.

## Monitoring Variables

To monitor the values of your variables while debugging, do the following:

- 1 In the SQL Editor, hold the pointer over the target variable.  
Rapid SQL opens a ScreenTip displaying the current value of that variable.

## CALL STACK WINDOW

Embarcadero SQL Debugger for Oracle provides a Call Stack window that displays the stack of currently active function calls. The Call Stack window is only visible when the PL/SQL Debugger is active.

### Opening and Closing the Call Stack Window

To open and close the Call Stack Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views sub-menu*, select or clear *Call Stack*.  
OR  
Press ALT+5.

## Using the Call Stack Window

To display a line of code that references the call in the *SQL Editor*, do the following:

- 1 In the *Call Stack window*, double-click the target line.

In the *SQL Editor*, Rapid SQL displays a green arrow on the line of the referenced call.

## DEPENDENCY TREE WINDOW

Embarcadero SQL Debugger for Oracle provides a Dependency Tree window that displays any external database objects the script accesses. Rapid SQL displays these database objects in a hierarchical tree, with the child objects as database objects accessed by the parent objects. You can use this window to display the code for a dependent database object in the *SQL Editor* window. This window is only visible when the Debugger is active.

## Opening and Closing the Dependency Tree Window

To open and close the Dependency Tree Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views* sub-menu, select or clear *Dependencies*.

OR

Press ALT+6.

## Displaying Dependencies

To display the code for a dependent database object in the *SQL Editor* window, do the following:

- 1 In the *Dependency Tree* window, double-click the target object.

Rapid SQL displays the SQL of the target object in the *SQL Editor* window.

## EMBARCADERO SQL DEBUGGER FOR ORACLE FUNCTIONALITY

Embarcadero SQL Debugger for Oracle offers you the following functionality:

- [Input Parameters](#)
- [Step Into](#)
- [Step Out](#)
- [Step Over](#)
- [Run to Cursor](#)
- [Insert or Remove Breakpoint](#)
- [Toggle Breakpoint](#)

- [Go](#)
- [Stop](#)
- [Restart](#)
- [Break](#).
- [Close](#)

To make use of the above functionality, you must first [open a debugging session](#).

## INPUT PARAMETERS

Input parameters are set when you first create an object. If the object you want to debug requires input parameters, Rapid SQL opens a Function, Procedure, or Trigger Execution dialog box and prompts you for the input parameters when you [open a debugging session](#).

This dialog box also lets you:

- Save input parameters as \*.prm files to preserve specific input parameter configurations.
- Open \*.prm files to save the effort of reentering specific input parameters.
- Reset parameters to their default setting.

The table below describes the options and functionality on Procedure Execution dialog box:

The following table describes the options available in this dialog box:

Dialog box component	Description
Owner drop-down list	Displays the current procedure's owner
Procedure drop-down list	Displays the name of the current procedure.
Parameter window	Specify the required input parameters in this window. If input parameters are not required for the execution of the target procedure, a message appears in this window, stating that the procedure "has no input parameters. Press execute to run it."
Open button	Click to open an Open dialog box, from which you can open an existing *.prm file. The saved parameters immediately populate the dialog box upon opening.
Save button	Click to save the values of your input parameters as a *.prm file. You can reopen a saved *.prm file from this dialog box at any time.
Reset button	Click to reset the parameters in the Parameter window to their default values.
Execute button	Click to execute the procedure once you have entered values for all required parameters in the Parameter window.

Option	Description
Owner	Displays the current procedure's owner.
Procedure	Displays the name of the current procedure.
Parameter	Specify the required input parameters in this window. If input parameters are not required for the execution of the target procedure, a message displays in this window, stating that the procedure "has no input parameters. Press execute to run it."
Open	Click to open an existing *.prm file. The saved parameters immediately populate the dialog box upon opening.
Save	Click to save the values of your input parameters as a *.prm file. You can reopen a saved *.prm file from this dialog box at any time.
Reset	Click to reset the parameters in the Parameter window to their default values.
Continue	Click to execute the procedure once you have entered values for all required parameters in the Parameter window.

**NOTE:** You cannot debug a script that requires input parameters until you provide input parameters.

## STEP INTO

After you [open a debugging session](#), *Step Into* lets you execute the current instruction. If the current instruction makes a call to a stored Oracle object, Embarcadero SQL Debugger for Oracle steps inside the nested child object.

**NOTE:** Oracle 7.3 has problems running the debugger on an object with cursors.

## Step Into

To use the Step Into facility, do the following:

- 1 On the *Debug* menu, click *Step Into*.

OR

On the *SQL Editor* toolbar, click Step Into.

OR

In the *SQL Editor* window, right-click and then click *Step Into*.

OR

Press F11.

Embarcadero SQL Debugger for Oracle moves the arrow to execute the current instruction.

## STEP OUT

After you [open a debugging session](#), Step Out lets you execute the remainder of the dependent child object and resumes line-by-line, step-debugging in the parent object.

**NOTE:** Step Out is only active when the pointer indicates a child-dependent instruction.

**NOTE:** Oracle 7.3 has problems running the debugger on an object with cursors.

## Step Out

To use the Step Out facility, do the following:

- 1 On the *Debug* menu, click *Step Out*.

OR

On the *SQL Editor* toolbar, click Step Out.

OR

In the *SQL Editor* window, right-click and then click *Step Out*.

OR

Press SHIFT+F11.

Embarcadero SQL Debugger for Oracle stops stepping through the current object and executes the remainder of the script.

## STEP OVER

After you [open a debugging session](#), Step Over lets you execute the current instruction without stepping into a nested child object if the instruction makes a call to a dependent object.

**NOTE:** Oracle 7.3 has problems running the debugger on an object with cursors.



**Step Over**

To use the Step Over facility, do the following:

- 1 On the *Debug* menu, click *Step Over*.

OR

On the *SQL Editor* toolbar, click *Step Over*.

OR

In the *SQL Editor* window, right-click and then click *Step Over*.

OR

Press F10.

Embarcadero SQL Debugger for Oracle executes the current instruction.

**RUN TO CURSOR**

After you [open a debugging session](#), Run to Cursor lets you execute all instructions between the yellow arrow and your cursor.

**Run to Cursor**

To use the Run to Cursor facility, do the following:

- 1 Scroll down from the yellow arrow to the target line.

- 2 Click the target line.

Embarcadero SQL Debugger for Oracle places the cursor on the target line.

- 3 On the *Debug* menu, click *Run to Cursor*.

OR

On the *SQL Editor* toolbar, click *Run to Cursor*.

OR

In the *SQL Editor* window, right-click and then click *Run to Cursor*.

OR

Press CTRL+F10.

Embarcadero SQL Debugger for Oracle executes all instructions between the pointer and your cursor.

## INSERT OR REMOVE A BREAKPOINT

A breakpoint is a position in a program where a debugger stops execution. When you start debugging, Embarcadero SQL Debugger for Oracle opens the script in an SQL Editor window. A yellow pointer indicates which line the Debugger executes next. Embarcadero SQL Debugger for Oracle executes all lines of code between the pointer and the first breakpoint. If no breakpoints are present, Embarcadero SQL Debugger for Oracle debugs the entire script.

While debugging, you can set one or more breakpoints in the currently executing object or in any object in the program call stack. Breakpoints can be [toggled](#), temporarily disabled, or enabled, without having to add or remove them. Embarcadero SQL Debugger for Oracle displays each enabled breakpoint as a red dot in the left margin of the SQL Editor Window, and each disabled breakpoint as a red circle.

Rapid SQL stores all breakpoints you set, so that when you debug the same script on separate occasions, you can reuse the same breakpoints.

After you [open a debugging session](#), you can insert a breakpoint on the line where your cursor is located, and you can remove a breakpoint on the line where your cursor is located.

**NOTE:** Script execution stops at the first breakpoint.

### Inserting or Removing a Breakpoint

To insert and remove a breakpoint, do the following:

- 1 In the SQL Editor window, click the target line of SQL.
- 2 On the *Debug* menu, click *Breakpoint*.

OR

On the *SQL Editor* toolbar, click Breakpoint.

OR

In the *SQL Editor* window, right-click and then click *Breakpoint*.

OR

Press F9.

Embarcadero SQL Debugger for Oracle inserts a new breakpoint or removes an existing breakpoint on the target line of code.

## TOGGLE BREAKPOINT

After you [open a debugging session](#) and [insert a breakpoint](#), Toggle Breakpoint lets you enable or disable that breakpoint. Embarcadero SQL Debugger for Oracle displays each enabled breakpoint as a red dot in the left margin of the SQL Editor Window, and each disabled breakpoint as a red circle. You can toggle any breakpoint in the SQL Editor window. When you exit a debugging session and reenter it, the Embarcadero SQL Debugger for Oracle retains any breakpoints you set.

## Toggling a Breakpoint

To toggle a breakpoint, do the following:

- 1 In the SQL Editor window, click the line of the target breakpoint.
- 2 On the *Debug* menu, click *Enable/Disable Breakpoint*.

OR

On the *SQL Editor* toolbar, click *Enable/Disable Breakpoint*.

OR

In the *SQL Editor* window, right-click and then click *Enable/Disable Breakpoint*.

OR

Press CTRL+F9.

Embarcadero SQL Debugger for Oracle toggles the breakpoint indicated by the pointer.

## Go

After you [open a debugging session](#), Go lets you execute all instructions, stopping only when it encounters a breakpoint or when the program is complete.

## Go

Go use the Go facility, do the following:

- 1 On the *Debug* menu, click *Go*.

OR

On the *SQL Editor* toolbar, click *Go*.

OR

In the *SQL Editor* window, right-click and then click *Go*.

OR

Press F5.

Embarcadero SQL Debugger for Oracle executes all instructions.

## STOP

After you [open a debugging session](#), Stop lets you halt the script execution and terminate the session.

## Stop

To stop the debugger, do the following:

- 1 On the *Debug* menu, click *Stop Debugging*.

OR

On the *SQL Editor* toolbar, click *Stop Debugging*.

OR

In the *SQL Editor* window, right-click and then click *Stop Debugging*.

OR

Press SHIFT+F5.

Embarcadero SQL Debugger for Oracle stops the script execution and terminates the session.

## RESTART

After you [open a debugging session](#), Restart lets you terminate the current debug session and open a new one. When the new session opens, Rapid SQL prompts you for new input parameters.

### Restart

To restart the debugger, do the following:

- 1 On the *Debug* menu, click *Restart*.

OR

On the *SQL Editor* toolbar, click Restart.

OR

In the *SQL Editor* window, right-click and then click *Restart*.

OR

Press CTRL+SHIFT+F5.

Embarcadero SQL Debugger for Oracle restarts the debug session.

## BREAK

After you [open a debugging session](#), Break lets you pause the debug session.

**Break**

To pause the session, do the following:

- 1 On the *Debug* menu, click *Break*.

OR

On the *SQL Editor* toolbar, click *Break*.

OR

In the *SQL Editor* window, right-click and then click *Break*.

Embarcadero SQL Debugger for Oracle suspends the debug session.

**CLOSE**

After you [open a debugging session](#), Close lets you close the SQL Editor and the Embarcadero SQL Debugger for Oracle.

**Close**

To close the SQL Editor and the debugger, do the following:

- 1 On the *SQL Editor* toolbar, click *Close*.

OR

In the upper right corner of the window, click *Close*.

OR

In the *SQL Editor* window, right-click and then click *Close*.

Embarcadero SQL Debugger for Oracle closes the debug session.

## USING THE EMBARCADERO SQL DEBUGGER FOR ORACLE

This section offers a general overview of how to use Embarcadero SQL Debugger for Oracle's full range of debugging [functionality](#). After you [open a debugging session](#) for any Oracle procedure or function, you can [begin debugging](#).

For more detailed information, see [Debugging a Sample Script](#).

## OPENING A DEBUGGING SESSION

When you open a debugging session, Rapid SQL opens the [five windows](#) of the Embarcadero SQL Debugger for Oracle interface. If the target script requires input parameters, Rapid SQL opens a Procedure or Function Execution dialog box and prompts you for the necessary input parameters before displaying the target code in the SQL Editor window. When Rapid SQL displays the target script in the SQL Editor window, you can [begin debugging](#).

**NOTE:** Embarcadero SQL Debugger for Oracle only lets you debug the SQL script of functions, triggers and procedures.

### Opening a Debugging Session

To debug a function, trigger or procedure, do the following:

- 1 On the *Explorer* Tab, click the node of the target function, trigger, or procedure.

Rapid SQL opens the node and displays two items: Code and Privileges.

- 2 Under the target object node, double-click Code.

Rapid SQL opens an SQL editor displaying the code of the target object.

- 3 On the Debug menu, click Start Debugging.

OR

On the SQL Editor toolbar, click Debug.

OR

In the SQL Editor window, right-click and then click *Debug*.

OR

Press CTRL+F5.

If the script requests input parameters, Rapid SQL opens the Procedure or Function Execution dialog box. If the script does not require input parameters, Rapid SQL displays the script in the SQL Editor window for you to [begin debugging](#).

**NOTE:** You cannot use the Embarcadero SQL Debugger for Oracle until it has fully initialized.

- 4 In the Procedure or Function Execution dialog box, specify the appropriate parameters, and then click Continue.

Rapid SQL displays the script in the SQL Editor window for you to [begin debugging](#).

**NOTE:** If the script requires Oracle types (tables, records, or Booleans) as input parameters, the PL/SQL Debugger generates an anonymous block.

## DEBUGGING AN SQL SCRIPT

After you [open a debugging session](#) and enter any required input parameters, you can begin working with your script in the Debugger.

## Debugging an SQL Script

To debug an SQL script, do the following:

- 1 On the *Debug* menu, click one of the PL/SQL Debugger options ([Step Into](#), [Step Over](#), and so forth) or click *Go*.

OR

On the *SQL Editor* toolbar, click one of the PL/SQL Debugger options ([Step Into](#), [Step Over](#), and so on) or click *Go*.

**NOTE:** You can monitor the progress of your debug session in the Variables window.

- 2 On the *Debug* menu, click *Breakpoint*.

OR

On the *SQL Editor* toolbar, click *Breakpoint*.

OR

Press F9.

**NOTE:** When you set a breakpoint, the Call Stack window shows what was called before the breakpoint.

**NOTE:** You can use the [Run to Cursor](#) option to test the lines of code between a breakpoint and your cursor (indicated by the yellow arrow in the SQL Editor).

- 3 To check your variables, do the following:

- 1) In the **SQL Editor**, click a variable in your script and drag it to the **Watch** window.
- 2) In the **Watch** window, change the value of the watch variable and then click **Go** to run your script and see the results of the new value.

- 4 To check record in stored objects, do the following:

- 1) Drag the record to the **Watch** window.
- 2) In the **Watch** window, change the value of the record, then click **Go** to run your script and see the results of the new value.

- 5 To check the dependencies, do the following:

- 1) In the **Dependency Tree** window double-click the target dependent object to extract the code into a new **SQL Editor**.
- 2) **Step through** the script while monitoring the [Dependency Tree window](#).

- 6 When you finish debugging the script, click *Close*.

Rapid SQL closes the PL/SQL Debugger SQL Editor.

**NOTE:** When you exit a debug session and reenter it, the Embarcadero SQL Debugger for Oracle retains any watch variables or breakpoints you have set.

## DEBUGGING A SAMPLE SCRIPT

The Rapid SQL installation includes a sample script intended to walk you through basic debugging functionality. The sample script creates a package that includes functions and procedures that you debug.

**NOTE:** To create the sample package, you must have CREATE privileges.

### Overview

Debugging a Sample Script is divided into three sections that will familiarize you with basic debugging features and functionality. These sections are:

- [Getting Started](#), which guides you through creating the package you will use in Debugging Sample Script 1 and Debugging Sample Script 2.
- [Debugging Sample Script 1](#), which guides you through debugging functionality and demonstrates the Embarcadero SQL Debugger for Oracle interface features.
- [Debugging Sample Script 2](#), which guides you through debugging functionality and error correction.

**NOTE:** For the purposes of this walk-through we have created this package under the user name DEMO\_SPENCE.

## GETTING STARTED

The Rapid SQL installation includes a sample script that you execute to create a package containing functions and procedures. These functions and procedures demonstrate basic debugging features available in the Embarcadero SQL Debugger for Oracle

**NOTE:** To create the sample package, you must have CREATE privileges.

The Rapid SQL installation places the script in the C:\Program Files\Embarcadero\RSQL600\UsrScrip directory.

**NOTE:** The default for the Rapid SQL directory is C:\Program Files\Embarcadero. If you changed the default, the sample script will still be located in the RSQL600\UsrScrip directory.

If you create the package included with the Rapid SQL installation, you can delete it and its objects from your system when you finish working with them. The objects to delete are as follows:

- The package COUNT\_TIME\_INTERVAL
- The package function WEEKEND\_DAYS\_()
- The package function WORKING\_DAYS\_()
- The package function YEARS\_ELAPSED\_BETWEEN\_()
- The procedure YEARS\_ELAPSED



- The procedure YEARS\_ELAPSED\_Y2K

## Overview

The Getting Started section guides you through:

- Opening the sample debug script.
- Executing the sample debug script.
- Changing the *Explorer* Tab display.
- Confirming the creation of the package, including its functions and procedures.

## Getting Started

- 1 Start Rapid SQL.

- 2 On the *File* menu, click *Open*.

Rapid SQL opens the Open File(s) dialog box.

- 3 In the *Open File(s)* dialog box, go to RSQL600\UsrScpt\DEBUGGER\_DEMO.sql, and then click *Open*.

**NOTE:** During the installation Rapid SQL places DEBUGGER\_DEMO.sql in the following directory: C:\Program Files\Embarcadero\RSQL600\UsrScpt.

Rapid SQL opens the What type of file dialog box.

- 4 On the What type of file dialog box, click The file includes the DDL to create a database object, and then click OK.

Rapid SQL opens the target script in an SQL Editor.

- 5 On the *SQL Editor* toolbar, click *Execute* to execute the script and create the package.

Rapid SQL executes the target script and opens the SQL Editor Results Tab, displaying the results of the script execution. If you were not able to create the package, check the error messages to determine the problem.

- 6 On the *Explorer* Tab list, click *Organize by Owner*.

Rapid SQL displays a list of owners in the Database Explorer.

- 7 On the *Explorer* Tab, double-click your owner name.

Rapid SQL displays a list of your schema objects.

- 8 Under your owner node, double-click the *Packages* node.

Rapid SQL displays *COUNT\_TIME\_INTERVAL*, confirming the package's creation. You are now ready to begin debugging [Sample Script 1](#) and [Sample Script 2](#).

## DEBUGGING SAMPLE SCRIPT 1

Sample Script 1 demonstrates Embarcadero SQL Debugger for Oracle's basic features and functionality with the function `WORKING_DAYS()`, which counts the number of business days between two dates.

### Overview

Debugging Sample Script 1 is divided into five parts:

- [Starting the Debug Session](#)
- [Entering Input Parameters](#)
- [Inserting Breakpoints](#)
- [Stepping Into](#)
- [Viewing Debug Session Results](#)

#### SAMPLE SCRIPT 1 - STARTING THE DEBUG SESSION

After you [open and execute](#) `DEBUGGER_DEMO.sql`, you can begin debugging Sample Script 1. To begin debugging the function `WORKING_DAYS()`, you must start a debug session.

### Starting the Debug Session

To start the debug session, do the following:

- 1 On the *Explorer* Tab, under the *Packages* node, double-click the *COUNT\_TIME\_INTERVAL* node.

Rapid SQL opens the *COUNT\_TIME\_INTERVAL* node and displays the following items:

- 2 Under the *COUNT\_TIME\_INTERVAL* node, double-click *Functions*.

Rapid SQL opens the *Functions* node and displays the following items:

- 3 Under the *Functions* node, right-click *WORKING\_DAYS()*, and then click *Debug* to start the debug session.

Rapid SQL opens the Function Execution dialog box with the current date in the boxes. You are now ready to begin working with [input parameters](#).

#### SAMPLE SCRIPT 1 - ENTERING INPUT PARAMETERS

After you [start a debugging session](#), you can enter input parameters. You cannot debug a script that requires input parameters until you input those parameters in the Function Execution dialog box.

### Input Parameters

To input parameters, do the following:

- 4 At the end of the *P\_START\_DATE DATE* row, click the drop-down arrow.  
Rapid SQL opens a calendar.
- 5 On the calendar, click left arrow to set the month to *November 1999*.
- 6 Click 1.  
Rapid SQL displays 11/01/1999 in the Value column of P\_START\_DATE.
- 7 Click the *P\_END\_DATE DATE* box, and then click the drop-down arrow.  
Rapid SQL opens a new calendar.
- 8 On the calendar, click left arrow to set the month to *November 1999*.
- 9 Click 8.  
Rapid SQL displays 11/08/1999 in the Value column of P\_END\_DATE.
- 10 Click Continue.  
Rapid SQL closes the *Function Execution* dialog box, and then opens the following five Embarcadero SQL Debugger for Oracle interface windows:
  - [SQL Editor](#), which displays the SQL code for the function.
  - [Watch window](#).
  - [Variables window](#).
  - [Call Stack window](#).
  - [Dependency Tree window](#), which displays the dependent objects.
 You are now ready to begin [inserting breakpoints](#).

#### SAMPLE SCRIPT 1- INSERTING BREAKPOINTS

After you [input parameters](#) in the Input Parameters dialog box, you can begin inserting breakpoints. In this example, the breakpoints must be inserted in the extracted dependent object code. After you extract this code, you must locate the target breakpoint lines by searching for the text DBMS\_OUTPUT.

### Breakpoints

To insert breakpoints, do the following:

- 1 In the *Dependency Tree window*, double-click the *COUNT\_TIME\_INTERVAL* package body.  
Rapid SQL displays the SQL code for the package body in the SQL Editor window.
- 2 On the *Edit* toolbar, click Find.  
Rapid SQL opens the *Find* dialog box.
- 3 On the Find dialog box, in the *Find What* box, type *DBMS\_OUTPUT*.

- 4 Click *Find Next*.

In the SQL Editor, Rapid SQL highlights the first occurrence of `DBMS_OUTPUT`, on line 22.

- 5 On the *SQL Editor* toolbar, click Breakpoint.

Rapid SQL inserts a breakpoint next to the target line number.

- 6 On the *Find dialog box*, click Find Next.

Rapid SQL highlights the next occurrence of `DBMS_OUTPUT`.

- 7 Click *Find Next* a third time.

Rapid SQL highlights the next occurrence of `DBMS_OUTPUT`, on line 35.

- 8 On the Find dialog box, click Cancel.

Rapid SQL closes the Find dialog box.

- 9 On the *SQL Editor* toolbar, click Breakpoint to insert a second breakpoint.

You should now have breakpoints set at lines 22 and 35. You are now ready to begin [stepping into](#) the code.

#### SAMPLE SCRIPT 1- STEPPING INTO

After you [insert breakpoints](#), you can step into the function code.

### Step Into

To Step Into the code, do the following:

- 1 On the *SQL Editor* toolbar, click Go.

Embarcadero SQL Debugger for Oracle begins debugging and runs to the first breakpoint, placing the yellow arrow on line 22.

- 2 On the *SQL Editor* toolbar, click Step Into .

Embarcadero SQL Debugger for Oracle moves the yellow arrow to the next line of the code.

- 3 Click Step Into again to enter the LOOP block.

Embarcadero SQL Debugger for Oracle displays the value of the variables in the Variables window.

- 4 Click Step Into again to start moving through the LOOP block.

In the Variables window, Embarcadero SQL Debugger for Oracle updates the value of variable `v_currdate` from 01-NOV-1999 to 02-NOV-1999.

- 5 Click Step Into two more times.

In the Variables window, Embarcadero SQL Debugger for Oracle updates the value of `v_theday` from NULL to Tuesday.

**NOTE:** If you continued stepping through the LOOP block, the Embarcadero SQL Debugger for Oracle would continue to update `v_currdate` and `v_theday` until `v_currdate` is greater than `p_end_date`.

- 6 On the *SQL Editor* toolbar, click Go.

Embarcadero SQL Debugger for Oracle runs to the next breakpoint.

- 7 On the *SQL Editor* toolbar, click Go once more.

Rapid SQL PL/SQL concludes the debug session and displays the [Debug Session Results box](#).

#### SAMPLE SCRIPT 1 - VIEWING DEBUG SESSION RESULTS

After [Stepping Into](#) and running to the end of the code, Embarcadero SQL Debugger for Oracle displays a Debug Session Results box containing the following information:

- Variable output
- DBMS\_OUTPUT

**NOTE:** In this example, the Embarcadero SQL Debugger for Oracle displays a Debug Session Results box because the sample program includes DBMS\_OUTPUT.

#### Debug Session Results

- 1 Click OK.

Rapid SQL closes the *Debug Session Results* box and terminates your debug session.

#### DEBUGGING SAMPLE SCRIPT 2

Sample Script 2 demonstrates Embarcadero SQL Debugger for Oracle's functionality when used on a function containing a bug which prevents it from executing successfully. The buggy function, `WEEKEND_DAYS()`, requires input parameters and counts the number of weekend days between two dates. In this section you must use Embarcadero SQL Debugger for Oracle to identify the bug, and then correct the script so that it can execute successfully.

#### Overview

Debugging Sample Script 2 is divided into six parts:

- [Executing the Function](#)
- [Starting the Debug Session](#)

- [Entering Input Parameters](#)
- [Inserting Breakpoints](#)
- [Stepping Into](#)
- [Correcting the Function](#)

#### SAMPLE SCRIPT 2 - EXECUTING THE FUNCTION

After you [open and execute](#) `DEBUGGER_DEMO.sql`, you can begin debugging Sample Script 2. To begin debugging the function `WEEKEND_DAYS ( )`, you must first execute the function to discover the type of error it returns when it fails to execute.

#### Executing the Function

To execute the function, do the following:

- 1 On the *Explorer* Tab, under the *Packages* node, double-click the *COUNT\_TIME\_INTERVAL* node.

Rapid SQL opens the *COUNT\_TIME\_INTERVAL* node and displays the following items:

- 2 Double-click the *Functions* node.

Rapid SQL opens the *Functions* node and displays the following items:

- 3 Click *WEEKEND\_DAYS ( )*, then right-click it and click *Execute*.

Rapid SQL opens the Function Execution dialog box.

- 4 In the Value column of the *P\_START\_DATE* row, type *11/01/1999*.

- 5 In the Value column of the *P\_END\_DATE* row, type *11/30/1999*.

- 6 Click *Execute*.

Rapid SQL attempts to execute the function but returns an error indicating that the character string buffer is too small.

You are now ready to [start the debug session](#).

#### SAMPLE SCRIPT 2 - STARTING THE DEBUG SESSION

After you unsuccessfully [execute the function](#) `WEEKEND_DAYS( )` and Rapid SQL displays the nature of its execution error, you can start a debugging session to determine the actual cause of the error.

## Starting the Debugging Session

To start the debugging session, do the following:

- 1 On the *Explorer* Tab, under the *COUNT\_TIME\_INTERVAL* node, under the *Functions* node, right-click *WEEKEND\_DAYS ()*, and then click *Debug* to start the debug session.

Rapid SQL opens the Function Execution dialog box.

You are now ready to begin [entering input parameters](#).

### SAMPLE SCRIPT 2 - ENTERING INPUT PARAMETERS

After you [start the debug session](#), you can enter input parameters in the Function Execution dialog box.

## Entering Input Parameters

To enter the input parameters, do the following:

- 1 At the end of the *P\_START\_DATE* row, click the drop-down arrow.

Rapid SQL opens a calendar.

- 2 On the calendar, click left arrow to set the month to *November 1999*.

- 3 Click *1*.

Rapid SQL displays 11/01/1999 in the Value column of the *P\_START\_DATE* row.

- 4 At the end of the *P\_END\_DATE* row, click the drop-down arrow.

Rapid SQL opens a new calendar.

- 5 On the calendar, click left arrow to set the month to *November 1999*.

- 6 Click *30*.

Rapid SQL displays 11/08/1999 in the Value column of the *P\_END\_DATE* row.

- 7 Click *Continue*.

Rapid SQL closes the *Function Execution* dialog box, and then opens the following five Embarcadero SQL Debugger for Oracle interface windows:

- [SQL Editor](#), which displays the SQL code for the function.
- [Watch window](#).
- [Variables window](#).
- [Call Stack window](#).
- [Dependency Tree window](#), which displays the dependent objects.

You are now ready to begin [inserting breakpoints](#).

## SAMPLE SCRIPT 2- INSERTING BREAKPOINTS

After you [enter input parameters](#), you can begin inserting breakpoints. In this example, the breakpoints must be inserted in the extracted dependent object code. After you extract this code, you must locate the target breakpoint lines by searching for a particular line of code.

**Breakpoints**

To insert breakpoints, do the following:

- 1 In the *Dependency Tree* window, double-click the *COUNT\_TIME\_INTERVAL* package body.  
Rapid SQL displays the SQL code for the package body in the SQL Editor window.
- 2 On the *Edit* toolbar, click Find.  
Rapid SQL opens the *Find* dialog box.
- 3 On the Find dialog box, in the *Find What* box, type *Function weekend\_days*, and then click *Find Next*.  
Embarcadero SQL Debugger for Oracle highlights the first occurrence of Function *weekend\_days*.
- 4 On the Find dialog box, click Cancel.  
Rapid SQL closes the Find dialog box.
- 5 Click line 60, the first line of executable code:
- 6 On the *SQL Editor* toolbar, click Breakpoint.  
Rapid SQL inserts a breakpoint next to the line number.
- 7 Click Go to start debugging and run to the breakpoint.  
Embarcadero SQL Debugger for Oracle places the yellow arrow on line 60 and populates the Variables window with the first set of variables in the function code.  
Embarcadero SQL Debugger for Oracle also populates the Call Stack window with everything called before the breakpoint.  
You are now ready to begin [stepping into](#) the function.

## SAMPLE SCRIPT 2- STEPPING INTO

After you [set and run to the breakpoint](#), you can step into the function to locate the cause of the error. To locate the cause of the error, you must monitor the Variables window. As you step through the code, the Variables window updates with the value of the variables.



## Step Into

To step into the function, do the following:

- 1 On the *SQL Editor* toolbar, click Step Into.

Rapid SQL moves the yellow arrow to the next line of the code, line 64.

- 2 On the *SQL Editor* toolbar, click Step Into.

Embarcadero SQL Debugger for Oracle's Variables window updates the value of `v_currdate` to 02-NOV-1999.

- 3 On the *SQL Editor* toolbar, click Step Into.

Rapid SQL moves the yellow arrow to the next line of the code, line 66.

- 4 On the *SQL Editor* toolbar, click Step Into.

Rapid SQL moves the yellow arrow to the next line of the code, line 67, and, in the Variables window, updates the value of `v_theday` to Tuesday.

- 5 On the *SQL Editor* toolbar, click Step Into.

Rapid SQL moves the yellow arrow back to line 64 to repeat the loop.

- 6 On the *SQL Editor* toolbar, click Step Into.

Embarcadero SQL Debugger for Oracle's Variables window updates the value of `v_currdate` to 03-NOV-1999.

- 7 On the *SQL Editor* toolbar, click Step Into.

Rapid SQL moves the yellow arrow to the next line of the code, line 66.

- 8 On the *SQL Editor* toolbar, click Step Into.

The Embarcadero SQL Debugger for Oracle locates the error. The application terminates the debug session, returns an error indicating that the numeric or value character string buffer is too small, extracts the `COUNT_TIME_INTERVAL` package code into an SQL Editor, and returns an error indicating the line on which the code failed. You are now ready to [correct the script](#).

## SAMPLE SCRIPT 2 - CORRECTING THE SCRIPT

After you [step through the SQL code and locate the error](#), you can correct the bug in Sample Script 2. When Embarcadero SQL Debugger for Oracle locates an error, it extracts the target package body into an SQL Editor. To correct this script, you must do the following:

- Scroll to the incorrect line in the script
- Analyze the code
- Correct the error
- Execute the corrected SQL script

- Execute the `WEEKEND_DAYS ( )` function

The code in Sample Script 2 fails on line 66, returning an error when the variable `v_theday` increments from the value Tuesday to the value Wednesday. The cause of this error is found in the declarations section of the function script, where the width of the `VARCHAR2` variable `v_theday` is set to 8. Because "Wednesday" includes nine characters, the value of the variable `v_theday` fails when it attempts to place a nine-character value in an eight-character variable. To correct this error, you must increase the width of the variable `v_theday` to accommodate nine characters.

### Correcting the Script

To correct the script, do the following:

- 1 On the *Explorer* Tab, under the *Packages* node, under the *COUNT\_TIME\_INTERVAL* node, right-click *Package Body*, and then click *Extract*.

Rapid SQL extracts the package body into an SQL Editor.

- 2 In the *SQL Editor*, scroll to line 57, the line defining the variable `v_theday`.
- 3 On line 57, change the value of the width from 8 to 9.
- 4 On the *SQL Editor* toolbar, click *Execute* to execute the script.

Rapid SQL successfully executes the script.

- 5 On the *Explorer* Tab, under the *COUNT\_TIME\_INTERVAL* package node, under the *Functions* node, click *WEEKEND\_DAYS ( )*.
- 6 Right-click *WEEKEND\_DAYS ( )*, and then click *Execute*.

Rapid SQL opens the Function Execution dialog box.

- 7 In the Value column of the *P\_START\_DATE* row, type *11/01/1999*.
- 8 In the Value column of the *P\_END\_DATE* row, type *11/30/1999*.
- 9 Click *Execute*.

Rapid SQL successfully executes the corrected function.

## EMBARCADERO SQL DEBUGGER FOR SYBASE ASE

Embarcadero SQL Debugger for Sybase is a programming tool that lets you debug Sybase objects in the following Sybase versions:

- 12
- 12.0 (special version)
- 12.5

## Objects

Using Embarcadero SQL Debugger for Sybase, you can debug the following objects:

- Procedures
- Triggers

You can only debug triggers by debugging the procedures that call them.

**NOTE:** Availability of this feature depends on your Rapid SQL licensing. For more information, see [Licensing](#).

**TIP:** The [Code Analyst](#) is a tool to identify time-consuming lines of code. Code Analyst lets you perform detailed response time analysis on the execution of Procedures and Functions.

The table below describes the sections of this chapter:

Section	Description
<a href="#">Features</a>	This section describes how Embarcadero SQL Debugger for Sybase helps you identify problems within your code.
<a href="#">Interface</a>	This section describes Embarcadero SQL Debugger for Sybase's graphical interface, which includes an editor window and four debug view windows.
<a href="#">Functionality</a>	This section describes the way in which Embarcadero SQL Debugger for Sybase functions.
<a href="#">Using Embarcadero SQL Debugger for Sybase</a>	This section describes how to run a debug session.

## EMBARCADERO SQL DEBUGGER FOR SYBASE FEATURES

Embarcadero SQL Debugger for Sybase is designed to help identify problems within your code. Embarcadero SQL Debugger for Sybase lets you:

- Interactively step through the flow of script execution.
- Examine the value of variables.
- Solve logical problems with your script design.

Embarcadero SQL Debugger for Sybase offers fundamental debugging features and several options to help fine tune debugging, as listed in the table below:

Debugging Feature	Description
<a href="#">Step Into</a>	Lets you execute each instruction step-by-step and step inside a stored object.

Debugging Feature	Description
<a href="#">Step Out</a>	Lets you stop stepping through the current object and execute the remainder of the script. This option is only active when the pointer indicates a child-dependent instruction.
<a href="#">Step Over</a>	Lets you execute the current instruction without stepping into any child dependents.
<a href="#">Breakpoints</a>	A position in a program where the debugger stops execution.

To set specific Debugger values on Rapid SQL's Options Editor, see [Debugger Options](#).

## EMBARCADERO SQL DEBUGGER FOR SYBASE OPTIONS

You can specify Debugger options from the Debug Tab of Rapid SQL's Options editor. The Debug Tab of the Options Editor lets you set the duration of your debug initialization and debug session, enable DBMS output, and refresh dependencies.

### Setting Debugger Options

To set debugger options, do the following:

- 1 On the *File* menu, click *Options*.  
OR  
On the *Main* toolbar, click Options.  
Rapid SQL opens the Options Editor.
- 2 On the Debug Tab, specify debugger options. The table below describes the options available:

Option	Description	Default
Initialization Timeout (seconds)	Specifies the number of seconds Rapid SQL tries to initialize the debugger. If it cannot initialize the debugger in the specified time, a message displays in the Debug Output window.	60
Debug Session Timeout (seconds)	Specifies, in seconds, the length of your debug session.	7200
Enable DBMS Output	Toggles the print output. Enable this option if you use <code>dbms_output.put_line</code> calls in your procedures and you want these lines displayed.	Selected
Refresh Dependencies for each run	Refreshes dependencies for each run. This potentially time-consuming process is useful if the target procedure has rapidly varying dependencies that can require updating during the debugging process.	Cleared

- 3 Click Close.

Rapid SQL closes the Options Editor.

## EMBARCADERO SQL DEBUGGER FOR SYBASE INTERFACE

Embarcadero SQL Debugger for Sybase has a graphical interface that includes an editor window and four debug view windows. When you open a debug session, Rapid SQL extracts the code for the object into a DDL Editor and opens four debug view windows at the bottom of the screen.

**TIP:** All Embarcadero debuggers display Performance Metrics that let you measure the execution time of each statement in the debug session.

The four debug view windows are optional, dockable windows designed to help debug your script. Embarcadero SQL Debugger for Sybase's five windows are:

- [DDL Editor window](#)
- [Watch window](#)
- [Variables window](#)
- [Call Stack window](#)
- [Dependency Tree window](#)

### Working with Debugger Windows

Rapid SQL lets you resize, move, dock and float the following Debugger windows:

- Watch
- Variables
- Call Stack
- Dependency Tree

To work with the above windows, do the following:

- 1 To resize the target window, click its frame and drag it.  
Rapid SQL resizes the window.
- 2 To move and dock the target window, click its grab bar and drag it.  
Rapid SQL moves the window to its new location and docks it with surrounding windows.
- 3 To float the target window, press Shift, then click its grab bar and drag it.  
Rapid SQL frames the window in its own floating frame and moves the window to its new location.

## DDL EDITOR WINDOW

Embarcadero SQL Debugger for Sybase provides a DDL Editor window that displays your code in read-only format. When you start debugging, Embarcadero SQL Debugger for Sybase extracts your code into a DDL Editor window, making it editable. The DDL Editor uses the default Rapid SQL syntax coloring.

## WATCH WINDOW

Embarcadero SQL Debugger for Sybase provides a Watch window that displays the watch variables for the database object you are debugging and lets you specify variables you want to evaluate or modify while debugging your program. For example, to check what happens when a variable (x) has a value of 100, you can double-click that variable in the DDL Editor and drag it into the Watch window. In the Watch window, change the value to 100. When you execute the script, the debugger uses the value x=100. This window is only visible when Embarcadero SQL Debugger for Sybase is active.

**NOTE:** You can type a fully qualified record variable into the Watch window.

**NOTE:** When you exit a debug session and reenter it, Embarcadero SQL Debugger for Sybase retains any watch variables or breakpoints you have set.

### Opening and Closing the Watch Window

To open and close the Watch Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views sub-menu*, select or clear *Watch*.

OR

Press ALT+3.

### Setting a Watch Variable

To set a Watch Variable, do the following:

- 1 In the *DDL Editor*, double-click the target variable and drag it to the *Watch window*.
- 2 In the *Watch window*, in the *Value column*, change the value of the variable.
- 3 On the *DDL Editor* toolbar, click *Go* to execute the script using the new value of the variable.  
Rapid SQL executes the script using the new value of the variable.

## Removing a Watch Variable

To remove a Watch Variable, do the following:

- 1 In the *Watch window*, delete the variable.

## VARIABLES WINDOW

Embarcadero SQL Debugger for Sybase provides a Variables window that displays the local variables and their current values during script execution. You cannot edit variables in the Variables window. This window is only visible when the Debugger is active. If the DDL Editor displays an external database object, and that object is a dependent of the object you are debugging, then the Variables window automatically refreshes and displays the variables for that external object.

Embarcadero SQL Debugger for Sybase also lets you monitor the value of your variables while debugging.

## Opening and Closing the Variables Window

To open and close the Variables Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views sub-menu*, select or clear *Variable*.

OR

Press ALT+4.

## Monitoring Variables

To monitor the values of your variables while debugging, do the following:

- 1 In the DDL Editor, hold the pointer over the target variable.

Rapid SQL opens a ScreenTip displaying the current value of that variable.

## CALL STACK WINDOW

Embarcadero SQL Debugger for Sybase provides a Call Stack window that displays the stack of currently active calls. The Call Stack window is only visible when the Debugger is active.

## Opening and Closing the Call Stack Window

To open and close the Call Stack Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views sub-menu*, select or clear *Call Stack*.

OR

Press ALT+5.

## Using the Call Stack Window

To display a line of code that references the call in the *DDL Editor*, do the following:

- 1 In the *Call Stack window*, double-click the target line.

In the DDL Editor, Rapid SQL displays a green arrow on the line of the the referenced call.

## DEPENDENCY TREE WINDOW

Embarcadero SQL Debugger for Sybase provides a Dependency Tree window that displays any external database objects accessed by the script. Rapid SQL displays these database objects in a hierarchical tree, with the child objects as database objects accessed by the parent objects. You can use this window to display the code for a dependent database object in the *DDL Editor* window. This window is only visible when the Debugger is active.

## Opening and Closing the Dependency Tree Window

To open and close the Dependency Tree Window, do the following:

- 1 On the *Debug Menu*, on the *Debug Views sub-menu*, select or clear *Dependencies*.

OR

Press ALT+6.

## Displaying Dependencies

To display the code for a dependent database object in the *DDL Editor* window, do the following:

- 1 In the *Dependency Tree* window, double-click the target object.

Rapid SQL displays the SQL of the target object in the DDL Editor window.

## EMBARCADERO SQL DEBUGGER FOR SYBASE FUNCTIONALITY

Embarcadero SQL Debugger for Sybase offers you the following functionality:

- [Input parameters](#)
- [Step Into](#)
- [Step Out](#)
- [Step Over](#)
- [Run to Cursor](#)
- [Insert or Remove Breakpoint](#)
- [Toggle Breakpoint](#)



- [Go](#)
- [Stop](#)
- [Restart](#)
- [Break](#)
- [Close](#)

To make use of the above functionality, you must first [open a debugging session](#).

## INPUT PARAMETERS

Input parameters are set when you first create an object. If the object you want to debug requires input parameters, Rapid SQL opens a Procedure Execution dialog box and prompts you for the input parameters when you [open a debugging session](#).

This dialog box also lets you:

- Save input parameters as \*.prm files to preserve specific input parameter configurations.
- Open \*.prm files to save the effort of reentering specific input parameters.
- Reset parameters to their default setting.

The table below describes the options and functionality on Procedure Execution dialog box:

The following table describes the options available in this dialog box:

Dialog box component	Description
Owner drop-down list	Displays the current procedure's owner
Procedure drop-down list	Displays the name of the current procedure.
Parameter window	Specify the required input parameters in this window. If input parameters are not required for the execution of the target procedure, a message appears in this window, stating that the procedure "has no input parameters. Press execute to run it."
Open button	Click to open an Open dialog box, from which you can open an existing *.prm file. The saved parameters immediately populate the dialog box upon opening.
Save button	Click to save the values of your input parameters as a *.prm file. You can reopen a saved *.prm file from this dialog box at any time.
Reset button	Click to reset the parameters in the Parameter window to their default values.
Execute button	Click to execute the procedure once you have entered values for all required parameters in the Parameter window.

Option	Description
Owner	Displays the current procedure's owner.
Procedure	Displays the name of the current procedure.
Parameter	Specify the required input parameters in this window. If input parameters are not required for the execution of the target procedure, a message displays in this window, stating that the procedure "has no input parameters. Press execute to run it."
Open	Click to open an existing *.prm file. The saved parameters immediately populate the dialog box upon opening.
Save	Click to save the values of your input parameters as a *.prm file. You can reopen a saved *.prm file from this dialog box at any time.
Reset	Click to reset the parameters in the Parameter window to their default values.
Continue	Click to execute the procedure once you have entered values for all required parameters in the Parameter window.

**NOTE:** You cannot debug a script that requires input parameters until you provide input parameters.

## STEP INTO

After you [open a debugging session](#), *Step Into* lets you execute the current instruction. If the current instruction makes a call to a stored Sybase object, Embarcadero SQL Debugger for Sybase steps inside the nested child object.

### Step Into

To use the Step Into facility, do the following:

- 1 On the *Debug* menu, click *Step Into*.

OR

On the *DDL Editor* toolbar, click Step Into.

OR

In the *DDL Editor* window, right-click and then click *Step Into*.

OR

Press F11.

Embarcadero SQL Debugger for Sybase moves the arrow to execute the current instruction.

## STEP OUT

After you [open a debugging session](#), Step Out lets you execute the remainder of the dependent child object and resumes line-by-line, step-debugging in the parent object.

**NOTE:** Step Out is only active when the pointer indicates a child-dependent instruction.

### Step Out

To use the Step Out facility, do the following:

- 1 On the *Debug* menu, click *Step Out*.

OR

On the *DDL Editor* toolbar, click Step Out.

OR

In the *DDL Editor* window, right-click and then click *Step Out*.

OR

Press SHIFT+F11.

Embarcadero SQL Debugger for Sybase stops stepping through the current object and executes the remainder of the script.

## STEP OVER

After you [open a debugging session](#), Step Over lets you execute the current instruction without stepping into a nested child object if the instruction makes a call to a dependent object.

### Step Over

To use the Step Over facility, do the following:

- 1 On the *Debug* menu, click *Step Over*.

OR

On the *DDL Editor* toolbar, click Step Over.

OR

In the *DDL Editor* window, right-click and then click *Step Over*.

OR

Press F10.

Embarcadero SQL Debugger for Sybase executes the current instruction.

## RUN TO CURSOR

After you [open a debugging session](#), Run to Cursor lets you execute all instructions between the yellow arrow and your cursor.

### Run to Cursor

To use the Run to Cursor facility, do the following:

- 1 Scroll down from the yellow arrow to the target line.

- 2 Click the target line.

Embarcadero SQL Debugger for Sybase places the cursor on the target line.

- 3 On the *Debug* menu, click *Run to Cursor*.

OR

On the *DDL Editor* toolbar, click Run to Cursor.

OR

In the *DDL Editor* window, right-click and then click *Run to Cursor*.

OR

Press CTRL+F10.

Embarcadero SQL Debugger for Sybase executes all instructions between the yellow arrow and your cursor.

## INSERT OR REMOVE A BREAKPOINT

A breakpoint is a position in a program where a debugger stops execution. When you start debugging, Embarcadero SQL Debugger for Sybase opens the script in a DDL Editor window. A yellow pointer indicates which line the Debugger executes next. Embarcadero SQL Debugger for Sybase executes all lines of code between the pointer and the first breakpoint. If no breakpoints are present, Embarcadero SQL Debugger for Sybase debugs the entire script.

While debugging, you can set one or more breakpoints in the currently executing object or in any object in the program call stack. Breakpoints can be [toggled](#) (temporarily disabled or enabled) without having to remove or reinsert them. Embarcadero SQL Debugger for Sybase displays each enabled breakpoint as a red dot in the left margin of the DDL Editor Window, and each disabled breakpoint as a red circle.

Rapid SQL stores all breakpoints you set so that when you debug the same script on separate occasions, you can reuse the same breakpoints each time.

After you [open a debugging session](#), you can insert a breakpoint on the line where your cursor is located, and you can remove a breakpoint on the line where your cursor is located.

**NOTE:** Script execution stops at the first breakpoint.

## Inserting or Removing a Breakpoint

To insert or remove a breakpoint, do the following:

- 1 In the DDL Editor window, click the target line of SQL.
- 2 On the *Debug* menu, click *Breakpoint*.

OR

On the *DDL Editor* toolbar, click Breakpoint.

OR

In the *DDL Editor* window, right-click and then click *Breakpoint*.

OR

Press F9.

Embarcadero SQL Debugger for Sybase inserts a new breakpoint or removes an existing breakpoint on the target line of code.

## TOGGLE BREAKPOINT

After you [open a debugging session](#) and [insert a breakpoint](#), Toggle Breakpoint lets you enable or disable that breakpoint. Embarcadero SQL Debugger for Sybase displays each enabled breakpoint as a red dot in the left margin of the DDL Editor Window, and each disabled breakpoint as a red circle. You can toggle any breakpoint in the DDL Editor window. When you exit a debugging session and reenter it, the Embarcadero SQL Debugger for Sybase retains any breakpoints you set.

## Toggling a Breakpoint

To toggle a breakpoint, do the following:

- 1 In the DDL Editor window, click the line of the target breakpoint.
- 2 On the *Debug* menu, click *Enable/Disable Breakpoint*.

OR

On the *DDL Editor* toolbar, click Enable/Disable Breakpoint.

OR

In the *DDL Editor* window, right-click and then click *Enable/Disable Breakpoint*.

OR

Press CTRL+F9.

Embarcadero SQL Debugger for Sybase toggles the breakpoint indicated by the pointer.

## Go

After you [open a debugging session](#), Go lets you execute all instructions, stopping only when it encounters a breakpoint or when the program is complete.

### Go

To use the Go facility, do the following:

- 1 On the *Debug* menu, click *Go*.

OR

On the *DDL Editor* toolbar, click *Go*.

OR

In the *DDL Editor* window, right-click and then click *Go*.

OR

Press F5.

Embarcadero SQL Debugger for Sybase executes all instructions.

## STOP

After you [open a debugging session](#), Stop lets you halt the script execution and terminate the session.

### Stop

To stop the debugger, do the following:

- 1 On the *Debug* menu, click *Stop Debugging*.

OR

On the *DDL Editor* toolbar, click *Stop Debugging*.

OR

In the *DDL Editor* window, right-click and then click *Stop Debugging*.

OR

Press SHIFT+F5.

Embarcadero SQL Debugger for Sybase stops the script execution and terminates the session.

## RESTART

After you [open a debugging session](#), Restart lets you terminate the current debug session and open a new one. When the new session opens, Rapid SQL prompts you for new input parameters.

### Restart

To restart the debugger, do the following:

- 1 On the *Debug* menu, click *Restart*.
- OR
- On the *DDL Editor* toolbar, click *Restart*.
- OR
- In the *DDL Editor* window, right-click and then click *Restart*.
- OR
- Press CTRL+SHIFT+F5.
- Embarcadero SQL Debugger for Sybase restarts the debug session.

## BREAK

After you [open a debugging session](#), Break lets you pause the debug session.

### Break

To pause the debugger, do the following:

- 1 On the *Debug* menu, click *Break*.
- OR
- On the *DDL Editor* toolbar, click *Break*.
- OR
- In the *DDL Editor* window, right-click and then click *Break*.
- Embarcadero SQL Debugger for Sybase suspends the debug session.

## CLOSE

After you [open a debugging session](#), Close lets you close the DDL Editor and Embarcadero SQL Debugger for Sybase.

**Close**

To close the DDL Editor and debugger, do the following:

- 1 On the *DDL Editor* toolbar, click Close.

OR

In the upper right corner of the window, click Close.

OR

In the *DDL Editor* window, right-click and then click *Close*.

Embarcadero SQL Debugger for Sybase closes the debug session.

## USING EMBARCADERO SQL DEBUGGER FOR SYBASE

This section offers a general overview of how to use Embarcadero SQL Debugger for Sybase's full range of debugging [functionality](#). After you [open a debugging session](#) for any Sybase procedure or trigger, you can [begin debugging](#).

### OPENING A DEBUGGING SESSION

When you open a debugging session, Rapid SQL opens the [five windows](#) of the Embarcadero SQL Debugger for Sybase interface. If the target script requires input parameters, Rapid SQL opens a Procedure Execution dialog box and prompts you for the necessary input parameters before displaying the target code in the DDL Editor window. When Rapid SQL displays the target script in the DDL Editor window, you can [begin debugging](#).

**NOTE:** Embarcadero SQL Debugger for Sybase only lets you debug the SQL script of triggers and procedures.

### Opening a Debugging Session

To debug a trigger or procedure, do the following:

- 1 On the *Explorer* Tab, click the node of the target procedure.

Rapid SQL opens the node and displays two items: Code and Privileges.

- 2 Under the target object node, double-click Code.

Rapid SQL opens an DDL Editor *displaying the code of the target object*.



- 3 On the Debug menu, click Start Debugging.

OR

On the DDL Editor toolbar, click Debug.

OR

In the *DDL Editor* window, right-click and then click *Debug*.

OR

Press CTRL+F5.

If the script requests input parameters, Rapid SQL opens a Procedure Execution dialog box. If the script does not require input parameters, Rapid SQL displays the script in the DDL Editor window for you to [begin debugging](#).

**NOTE:** You cannot use Embarcadero SQL Debugger for Sybase until it has fully initialized.

- 4 In the Procedure Execution dialog box, type the parameters, and then click OK.

Rapid SQL displays the script in the DDL Editor window for you to [begin debugging](#).

**NOTE:** If the script requires Sybase types (tables, records, or Booleans) as input parameters, Embarcadero SQL Debugger for Sybase generates an anonymous block.

## DEBUGGING AN SQL SCRIPT

After you [open a debugging session](#) and enter any required input parameters, you can begin working with your script in Embarcadero SQL Debugger for Sybase.

### Debugging an SQL Script

To debug a SQL script, do the following:

- 1 On the *Debug* menu, click an Embarcadero SQL Debugger for Sybase option ([Step Into](#), [Step Over](#), and so forth) or click Go.

OR

On the *DDL Editor* toolbar, click an Embarcadero SQL Debugger for Sybase option ([Step Into](#), [Step Over](#), and so on) or click Go.

**NOTE:** You can monitor the progress of your debug session in the Variables window.

- 2 On the *Debug* menu, click *Breakpoint*.

OR

On the *DDL Editor* toolbar, click *Breakpoint*.

**NOTE:** When you set a breakpoint, the Call Stack window shows what was called before the breakpoint.

**NOTE:** You can use the [Run to Cursor](#) option to test the lines of code between a breakpoint and your cursor (indicated by the yellow arrow in the DDL Editor).

- 3 To check your variables, do the following:
  - 1) In the **DDL Editor**, click a variable in your script and drag it to the **Watch** window.
  - 2) In the **Watch** window, change the value of the watch variable and then click **Go** to run your script and see the results of the new value.
- 4 To check record in stored objects, do the following:
  - 1) Drag the record to the **Watch** window.
  - 2) In the **Watch** window, change the value of the record, then click **Go** to run your script and see the results of the new value.
- 5 To check the dependencies, do the following:
  - 1) In the **Dependency Tree** window double-click the target dependent object to extract the code into a new **DDL Editor**.
  - 2) **Step through** the script while monitoring the [Dependency Tree window](#).
- 6 When you finish debugging the script, click Close.

Rapid SQL closes an Embarcadero SQL Debugger for Sybase's DDL Editor.

**NOTE:** When you exit a debug session and reenter it, Embarcadero SQL Debugger for Sybase retains any watch variables or breakpoints you have set.

## RAPID SQL PL/SQL PROFILER

The Rapid SQL PL/SQL Profiler module lets you capture metrics of various PL/SQL programmable objects as they are executed in the database. Developers can use data collected in profile sessions to improve performance of PL/SQL code execution. Rapid SQL PL/SQL Profiler collects and stores data in database tables and helps identify and isolate performance problems, and provide code coverage information. The Rapid SQL PL/SQL Profiler lets you:

- Graphically browse PL/SQL profiling data within the Explorer Tab
- View profiling data in the right pane of the application, which is populated as you navigate the Explorer Tab
- Start and stop PL/SQL profiling sessions with a single click
- Graphically analyze time spent in each programmable object (unit)

- Graphically analyze time spent in each source code line of a unit

The table below describes the sections of this chapter:

Section	Description
<a href="#">Setting Up the Profiler</a>	This section describes the process of setting up Rapid SQL PL/SQL Profiler.
<a href="#">Profiler Functionality</a>	This section describes the functionality of Rapid SQL PL/SQL Profiler.
<a href="#">Using the Profiler</a>	This section describes how to run a profile session.

**NOTE:** Availability of this feature depends on your Rapid SQL licensing. For more information, see [Licensing](#).

For more information, see:

- [Setting Up Rapid SQL PL/SQL Profiler](#)
- [Rapid SQL PL/SQL Profiler Explorer](#)
- [Rapid SQL PL/SQL Profiler Functionality](#)
- [Using Rapid SQL PL/SQL Profiler](#)

## SETTING UP RAPID SQL PL/SQL PROFILER

The profiling tables must be on the server before using the Rapid SQL PL/SQL Profiler. The first time you open the PL/SQL Profiler, Rapid SQL checks the server for the profiling tables. If the profiling tables are not on the server, Rapid SQL automatically installs profiling tables on the server.

## RAPID SQL PL/SQL PROFILER EXPLORER

The Rapid SQL PL/SQL Profiler displays profiling data in the right pane of the application, which is populated as you navigate the Explorer Tab.

The table below describes the nodes of the Rapid SQL PL/SQL Profiler Explorer and the corresponding information in the right pane of the application:

Node	Right pane information
PL/SQL Code Profiling	Contain all Comment, Run ID and Run Date\time data that is current stored in the Profiling tables.
Label\Comment level	Contains all Run ID and Run Date\time data for the specific Label\Comment.

Node	Right pane information
Run level	Contains all Unit, Unit Name, Unit Type, Run Date\time data for the specific Run ID.

## RAPID SQL PL/SQL PROFILER FUNCTIONALITY

Rapid SQL PL/SQL Profiler offers you the following functionality:

- [Start](#)
- [Flush](#)
- [Run Summary](#)
- [Run Detail](#)
- [Unit Summary](#)
- [Unit Detail](#)
- [Clear Profile Table](#)
- [Stop](#)

### START

Rapid SQL PL/SQL Profiler lets you begin a new profiling session or open a previous profiling session with the Start command.

### Starting a New Profile Session

To start a new Profiler session, do the following:

- 1 On the *Tools* menu, click *PL/SQL Profiler* and then click *Start*.

OR

On the *PL/SQL Profiler* toolbar, click *Execute*.

OR

On the *Explorer* Tab, right-click the *PL/SQL Code Profiling* node, and then click *Start*.

Rapid SQL opens the PL/SQL Profiler - Start dialog box.

- 2 In the *Profile Label* box, type the name of the new profile.

**NOTE:** Each user can own one or more Profiles.

- 3 Click *OK* to begin profiling.

## Starting an Existing Profile Session

To start an existing Profiler session, do the following:

- 1 On the Explorer Tab, expand the PL/SQL Code Profiling node.  
Rapid SQL displays the list of existing Profiles.
- 2 On the *Tools* menu, click PL/SQL Profiler, and then click Start.  
OR  
On the *PL/SQL Profiler* toolbar, click Execute.  
OR  
On the Explorer Tab, right-click the PL/SQL Profiler node, and then click Start.  
OR  
On the Explorer Tab, right-click the target *Profile*, and then click *Start*.  
Rapid SQL opens the PL/SQL Profiler - Start dialog box.
- 3 Click the *Profile Label* list and then click the existing profile.
- 4 Click *OK* to begin profiling.

## FLUSH

Rapid SQL PL/SQL Profiler lets you move the data from the dynamic tables into Analysis tables with the Flush command.

## Flushing a Profile

To flush a Profile, do the following:

- 1 On the *Tools* menu, click *PL/SQL Profiler* and then click *Flush*.  
OR  
On the *PL/SQL Profiler* toolbar click Flush.  
OR  
On the *Explorer Tab*, right-click the *PL/SQL Code Profiling* node, and then click *Flush*.  
Rapid SQL opens the PL/SQL Profiler - Flush dialog box.
- 2 In the *PL/SQL Profiler - Flush* dialog box:
  - Click *Flush* to delete the data in a running profile.
  - Click *Flush & Analyze* to open the [PL/SQL Profiler Run Detail](#) window.
  - Click *Cancel* to abort the flush and continue the profiling session.

**NOTE:** You can only Flush a running Profile.

## RUN SUMMARY

The Rapid SQL PL/SQL Profiler Run Summary window lets you view the following information for each of your profiles:

- Run ID
- Run Date
- Total Time

### Opening the Run Summary Window

To open the Run Summary Window, do the following:

- 1 On the Explorer Tab, expand the PL/SQL Code Profiling node.  
Rapid SQL displays the list of existing Profiles.
- 2 On the *Tools* menu, click PL/SQL Profiler, and then click Run Summary.

On the *PL/SQL Profiler* toolbar click Run Summary.

On the Explorer Tab, right-click the PL/SQL Profiler node, and then click Run Summary.

On the Explorer Tab, right-click the target *Profile*, and then click *Run Summary*.

3

Rapid SQL opens the PL/SQL Profiler - Run Summary window.

- 4 Click the *Profile Label* list, and then click the target profile.

## RUN DETAIL

The Rapid SQL PL/SQL Profiler Run Detail window lets you view the following information for each of your profiles:

- Run Number
- Run Date
- Run Time

The Rapid SQL PL/SQL Profiler lets you view the information for all runs or you can view profile information based on the unit type or unit owner.

The Rapid SQL PL/SQL Profiler Run Detail window lets you view results in milliseconds, seconds and minutes. The Run Detail window also contains graphical displays of the profiling data that you can go to the specific unit within the summary portion of the window.

## Opening the Run Detail Window

To open the Run Detail Window, do the following:

- 1 On the Explorer Tab, expand the PL/SQL Code Profiling node.  
Rapid SQL displays the list of existing Profiles.
- 2 On the *Tools* menu, click *PL/SQL Profiler* and then click *Run Detail*.  
OR  
On the *PL/SQL Profiler* toolbar click Run Detail.  
OR  
On the *Explorer* Tab, right-click the *PL/SQL Code Profiling* node and then click *Run Detail*.  
OR  
Right-click the target run and then click *Run Detail*.  
OR  
In a *PL/SQL Profiler* - [Run Summary](#) window right-click and then click *Detail*.  
OR  
In a *PL/SQL Profiler* - [Unit Summary](#) window right-click and then click *Detail*.  
Rapid SQL opens the PL/SQL Profiler - Run Detail window.
- 3 In the *PL/SQL Profiler* - *Run Detail* window:
  - Click the *Label* list box and then click the target profile.
  - Click the *Run* list and then click the target run.
  - Click the *Unit Type* list and then click the target unit type(s).
  - Click the *Unit Owner* list and then click the target unit owner(s) to populate the table.

## UNIT SUMMARY

The Rapid SQL PL/SQL Profiler Unit Summary window lets you view the following information for each of your profiles:

- Run ID
- Run Date
- Run Time
- Unit Time
- Percentage of Run Time

The Rapid SQL PL/SQL Profiler Unit Summary window lets you view results in milliseconds, seconds and minutes. The Unit Summary window also displays graphs of execution statistics for the top N runs and associated units. You can use the graphical displays to go to the specific run within summary portion of the window.

### Opening the Unit Summary Window

To open the Unit Summary Window, do the following:

- 1 On the *Tools* menu, click *PL/SQL Profiler*, and then click *Unit Summary*.  
OR  
On the *PL/SQL Profiler* toolbar click Unit Summary.  
OR  
On the *Explorer* Tab, right-click the *PL/SQL Code Profiling* node, and then click *Unit Summary*.  
Rapid SQL opens the PL/SQL Profiler - Unit Summary window.
- 2 In the *PL/SQL Profiler - Unit Summary* window:
  - Click the *Unit Owner* list and then click the target unit owner.
  - Click the *Unit Name* list and then click the target unit name to populate the table.

### CLEAR PROFILE TABLE

Rapid SQL PL/SQL Profiler lets you delete data from the user's profile tables with the command Clear Profile Table.

### Clearing a Profile Table

To clear a profile table, do the following:

- 1 On the *Tools* menu, click *PL/SQL Profiler*, and then click *Clear Profile Table*.  
OR  
On the *PL/SQL Profiler* toolbar, click *Clear Profile Table*.  
OR  
On the *Explorer* Tab, right-click the *PL/SQL Code Profiling* node, and then click *Clear Profile Table*.  
Rapid SQL clears the profile table.
- 2 In the dialog box, if you are sure that you want to clear out the profiler tables, click *Yes*.



## UNIT DETAIL

The Rapid SQL PL/SQL Profiler Unit Detail window lets you view the following information for each of your profiles:

- Average Time
- Source
- PL/SQL Script

The Rapid SQL PL/SQL Profiler Unit Detail window lets you view results in milliseconds, seconds and minutes. The Unit Detail window also provides two calculation options for viewing unit execution time as a percentage of total execution time (total run vs unit run). The Rapid SQL PL/SQL Profiler Unit Detail window also displays graphs of execution statistics for the top N run. You can use the graphical displays to go to the specific line within source code portion of the window. The graphical display portion of the window contains options for viewing advanced statistics.

The [Advanced View](#) of the Rapid SQL PL/SQL Profiler Unit Detail window, lets you view the following information for each of your profiles:

- Hit Lines
- Missed Lines
- Line Number
- Calls
- Total Time
- Percentage of the Total Time
- Average Time
- Minimum Time
- Maximum Time

### Opening the Unit Detail Window

To open the Unit Detail Window, do the following:

- 1 On the Explorer Tab, expand the PL/SQL Code Profiling node.

Rapid SQL displays the list of existing Profiles.

- 2 On the *Tools* menu, click *PL/SQL Profiler* and then click *Unit Detail*.

OR

On the *PL/SQL Profiler* toolbar click *Unit Detail*.

OR

On the *Explorer* Tab, right-click the *PL/SQL Code Profiling* node and then click *Unit Detail*.

OR

Right-click the target run and then click *Unit Detail*.

OR

In a *PL/SQL Profiler* - [Run Summary](#) window right-click and then click *Unit Detail*.

OR

In a *PL/SQL Profiler* - [Unit Summary](#) window right-click and then click *Unit Detail*.

Rapid SQL opens the *PL/SQL Profiler* - *Unit Detail* window.

- 3 In the *PL/SQL Profiler* - *Unit Detail* window, do any of the following:
  - Click the *Label* list and then click the target profile.
  - Click the *Run* list and then click the target run.
  - Click the *Unit* list and then click the target unit to populate the table.
  - Right-click and then click *Show Only Hit Lines* to populate the table with the *Average Time* and *Source* for hit lines.
  - Right-click and then click *Show Only Missed Lines* to populate the table with the *Average Time* and *Source* for missed lines.

### Opening the Unit Detail Window Advanced View

To open the Unit Detail Window Advanced View, do the following:

- 1 In the Unit Detail table, right-click and then click *Advanced View* to populate the table with advanced view information.

### STOP

Rapid SQL PL/SQL Profiler command Stop pauses the data gathering operation. Stop & Analyze populates the summary tables so you can view the Unit Detail and Run Summary windows

## Stopping a Profiling Session

To stop a Profiling session, do the following:

- 1 On the *Tools* menu, click *PL/SQL Profiler* and then click *Stop*.  
OR  
On the *PL/SQL Profiler* toolbar click *Stop*.  
OR  
On the *Explorer* Tab, right-click the *PL/SQL Code Profiling* node, and then click *Stop*.  
Rapid SQL opens the *PL/SQL Profiler - Stop* dialog box.
- 2 In the *PL/SQL Profiler - Stop* dialog box:
  - Click *Stop* to stop the profiling session.
  - Click *Stop & Analyze* to open the [PL/SQL Profiler Run Detail](#) window.
  - Click *Cancel* to continue the profiling session.

## USING RAPID SQL PL/SQL PROFILER

The steps below provide a high level overview of running a profiling session:

- Starting the Session
- Executing the Sample Script
- Stopping and Analyzing the Session

For more information, see [Sample Profiling Session](#).

### Using the Rapid SQL PL/SQL Profiler

**NOTE:** The first execution of a PL/SQL unit can take more time to execute because the code is loading into memory; subsequent runs take less time.

To use the Profiler, do the following:

- 1 On the *Tools* menu, click *PL/SQL Profiler* and then click *Start*.  
OR  
On the *PL/SQL Profiler* toolbar, click *Start*.  
Rapid SQL opens the *PL/SQL Profiler - Start* dialog box.
- 2 In the *Profile Label* box, enter the name of the new profile.

**NOTE:** Each user can own one or more Profiles.

- 3 Click *OK*.

Rapid SQL begins profiling.

- 4 In the *Explorer* Tab, execute one of the following PL/SQL database objects:

- Procedure
- Function
- Package Procedure
- Package Function

Profiler displays profiling data in the right pane of the application.

- 5 On the *Tools* menu, click *PL/SQL Profiler* and then click [Stop](#).

Rapid SQL opens the PL/SQL Profiler - Stop dialog box.

- 6 In the *PL/SQL Profiler - Stop* dialog box, do any of the following:

- Click *Stop* to stop the profiling session.
- Click *Stop & Analyze* to open the [PL/SQL Profiler Run Detail](#) window.
- Click *Cancel* to continue the profiling session.

- 7 If you clicked *Stop & Analyze* do the following:

- Click the *Label* list and then click the target profile.
- Click the *Run* list and then click the target run.
- Click the *Unit Type* list and then click the target unit type(s).
- Click the *Unit Owner* list and then click the target unit owner(s) to populate the table.

- 8 Use the *Tools* menu to open any of the following PL/SQL Profiler windows:

- [PL/SQL Profiler Run Summary](#)
- [PL/SQL Profiler Unit Summary](#)
- [PL/SQL Profiler Unit Detail](#)

## SAMPLE PROFILING SESSION

The Rapid SQL installation includes two scripts for the sample profiling session:

- PROFILER\_BUILD\_DEMO.SQL
- PROFILER\_DEMO.SQL

The PROFILER\_BUILD\_DEMO.SQL creates the objects that you will profile in the walk-through, and the PROFILER\_DEMO.SQL is what you will profile during the walk-through.

**NOTE:** To create the objects in the PROFILER\_BUILD\_DEMO.SQL script, you must have CREATE privileges.

The sample script demonstrates the following features of the Rapid SQL PL/SQL Profiler:

- Unit Detail
- Run Detail
- Show Only Hit Lines
- Advanced View

During the installation, Rapid SQL places the scripts in the C:\Program Files\Embarcadero\RSQL600\UsrScript directory.

**NOTE:** The default for the Rapid SQL directory is C:\Program Files\Embarcadero. If you changed the default, the sample scripts will still be located in the RSQL600\UsrScript directory.

## Overview

Sample Profiling Session is divided into six parts:

- Getting Started
- Starting the Session
- Executing the Sample Script
- Stopping the Session
- Rerunning & Reexecuting the Session
- Stopping & Analyzing

### SAMPLE PROFILING SESSION - GETTING STARTED

In this step of Sample Profiling Session, you create the objects that you will profile in the walk-through.

## Overview

The Getting Started section guides you through:

- Opening PROFILER\_BUILD\_DEMO.SQL
- Changing the Explorer tab Display
- Confirming the Creation of the Package

## Getting Started

- 1 Start Rapid SQL.

- 2 On the *File* menu, click *Open*.

Rapid SQL opens the Open Files dialog box.

**NOTE:** The default for the Rapid SQL directory is C:\Program Files\Embarcadero. If you changed the default, the sample scripts will still be located in the RSQL600\UsrScript directory.

- 3 In the *Open Files* dialog box, type the path to the UsrScript directory, press ENTER and then double-click *PROFILER\_BUILD\_DEMO.SQL* to open the script in an SQL Editor window.

Rapid SQL opens the PROFILER\_BUILD\_DEMO.SQL script in an SQL Editor window

- 4 On the *SQL Editor* window, click *Execute*.

Rapid SQL executes the script and creates the package.

- 5 On the *Explorer* Tab list, click *Organize by Owner*.

- 6 On the *Explorer* Tab, click the node of your owner name.

Rapid SQL displays your schema objects.

- 7 Double-click the *Packages* node to display *PF\_COUNT\_TIME\_INTERVAL* and confirm its creation.

**NOTE:** If you were not able to create the package, check the error messages to determine the problem.

## SAMPLE PROFILING SESSION - STARTING THE SESSION

In this step of Sample Profiling Session, you start the profiling session.

### Sample Profiling Session - Starting the Session

To start the session, do the following:

- 1 On the *File* Menu, click *Open* to open the *Open Files* dialog box.

- 2 In the *Open Files* dialog box, type the path to the UsrScript directory, press *ENTER* and then double-click *PROFILER\_DEMO.SQL* to open the script in an SQL Editor.

Rapid SQL opens the PROFILER\_DEMO.SQL script in an SQL Editor.

- 3 On the *Tools* menu, click *PL/SQL Profiler* and then click *Start*.

Rapid SQL opens the PL/SQL Profiler - Start dialog box.

- 4 In the *Profile Label* list, type *DemoProfile*.

- 5 Click *OK* to begin the profiling session.

Rapid SQL begins the profiling session.

If this is the first time you start the Rapid SQL PL/SQL Profiler, Rapid SQL displays a message that user profiling tables need to be installed.

- 6 Click *Yes*.

Rapid SQL opens SQL\*Plus to create the tables. You will need to start the profiling session again (see Step 3).

#### SAMPLE PROFILING SESSION - EXECUTING THE SAMPLE SCRIPT

In this step of Sample Profiling Session, you execute the DEMO script.

#### **Sample Profiling Session - Executing the Sample Script**

To execute the sample script, do the following:

- 1 On the *SQL Editor* window toolbar, click *Execute*.

Rapid SQL executes the script and opens a Results Tab.

#### SAMPLE PROFILING SESSION - STOPPING THE SESSION

In this step of Sample Profiling Session, you stop the profiling run.

#### **Sample Profiling Session - Stopping the Session**

To stop the session, do the following:

- 1 On the *Tools* menu, click *PL/SQL Profiler* and then click *Stop*.

Rapid SQL opens the PL/SQL Profiler - Stop dialog box.

- 2 Click *Stop*.

#### SAMPLE PROFILING SESSION - RERUNNING & REEXECUTING THE SESSION

In this step of Sample Profiling Session, you run the same profile session and execute the DEMO script again.

#### **Sample Profiling Session - Rerunning & Reexecuting the Session**

To rerun and reexecute the session, do the following:

- 1 In the *SQL Editor*, click the *Query* Tab.
- 2 On the *Tools* menu, click *PL/SQL Profiler* and then click *Start*.

Rapid SQL opens the PL/SQL Profiler - Start dialog box again.

- 3 Click the down arrow on the *Profile Comment* list and then click *DemoProfile*.
- 4 Click *OK* to begin the profiling session.
- 5 On the *SQL Editor* toolbar, click *Execute*.

Rapid SQL executes the script again and opens the *Results Tab*.

#### SAMPLE PROFILING SESSION - STOPPING & ANALYZING

In this step of Sample Profiling Session, you stop profiling and analyze the runs.

#### Sample Profiling Session - Stopping & Analyzing

To stop and analyze, do the following:

- 1 On the *Tools* menu, click *PL/SQL Profiler* and then click *Stop*.  
Rapid SQL opens the *PL/SQL Profiler - Stop* dialog box again.

- 2 Click *Stop & Analyze*.  
Rapid SQL opens the *PL/SQL Profiler - Run Detail* window.

- 3 Click the *Run* list and then click *Run#x*.

**NOTE:** Rapid SQL assigns a number to each profiling session. These numbers increase incrementally each time you run a profiling session. x= the number that was assigned to your first run.

Rapid SQL PL/SQL Profiler populates the grid with information on the procedure, package body, and package specification.

**NOTE:** For the purposes of this walk-through we have created this package under the account SCOTT.

- 4 Click the *Run* list again and then click the *Run#* for your second run.

Notice this time there is no information on the package specification. It was created in the first run.

- 5 Right-click and then click *Detail*.

Rapid SQL PL/SQL Profiler opens the *PL/SQL Profiler - Unit Detail* window and populates the grid with the average time to execute each unit and the source code. Notice the time to execute *SELECT object\_name*, in the example is 126 ms.

- 6 In the *PL/SQL Profiler - Unit Detail* window, click the *Run* list and then click *Run#x* for your first run.

- 7 Click the *Unit* list and then click *user name.PF\_COUNT\_SYSTEM\_OBJECTS*.

Notice the time to execute *SELECT object\_name* is considerably greater: in the example it is 24476 ms.



8 Right-click and then click *Show Only Hit Lines*.

**THE RAPID SQL PL/SQL PROFILER SHOWS ONLY THE LINES OF CODE THAT EXECUTED.**

9 Right-click and then click [Advanced View](#).

The Rapid SQL PL/SQL Profiler opens the Advanced View window.

10 Continue clicking the *Run* and *Unit* lists to compare the performance of each run and each session.

This concludes the Sample Profiling Session. You can now delete the objects created during the Sample Profiling Session. They are:

- Check Constraints, PLSQL\_PROFILER\_UNITS, PLSQL\_PROFILER\_DATA
- Foreign Keys, PLSQL\_PROFILER\_UNITS, PLSQL\_PROFILER\_DATA
- Package, PF\_COUNT\_TIME\_INTERVAL
- Package functions, WEEKEND\_DAYS\_(), WORKING\_DAYS\_(), YEAR
- S\_ELAPSED\_BETWEEN\_()
- PL/SQL code Profiles, DemoProfile
- Primary Keys, PLSQL\_PROFILER\_RUNS, PLSQL\_PROFILER\_UNITS, PLSQL\_PROFILER\_DATA
- Procedure, PF\_COUNT\_SYSTEM\_OBJECTS
- Sequence, PLSQL\_PROFILER\_RUNNUMBER
- Tables, PLSQL\_PROFILER\_RUNS, PLSQL\_PROFILER\_UNITS, PLSQL\_PROFILER\_DATA



# Index

## Symbols

.mdf/.ndf/.ldf files 229

#define 591

#include 591  
    default file location 83

## Numerics

3D Rounded Docking Tabs. 69

## A

abort tran on log full 274

account, user 270

accounts/accounting  
    setting lock time 257

Activate Logging 447

active  
    transactions, rollback segments 259

Add Cluster Column 385, 461

Add Column 450

Add Database Fragment 449

Add Database Object File(s) to Project Wizard

    Panel 1 677

    Panel 2 678

    Panel 3 678

    Panel 4 for All Objects 679

    Panel 4 for Selected Objects 679

    Panel 5 680

    Panel 6 680

Add Datafile 462

Add Parameter 457

Add Partition 458

After 78

AFTER trigger timing 238

alarms/thresholds  
    code analyst 77

Aliases 130

aliases  
    creating, DB2 LUW 293  
    creating, Sybase 427  
    editor, DB@ LUW 182  
    editor, DB2 Z/OS 200  
    editor, Sybase 273

Alignment index partition setting 249, 255, 269, 277

ALL views 85

Allocating Extent 463

Allow Bit Data setting 198

Allow Duplicate Rows index setting 277

allow nulls  
    domains 346

Allow Nulls setting  
    tables 210

allow nulls setting  
    by default, databases 274  
    by default, databsaes 428  
    user datatype 286

allow parallel option 192

Allow Parallel setting 184, 299

Allow Reverse Scans index setting 186

ALLPAGES 283

alter, rename style 85

Analyze 465

Analyze Code 595

Analyzing  
    Index Partitions 465

ANSI null default 229, 363

ANSI nulls 229, 363

ANSI padding 229, 363

ANSI warnings 229, 363

Append Data setting 193

Append, materialized query tables 187  
applications

    bind to database 274

arithabort 229

array dimension, domains 346

as locator 185

As Locator function setting 203

As Locator option 193

Asc Sort, indexes 248

Ascending option  
    indexes, InterBase 351

asymmetric key 233

Audit table setting 210

auto allocation 263

auto close 229, 363

Auto Commit Status 596

auto create statistics 229, 363

auto format options, ISQL 91

auto format, DDL extract options 80

auto identity 274, 428

Auto Replace feature  
    enabling/configuring 778

auto shrink 229, 363

auto update statistics 229, 363

automatic segment management 263

auto-reload  
    files, ISQL Editor 90

auto-reload, options 89–90

Auto-Save  
    ISQL editor 90

auto-save, files 89–90

Average Size package setting 204

## B

batch files

    Java drop/load 94

Batch Mode execution 78

Begin Transaction ON/OFF 596

Bind Data Cache 469

Bind Package 470

Bind Plan 473

## INDEX

- bind size 96
- Bind Time package setting 204
- Binder package setting 188
- binding
  - application/login to database 274
  - default/rule 241
  - user datatype
    - default binding 286
- bit data, user datatypes 198
- bitmap index type 248
- blob filters
  - creating using wizard 344
  - editing 217
  - support for 131
- body
  - function 202
  - packages 254
  - type bodies 267
- Bookmarks 66, 626
  - Clearing 627
  - Goto 627
  - Moving 626
- Browser
  - Options 76
- Browsers
  - Copying 60
  - Displaying Dependencies 61
  - Extracting DDL 61
  - Main Menu 61
  - Object Types 56
  - Opening 56
  - Printing 59
  - Refreshing 62
  - Searching 60
  - Shortcut Menu 61
  - Sorting 60
  - Toolbar 56, 61
- buffer pools
  - tablespace specification 337
- Buffered Inset package setting 188
- bufferpools
  - indexes 204, 214
  - primary key setting 207
  - tablespace default 196
- Building a Query
  - Statements
    - UPDATE 741
- C
- C
  - method type, DB2 LUW 192
- cache
  - clusters 244
  - materialized view logs option 251, 253
  - property, tables 261
  - sequence setting 190, 259
- called on NULL input option 192
- Called on Null Input setting 299
- carriage return 90
- CASCADE
  - delete rule 246
  - update/delete rule 231
- cascade
  - unique key setting 268
- CASCADE delete rule 183
- Cascading Permissions 661
- case
  - object identifiers 84
- CAST FROM 192
- catalog, datasource
  - specifying storage method 80
- categories
  - stored outlines 253
- categories, datasource
  - assigning 118
  - customizing 121
  - introduced 127
  - user interface element treatment 80
- CCSID
  - database setting 201
  - table setting 193
  - user datatype setting 215
- Certificate setting
  - Asymmetric Key setting 240
- Change Category 478
- Change Container Size 479
- Change Database Comment 479
- Change Password 480
- Change Status 481
- change status object action 482
- CHAR
  - unicode support, Data editor 768
  - unicode support, Results editor 659
- character set, client 78
- character type
  - unicode support, Data editor 768
  - unicode support, Results editor 659
- Check Constraints 131
  - Changing Constraint Status 481, 573
  - Creating for IBM DB2 UDB 502
  - Creating for IBM DB2 UDB for OS/390 and z/OS 502
  - Creating for Microsoft SQL Server 502
  - Creating for Oracle 502
  - Creating for Sybase ASE 502
  - Modifying for IBM DB2 UDB 502
- check constraints
  - editor, DB2 LUW 182
  - editor, DB2 Z/OS 200
  - editor, Oracle 244
  - editor, SQL Server 228
  - editor, Sybase 273
  - table, DB2 LUW 194, 225
- Check Expiration setting 233
- Check Flag table setting 210
- Check Policy setting 233
- Checking Syntax 595
- Checkpoint 482
- clipboard, line end options 90

## INDEX

### CLOB

- unicode support, Data editor 768
- unicode support, Results editor 659
- xmltype display 87

Clone Partition 483

Clone Table 483

Close index setting 204, 214

Close primary key setting 207

Cluster Factor setting 186, 189, 197, 204, 207, 214, 248, 268

Cluster Factor, primary keys 255

Cluster Ration setting 186, 189, 197, 204, 207, 214

clustered indexes, SQL Server 233

clustered setting 286

- indexes 186

- primary keys 235, 239

Clusters 132

- Analyzing 465

clusters 245

- and materialized views 251

- creating, Oracle 383

- editor, Oracle 244

Coalesce 484

Code Analyst 782

- Clone 799

- Cloning Sessions 799

- Common Tasks 782

- Comparing Sessions 799

- Creating a Session 797

- DBMS Support 783

- Debuggers 787

- Deleting a Session 800

- Execute 800

- Executing a Session 800

- Executing SQL 804

- Extracting SQL 804

- IBM DB2 UDB Data 784

- Identifying and Fixing Bottlenecks 798

- Identifying Time-consuming Lines of Code 798

- Installing 787

- Killing a Session Execution 800

- Microsoft SQL Server Data 786

- Oracle Data 786

- Oracle Profiler 787

- Printing Results 802

- Privileges 787

- Product Design 789

- Refresh 801

- Refreshing a Session 801

- Requirements 787

- Saving Results 802

- Schedule 801

- Scheduling a Session 801

- Setting View Options for the Unit Detail Tab 803

- Stopping a Session Execution 800

- Sybase ASE 786

- Tabs 789

- Tutorial 791

- Uninstalling 788

- Unschedule 801

- Unscheduling a Session 801

- Using 796

- Viewing Run Details 802

- Viewing Unit Details 803

- Viewing Unit Summary Information 803

- Wizard 797

code analyst

- options/preferences 77

Code Formatting 595

Code Page package setting 188

Code Size setting 247, 256

- packages 254

Code Templates 779

Code Templates feature

- enabling/configuring 778

Code Workbench 778

- Creating Code Templates 779

- Modifying Code Templates 779

Collation Properties 454

Collecting Column Statistics 586

Collecting column statistics 583

Collecting distribution statistics 583

## INDEX

- Collecting Distribution Statistics for Tables 586
- Collecting Extended Index Statistics 586
- Collecting extended index statistics 583
- Collecting sample statistics 585
- Collecting Sample Statistics for Indexes 586
- colors
  - Query Builder options 97
- colors/fonts
  - color, editors 91
- column mapping
  - foreign keys, Sybase 276
- Column Rename 557
- columns
  - cluster 244
  - dimension 194
  - distinct values in 210
  - distribution key 194
  - filtering for materialized view logs 252
  - indexes, Oracle 248
  - managing, DB2 LUW indexes 186
  - managing, DB2 Z/OS indexes 203
  - return, DB2 functions 185
  - tables, DB2 LUW 193, 210, 224
  - unique keys editor, Sybase 285
  - views, SQL Server 242
- Compatible Level setting 229
- Compile 485
- compiler directives 591
- composite limit 257
- Compress setting 193
- compressed rows
  - percentage of 210
- concat null yields null 229, 363
- concurrent licenses
  - checkout, offline usage 21
- conditions
  - check constraint 273
  - check constraints, SQL Server 228
  - check constraints, DB2 LUW 182
  - check constraints, Oracle 244
- configuration parameters, DB2 LUW 183
- confirmation dialogs
  - data editor options 79
- connect string, database links 246
- connect time 257
- Connection Editor 488
- connections
  - options/preferences 78
  - packages 205
- constraints
  - table, DB2 LUW 194, 225
  - tables, Sybase 284
- containers
  - managing for tablespace 196
- CONTAINS SQL function setting 189
- Convert to Partitioned 489
- copy end of line options 90
- Copy index setting 204, 214
- Copy Object Names 491
- Copy primary key setting 207
- CPU
  - per call profile setting 257
  - per session profile setting 257
- Create Alias 491
- Create Insert Statements 494
- Create Like 495
- Create Synonym 503
- Created setting
  - packages 254
- Creating INSERT Statements 775
- Creating Objects 289
- CurrentData package setting 204
- Cursor Blocking package setting 188
- cursor close on commit 229, 363
- cycle numbers, IBM LUW sequences 190
- cycle numbers, sequences 190
- Cycle When Reach Max/Min setting 190, 259
- D
- DASD storage 210
- Data Blks/Key
  - indexes 248
  - primary keys 255
  - unique keys 268
- data block storage
  - clusters 244–245
  - materialized view logs 252
  - materialized views 251
  - unique keys 268
- DATA CAPTURE NONE/CHANGES 210
- data capture option 187
- Data Dictionary
  - role usage 85
- data dictionary
  - viewing usage permission 84
- Data Editor 525
  - Calendar 769
  - Data Editor Filter 767
  - Date Functions 769
  - Date/Time Format Builder 769
  - Edit Window 765
  - Editing Date and Time Functions 769
  - Editing Date/Time 769
  - Modifying Criteria 767
  - SQL Window 766
  - Time Functions 769
  - Toolbar 766
- Data editor
  - execution mode 78
  - options/preferences 75
  - type support 767
- data editor
  - default column value handling 773
- Data Editor Filter 767
- Data Pages setting 233, 235, 240, 277, 279, 286

## INDEX

- data sources
  - export/import 113
  - filtering objects/nodes view 47
  - managing 112
- Database Links 133
- database links
  - creating, Oracle 386
  - editor, Oracle 246
- Database Object Finder 125
- Database partition Group table setting 193
- Database Search
  - Results 714
- Database Search Wizard 713
  - Panel 1 713
  - Panel 2 714
- databasemanager configuration file parameters 183
- Databases 133
  - Attach 517
  - Checkpoint 482
  - Coalescing 484
  - DBCC 507
  - Detach 517
  - Setting Online/Offline 568
- databases
  - creating, DB2 LUW 294
  - creating, DB2 z/OS 320
  - creating, SQL Server 363
  - creating, Sybase 427
  - editor, DB2 LUW 183
  - editor, IBM DB2 Z/OS 201
  - editor, SQL Server 229
  - editor, Sybase 274
- datafiles
  - status 265, 418
  - tablespace, adding/modifying 418
- DataLink Options 505
- DATAPAGES 283
- Datasource
  - Edit Registration 116
- Datasource Catalog 106
- Datasource Explorer
  - Connections 51
  - DDF 52
  - filtering tree view 47
  - options/preferences 75, 83
  - Parameters 52
  - Subsystem Node 51
- Datasource Groups 127
  - Changing 108
  - Modifying 124
  - Removing 123
- Datasource Management 43
- Datasource Organizer 43
- Datasources 105
  - Auto-Discover 108
  - Connecting 109
  - Discovering 108, 111
  - Properties 115
  - Registering 125
  - Selecting 121
- datasources
  - categories 127
  - indexing 92
  - JDBC driver 120
  - options/preferences 75
  - specifying catalog storage 80
- datatypes
  - domain 346
  - function parameters 203
  - methods, DB2 LUW 192
  - sequence 190, 259
  - user-defined DB@ LUW 197
- datatypes, user
  - editor 286
- date formats
  - data editor 79
- Date Time Format package setting 188
- date/time format
  - results, ISQL 99
- db chaining 229, 363
- DB Protocol package setting 204
- DB2GENERAL parameter style 192
- DBA views 85
- DBAs
  - data dictionary view permission 84
- DBCC 505
  - Databases 507
  - Indexes 510
  - Tables 508
- DBID table setting 210
- DBINFO option 192
- DBInfo parameter 184, 189
- DBINFO setting 299
- DBMS\_JAVA 94
- DBMS\_JAVA package 94
- DBMS\_OUTPUT 82
- dbo use only 229, 274, 363, 428
- DBRM 135
- DDL
  - extraction options 80
- DDL Editors 603
  - Opening 603
- ddl in tran 274, 428
- Deallocate Unused Space 514
- debugger, SQL
  - options/preferences 82
  - options/preferencess 75
- Dec31 package setting 204
- Decimal Point package setting 204
- default binding 241
- default language, logins 233
- default to local cursor 229, 363

## INDEX

- Default Value setting 210
- Defaults 135
- defaults
  - creating, SQL Server 364
  - creating, Sybase 428
  - data editor handling 773
  - editor, SQL Server 230
  - editor, Sybase 275
  - package schema 188
  - segment 286
  - segment name 283
  - tablespace 568
- Defer Prepare package setting 204
- Deferable primary key setting 255
- Deferrable/Deferred unique key setting 268
- Deferred primary key setting 255
- define directive 591
- Definer package setting 188
- degree of parallelism 368, 371, 378
- Degree package setting 188, 204
- Degrees setting 251, 253
- delete
  - trigger firing 285
- delete rule 231, 246
- delete rule CASCADE delete rule 202
- delete rule, update rule 183, 220
- DELETES
  - trigger firing on 238
- delimited identifiers 78
- Density setting 186, 189, 197, 204, 207, 214
- dependencies
  - debugger options 82
- Descending option
  - indexes, InterBase 351
- deterministic option 192
- Deterministic setting 184, 189, 299
- DEV category 127
- Developer Network 19
- device
  - container type
    - files
      - container type 196
- device fragments 274
- Dialog Box
  - Grant Privileges To 662
  - Revoke Privileges From 663
- Dialog Boxes 64
- dictionary extent management 263
- dimensions, columns 194
- directories
  - creating, Oracle 386
  - ETSQLX files 83
  - Oracle utilities 96
  - oracle utilities 79
- disable
  - SQL Server indexes 519
- Disable Job Queues 520
- Disable Triggers 521
- Disk Resize 522
- Display
  - Customizing
    - Output Window 53
- Distinct Keys setting 248, 268
  - indexes 186, 189, 204
  - primary keys 207, 214
  - unique keys 197
- distinct values, column 210
- distribution key columns 194
- docking tab colors 69
- domains
  - creation wizard 345
  - editing 217
  - support for 136
- Drop 523
- DROP statements
  - extract/migrate defaults 81
- DROP STORAGE clause 582
- drop, Java 94
- dump utility, MySQL 96
- duplicate keys 233
- Dynamic Rules package setting 204
- Dynamic setting 250
- Dynamic Sizing settings 259
- E
- Edit Data 525
- Edit registration
  - Datasource 116
- Editing Result Sets 648
  - Adding a Column 652
  - Adding a Row 651
  - Columns 650
  - Copying 649
  - Cutting 649
  - Deleting a Column 653
  - Deleting a Row 653
  - Formatting Result Sets 654
  - Inserting a Column 652
  - Inserting a Row 652
  - Pasting 649
  - Resizing Columns 654
  - Resizing Rows 653
  - Rows 650
  - Sorting Data 654
- Editing SQL Scripts 617
  - Changing Case 623
  - Commenting Code 623
  - Copying and Pasting a Column 625
  - Copying and Pasting Text 624
  - Cutting and Pasting a Column 625
  - Cutting and Pasting Text 624
  - Pasting SQL Statements 619
  - Pasting SQL Syntax 618
  - Selecting a Block of Text 622
  - Selecting Text 621
  - Setting Bookmarks 626
  - Uncommenting Code 623



## INDEX

### editors

- aliases, DB@ LUW 182
- aliases, DB2 Z/OS 200
- aliases, Sybase 273
- blob filters 217
- check constraints, DB2 LUW 182
- check constraints, DB2 Z/OS 200
- check constraints, Oracle 244
- check constraints, SQL Server 228
- check constraints, Sybase 273
- clusters, Oracle 244
- database links, Oracle 246
- databases, DB2 LUW 183
- databases, IBM DB2 Z/OS 201
- databases, Sybase 229
- datbases, Sybase 274
- defaults, SQL Server 230
- defaults, Sybase 275
- domains 217
- encryption keys 218
- exceptions 219
- extended procedures, Sybase 275
- foreign key, Sybase 276
- foreign keys, DB2 LUW 183
- foreign keys, IBM DB2 Z/OS 201
- foreign keys, InterBase 220
- foreign keys, Oracle 246
- foreign keys, SQL Server 231
- functions, DB2 LUW 184
- functions, DB2 Z/OS 202
- functions, Oracle 247
- functions, SQL Server 232
- generators 220
- groups, Sybase 276
- indexes, DB2 z/OS 203
- indexes, IBM LUW 185
- indexes, InterBase 222
- indexes, Oracle 248
- indexes, SQL Server 232
- indexes, Sybase 221, 277
- JDBC Driver 120
- job queues, Oracle 249
- libraries, Oracle 250
- logins, Sybase 278
- materialized query tables 187
- options/preferences 89–90
- outline, Oracle 253
- packages, IBM DB2 Z/OS 204
- packages, IBM LUW 188
- packages, Oracle 254
- plans, IBM DB2 Z/OS 205
- primary keys, IBM LUW 188
- primary keys, Oracle 254
- primary keys, SQL Server 234
- primary keys, Sybase 279
- primary keys, z/OS 207
- procedures, IBM LUW 189, 222
- procedures, Oracle 256
- procedures, SQL Server 235

- procedures, Sybase 280
- profiles, Oracle 258
- roles, SQL Server 223
- roles, Sybase 280
- rollbacksegments, Oracle 258
- rules, SQL Server 236
- sequences, IBM LUW 190
- sequences, Oracle 259
- shadowss, InterBase 223
- stogroups, IBM DB2 Z/OS 209
- structured types, IBM LUW 191
- synonyms, DB2 Z/OS 209
- synonyms, Oracle 260
- tables, DB2 LUW 193
- tables, DB2 Z/OS 210
- tables, InterBase 224, 226
- tables, Sybase 282
- tablespaces, DB2 LUW 195
- tablespaces, DB2 Z/OS 212
- tablespaces, Oracle 262
- triggers, DB2 Z/OS 213
- triggers, IBM DB2 LUW 196
- triggers, Oracle 266
- triggers, SQL Server 238
- triggers, Sybase 285
- type bodies, Oracle 267
- types, Oracle 267
- unique keys, DB2 z/OS 213
- unique keys, InterBase 226
- unique keys, Oracle 267
- unique keys, SQL Server 239
- unique keys, Sybase 285
- user datatypes, DB2 LUW 197
- user datatypes, IBM DB2 Z/OS 214
- user datatypes, SQL Server 241
- user messages, SQL Server 241
- users, DB2 Z/OS 215
- users, IBM DB2 LUW 198
- users, Sybase 287
- views, DB2 LUW 198
- views, DB2 Z/OS 216
- views, SQL Server 242
- views, Sybase 288

EditProcedure table setting 210

Embarcadero SQL Debugger for IBM DB2 UDB 805

Embarcadero SQL Debugger for IBM DB2 UDB for

## INDEX

- Windows/Unix
  - Basics
    - Interface 809
  - DDL Editor 810
  - Features 809
  - Functionality 813, 815
    - Break 821
    - Breakpoints 818–819
    - Close 821
    - Go 819
    - Input Parameters 814
    - Restart 820
    - Run to Cursor 817
    - Step Into 815
    - Step Out 816
    - Step Over 816
    - Stop 820
  - Interface
    - Call Stack Window 812
    - Dependency Tree Window 813
    - Variables Window 811
    - Watch Window 810
  - Options 808
  - Requirements 806
  - Using 822
    - Debugging an SQL Script 823
    - Opening a Debugging Session 822
- Embarcadero SQL Debugger for Microsoft 824
  - Basics
    - Requirements
      - Client 829
      - Server 826
  - Break 842
  - Call Stack Window 834
  - DDL Editor 833
  - Debugging an SQL Script 844
  - Dependency Tree Window 835
  - Features 825
  - Functionality 835
    - Breakpoints 839–840
    - Close 842
    - Go 841
    - Input Parameters 836
    - Run to Cursor 839
    - Step Out 838
    - Step Over 838
  - Interface 832
    - Variables Window 834
  - Opening a Debugging Session 843
  - Options 831
  - Requirements 826
  - Restart 842
  - Step Into 837
  - Stop 841
  - Using 843
    - Debugging a Sample Script 845
  - Watch Window 833
- Embarcadero SQL Debugger for Oracle 855
  - Breakpoints
    - Insert 866
  - Debugging a Sample Script 872
    - Getting Started 872
    - Sample Script 1 874
    - Sample Script 2 877
  - Debugging a SQL Script 870
  - Debugging Sample Script 2 877
  - Exclusion List 856
  - Features 856
  - Functionality 861
    - Breakpoint 866
    - Breakpoints
      - Remove 866, 868
      - Toggle 866
    - Close 869
    - Go 867
    - Input Parameters 862
    - Pause 868
    - Run to Cursor 865
    - Step Into 863
    - Step Out 864
    - Step Over 864
    - Stop 867
  - Interface 858
    - Call Stack Window 860
    - Dependency Tree Window 861
    - Watch Window 859
  - Objects 855
  - Options 857
  - SQL Editor Window 859
  - Using 869
    - Opening a Debugging Session 870
  - Variables Window 860
- Embarcadero SQL Debugger for Sybase 882
  - Features 883
  - Functionality
    - 888
    - Breakpoint 892
    - Breakpoints
      - Insert 892
      - Remove 892, 895

## INDEX

- Toggle 893
  - Close 895
  - Go 894
  - Input Parameters 889
  - Pause 895
  - Run to Cursor 892
  - Step Into 890
  - Step Out 891
  - Step Over 891
  - Stop 894
- Interface
  - 885
  - Call Stack Window 887
  - DDL Editor Window 886
  - Dependency Tree Window 888
  - Variables Window 887
  - Watch Window 886
- Objects 883
- Options 884
- Using 896
  - Debugging an SQL Script 897
  - Opening a Debugging Session 896
- Enable Job Queues 520
- Enable Query Rewrite materialized view setting 251
- Enable Triggers 521
- Enabled primary key setting 255
- Enabled setting 238
- EnableQueryOptimization setting 199
- enabling/disabling
  - check constraints 229
  - materialized query tables 187
  - unique keys 268
- Encoding package setting 204
- Encoding Scheme database setting 201
- encoding, Java 94
- Encrypted setting 238
- encryption
  - rules 236
- encryption keys
  - editing 218
  - support for 136
- end-of-line options 90
- enforce password policy 233
- enforced foreign key 183
- Enforcing Index settings 255
- entry point, blob filters 345
- entry points
  - external functions 348
- environments, packages 205
- environments, plans 206
- Erase primary key setting 207
- Error 525
- error messages
  - SQL Server 378
- Error Size setting 247, 256
  - packages 254
- errors
  - maximum before abort, ISQL 86
- Errors materialized view setting 251
- Estimate Size 526
- ETSQLX files, directory 83
- events
  - trigger 196, 213
- Excel 2000 648
- Excel Spreadsheets
  - Importing 775
- exceptions
  - creation wizard 346–347
  - editing 219
  - support for 137
- Execute 528
- Executing Procedures 595
- Executing Scripts 627
  - Cancelling a Query 630
  - Compilation Error Tab 630
  - Executing from Command History 630
  - Executing Partial Scripts 628
  - Plan Facility 632
  - Script Execution Facility 628
  - Step Execution Facility 631
  - Token Matching Capability 642
- Expected Row Size setting 283
- expired password 425
- Explain Level/Snapshot/Mode setting 188
- Explain package setting 204
- explain plans
  - table creation 87
- Explicit Bind Time package setting 188
- Explicit Permissions 661
- Explorer Tab
  - Creating New Objects 46
  - Datasource Node 44
  - Extracting DDL 46
  - Organizing 44
    - By Object Owner 45
    - By Object Type 45
    - System Objects 45
- export
  - data sources 113
  - default directory 79
  - utility directory 96
- expression
  - check constraint 273
- expressions
  - check constraint 200
  - check constraints, SQL Server 228
- Extended Procedures 137
- extended procedures
  - creating, SQL Server 365
  - creating, Sybase 429
  - editor, Sybase 275
  - library name 275
- extending a container 196
- Extends setting 259
- Extent Size tablespace setting 193

## INDEX

- extents 245
  - assigned to clusters 244
  - clusters 245
  - LOB column usage 262
  - materialized view logs 252
  - materialized views 251
  - primary key settings 255
  - tablespaces 263
  - unique keys 268
- external action option 192
- External Action setting 184, 189, 299
- external applications 69
- external functions
  - creation wizard 348
  - support for 138
- External Name setting 184, 189
- Externally setting 258
- Extract 531
- extraction
  - DDL options 80
- F
- failed login attempts 257
- fenced option 192
- Fenced setting 184, 189, 299
- File Factor setting 235, 239
- File Group setting 233, 235, 239
- File Growth Rate setting 230
- files
  - auto-saving 89–90
  - data editor types 78
  - datasource stroage 80
  - mdf, ndf, ldf 229
  - SQL logging optionsLogging, SQL truncating 95
  - tracking 89–90
  - tracking, ISQL Editor 90
- Fill Factor setting 233, 279
- filtering, Datasource Explorer tree 47
- Final Cal setting 299
- final call option 192
- Final Call setting 184
- Fire On Delete/Update/Insert setting, rule 236
- Fire on Insert/Update/ Delete setting 285
- Fire On Insert/UpdateDelete 238
- First Key index setting 186, 189, 204
- First Key primary key setting 207, 214
- First n Keys setting 197
- Flush Cache 532
- fonts
  - data editor options 79
  - editor options 91
  - Qurey Builder options 97
- For Data setting 210
- For Data table setting 215
- Foreign Keys 138
  - Changing Constraint Status 481, 573
- foreign keys
  - creating, DB2 LUW 296
  - creating, DB2 z/OS 321
  - creating, Oracle 387
  - creating, SQL Server 366
  - creating, Sybase 430
  - editor 276
  - editor, DB2 LUW 183
  - editor, IBM DB2 Z/OS 201
  - editor, InterBase 220
  - editor, Oracle 246
  - editor, SQL Server 231
  - table, DB2 LUW 194, 225
- Format 595
- formatting
  - ISQL 91
- Formatting Code 595
- Formatting Result Sets 654
  - Changing Display Properties 657
  - Display Properties 656
  - Setting Alignment Properties 655
  - Setting Border Properties 656
  - Setting Colors 658
  - Setting Fonts 657
- fragments 274
- Free Packages 533
- Free Page primary key setting 207
- Free Plans 533
- freelists
  - cluster 244
  - clusters 245
  - in indexes 248, 262
  - in unique keys 268
  - primary key settings 255
- Full Keys
  - index setting 186, 189, 204
  - primary key setting 207, 214
  - unique keys setting 197
- Full Screen Mode 74
  - Activating 74
  - Dismissing 74
  - Toolbar 74
- Function Path package setting 188
- Function-Based index setting 248
- Functions 139
  - Compiling 486
  - Synonyms 503
- functions
  - alarms/thresholds, time consumption 77
  - creating, DB2 LUW 298
  - creating, DB2 z/OS 322
  - creating, Oracle 388
  - creating, SQL Server 367
  - editor, DB2 LUW 184
  - editor, DB2 Z/OS 202
  - editor, Oracle 247
  - editor, SQL Server 232
  - password verification 257
  - triggers, path 196, 213

## INDEX

### G

- GBP Cache primary key setting 207
- Generate Numbers in Order setting 190, 259
- Generate Package 534
- Generate Procedure 534
- Generate Reports 565
- Generate Select Statement 534
- Generate Statements 534
- generators
  - creation wizard 350
  - editing 220
  - support for 140
- Gets setting 259
- Globally setting 258
- grace period, password 257
- grid properties, results window 85
- Group Member database setting 201
- Group Member package setting 204
- groups
  - assigning users 287
  - assigning users to 287
  - creating, Sybase 431
  - editor, Sybase 276

### H

- hash
  - cluster type 244
  - partitions, Sybase tables 284
- hash composite 277
- header
  - type bodies 267
- headers
  - packages 254
- heap size
  - JDBC option 94
- heap size, JDBC 94
- HIDD\_DB2\_UIR\_TABLE\_COLUMNS\_PASTE\_DLG 544
- HIDD\_ORA\_ORU\_LOGIN\_CHANGEPASSWORD\_DLG 480
- HIDD\_ORA\_ORU\_WIZARDPAGE\_PARTITION\_PROPERTIES 271
- HIDD\_SQS\_SSU\_DEFAULT\_EDITOR\_BINDINGS\_TAB 230
- high watermarks
  - rollback segment setting 259
- Histogram Statistics 468
- HTML
  - reports, default directories 83
  - template directory 83

### I

- I/O
  - tablespace settings 196
- identity
  - user datatype 286
- Identity Column setting 210
- Identity Gap setting 283

- identity in nonunique index 274, 428
- idle time 257
- IGNORE DELETE TRIGGERS clause 582
- Ignore Duplicate Key setting 233, 235, 239
- Ignore Duplicate Rows index setting 277, 432
- immediate offline mode 481
- IMMEDIATE clause 582
- Impact Analysis 535
- import
  - data sources 113
  - default directory 79
  - utility directory 96
  - utility, MySQL 96
- Import Data Wizard 775
- include columns 186
- include directive 591
- Include Longvar Columns setting 187
- index
  - cluster type 244
- Index Buffer Pool database setting 201
- Index Constraint 536
- Index Level setting 186, 189, 197, 204, 207, 214, 248, 268
- Index Level, primary keys 255
- Index Page Ratio setting 277, 279, 286
- Index Tablespace setting 187
- Index Type index setting 186
- index, datasource 92
- Indexes 140
  - Analyzing 465
  - DBCC 510
  - Deallocating Unused Space 514
  - Partitions 142
  - Rebuilding 551
  - Reorganizing 564
- indexes
  - creating, DB2 LUW 301
  - creating, DB2 z/OS 324
  - creating, InterBase 350
  - creating, Oracle 389
  - creating, SQL Server 368
  - creating, Sybase 431
  - default segment warning 101
  - disabling, SQL Server 519
  - editor, DB2 z/OS 203
  - editor, IBM LUW 185
  - editor, Oracle 248
  - editor, SQL Server 232
  - editor, Sybase 221, 277
  - rebuilding, InterBase 554
  - rebuilding, SQL Server 553
  - reorganizing, SQL Server 564
  - same tablespace indexes warning 100
  - set statistics action 569
  - table, DB2 LUW 194, 225
  - tables, Sybase 283
- Inherit Special Registers setting 184, 189, 299

## INDEX

- Initial Extent
  - index setting 248
  - materialized view log setting 252
  - primary key settings 255
  - rollback segment setting 259
  - table setting 262
  - tablespace setting 263
  - unique key setting 268
- initial extent
  - clusters setting 245
  - materialized view setting 251
- initial heap size, JDBC option 94
- Initial Transactions
  - cluster setting 245
  - index setting 248
  - materialized view log setting 252
  - materialized view setting 251
  - primary key settings 255
  - table setting 262
  - unique key setting 268
- insert
  - trigger firing 285
- Insert Column 450
- Insert Parameter 457
- Insert Partition 458
- INSERTs
  - trigger firing on 238
- Instance 144
- instances
  - materialized view log setting 251, 253
- INSTEAD OF trigger timing 238
- Intra-partition package setting 188
- IP addresses
  - Performance Center 96
- isolation level 87
- Isolation Level package setting 188
- ISQL editor
  - results options 98
- ISQL WIndow
  - Code Templates 779
- ISQL window
  - setting query options 634
- ISQL Window Status Bar 596
- IsValid index setting 248
- J
- JAVA
  - method type, DB2 LUW 192
- Java
  - load/drop options 94
  - options/preferences 75
- Java Classes 144
- Java Resources 145
- Java Sources 145
- jConnect 94–95
- JDBC
  - options/preferences 75, 94
- JDBC driver 120
- JDBC Driver Editor 120
- job config files directory 83
- Job Queues 145
- job queues
  - editor, Oracle 249
- joins
  - type/index Query Builder restrictions 97
- K
- Katakana Charset package setting 204
- Keep Dynamic package setting 204
- Kerberos authentication 117
- key columns, cluster 244
- key size, clusters 244
- Keyboard Commands 71
- Keyboard Emulation Mode 596
- keyboard preferences 69
- L
- Label table setting 210
- language
  - logins 233
  - package setting 204
  - user messages 241
- Language Level package setting 188
- Last Bind Time package setting 188
- Last Modified setting 247
  - packages 254
- Last Runstats setting 210
- layout, query plans 86
- Leaf Blks/Key setting 248, 268
- Leaf Blks/Key, primary keys 255
- Leaf Blocks setting 248, 268
- Leaf Blocks, primary keys 255
- Leaf Pages setting 186, 189, 197, 204, 207, 214
- Libraries 146
- libraries
  - creating, Oracle 392
  - dynamic/status settings 250
  - editor, Oracle 250
  - extended procedure 275
  - modify path/file name 250
- library name, extended procedures 275
- licensing 19
  - offline/online, concurrent 21
- lifetime, password 257
- line feed 90
- line numbers 91
- list composite 277
- list partitions 284
- Live Mode execution 78
- load data
  - default directory 79
- Load Java 94, 536
- range- 336
- local extent management 263
- Locality index partition setting 249, 255, 269, 277
- locator, as 185
- lock connections
  - default settings 86

## INDEX

- Lock Isolation package setting 204
- Lock Scheme setting 283
- Lock Size setting 193
- Lock Table 538
- lock time, account 257
- locking
  - rows/pages 233, 369
- Locksize, materialized query tables 187
- Logged Initially setting 187
- logging
  - database files 229
  - index setting 248
  - materialized query table options 187
  - materialized view logs
    - 251, 253
  - primary key index constraints 255
  - table setting 261
  - unique keys 268
- logging,
  - SQL options 95
- logging, SQL
  - options/preferences 75
- logical reads 257
- LogIndexBuild setting 193
- Logins 146
- logins
  - associating with users 287
  - bind to database 274
  - creating, SQL Server 369
  - creating, Sybase 432
  - editor, SQL Server, editors
    - logins, SQL Server 233
  - editor, Sybase 278
  - extracting with user 80
  - locked/unlocked 278
  - password settings 278, 281
  - profile assignment 258
  - profile limits 257
  - role permissions 281
  - status 278
- LOGSEGMENT 286
- LONG
  - unicode support, Data editor 768
  - unicode support, Results editor 659
- Long Tablespace setting 187
- longvarchar types 659, 768
- M
- Mailing Result Sets 646
- Make Unique setting 186, 189, 197, 204, 207, 214
- manual segment management 263
- map, tablespace 263
- mapping
  - foreign keys, Sybase 276
- master database
  - object creation warning 101
- Master View materialized view setting 251
- Match Full setting 276
- Materialized Query Tables 147
- materialized query tables
  - creating, DB2 LUW 302
  - editor 187
- Materialized View Logs 149
  - creating, Oracle 397
- Materialized Views 148
  - Compiling 487
  - Synonyms 503
- materialized views
  - creating, Oracle 393
- mathc full 276
- Max Extents
  - index setting 248
  - materialized view log setting 252
  - primary key settings 255
  - rollback segment setting 259
  - table setting 262
  - tablespace setting 263
  - unique key setting 268
- Max Row Size setting 240
- Max Size of Non-Leaf Index Row setting 240
- Max Transactions
  - index setting 248
  - materialized view log setting 252
  - primary key settings 255
  - table setting 262
  - unique key setting 268
- MAXDOP 368, 371, 378
- Maximum
  - errors before abort, ISQL 86
- maximum extents 245
- maximum extentsmaterialized views 251
- Maximum Failed Logins setting 281
- Maximum Rows Per Page setting 279, 283, 286
- maximum transactions, clusters 245, 251
- MDI Tab Swapping 69
- Menu Contents 69
- menu preferences 69
- Menus 65
  - Main Menu 65
  - Shortcut Menus 66
- merge publish 229, 363
- messages
  - SQL Server 378
- messages, language 241
- methods
  - adding/editing, structured types 192
- Minimum Extent Size setting 263
- Minimum Extents
  - index setting 248
  - materialized view log setting 252
  - primary key settings 255
  - rollback segment setting 259
  - table setting 262
  - tablespace setting 263
  - unique key setting 268
- minimum extents
  - cluster setting 245
  - materialized views setting 251

## INDEX

- Minimum Password Length setting 281
- Mixed Character Set package setting 204
- Mode materialized view setting 251
- MODIFIED SQL DATA function setting 189
- Modify Cluster Column 385, 461
- Modify Column 450
- Modify Datafile 462
- Modify Parameter 457
- Modify Partition 458
- module
  - blob filters 345
  - external functions 348
- Move Log 539
- MRU Replacement Strategy setting 277, 283, 432
- Multi-node Bound package setting 188
- Must change setting 233
- MySQL
  - utilities options/preferences 75
  - utility options/preferences 96
- mysqldump 96
- mysqlimport 96
- N
- NCHAR
  - unicode support, Data editor 768
  - unicode support, Results editor 659
- nchar
  - unicode support, Data editor 768
  - unicode support, Results editor 659
- NCLOB
  - unicode support, Data editor 768
  - unicode support, Results editor 659
- New Datasource
  - Register 116
- New Project Reverse Engineering Wizard
  - Panel 2 669
  - Panel 3 669
  - Panel 4 670
  - Panel 4 for Selected Objects 670
  - Panel 5 671
- Next Extent
  - primary key settings 255
- Next Extent setting 248, 252, 259, 262–263, 268
- next extent, clusters 245
- next extent, materialized views 251
- NO ACTION delete rule 202
- NO ACTION delete/update rule 183
- no chkpt on recovery 274, 428
- no free space acctg 274, 428
- No Sort index setting 248
- No Sort unique key setting 268
- NO SQL function setting 189
- nodegroups
  - creating, DB2 LUW 304
  - tablespace default 196
- nodes
  - filtering 47
- non-clustered index, same segment warning 101
- NONE update/delete rule 231
- nonunique 186
- nonunique index type 248
- Norely primary key setting 255
- normal offline mode 481
- NOT ENFORCED option 183
- Not For Replication setting 231
- NoValidate index constraint status 255
- ntext
  - unicode support, Data editor 768
  - unicode support, Results editor 659
- Number of Extents unique key setting 268
- numeric roundabout 229, 363
- NVARCHAR
  - unicode support, Data editor 768
  - unicode support, Results editor 659
- nvarchar
  - unicode support, Data editor 768
  - unicode support, Results editor 659
- nvarchar(max)
  - unicode support, Data editor 768
  - unicode support, Results editor 659
- NVARCHAR2
  - unicode support, Data editor 768
  - unicode support, Results editor 659
- O
- OAM Page Ratio setting 277, 279, 286
- OBID table setting 210
- object editors
  - max per open operation 84
- Object Status trigger setting 196, 213
- object types
  - creating, Oracle 421
- objects, database
  - associated with users 287
  - creating 289
  - editors 181
  - filtering from tree view 47
  - owning user, DB2 Z/OS 215
  - tablespace-associated 196
  - transferring ownership 577
  - user dependencies 198
- offline
  - database setting 229, 363
  - datafiles 265, 418
  - license usage 21
  - tablespace status 481
- OLE methods 192
- OMF
  - Oracle Managed Files 125
- OneNote style tabs 69
- online
  - datafiles 265, 418
  - tablespace status 481
- Operative package setting 204
- Optimal Size setting 259
- Optimization Hint package setting 204
- optimization, view queries 198



## INDEX

- Options
  - Browser 76
  - Version Control 100
- options
  - code analyst 77
- options (application) 75
- options/preferences
  - Datasource Explorer 83
  - datasource indexing 92
  - Java 94
  - JDBC 94
  - logging 95
  - Performance Center 96
  - results (ISQL) 98
  - results window grid 85
  - semantic validation 93
  - SMTP 99
- Oracle
  - Data Dictionary usage 85
  - utilities options/preferences 75
- Oracle job queue 249
- Oracle Job Queues 149
- Outlines 150
- outlines
  - creating, Oracle 399
  - editor, Oracle 253
- Output Window
  - maximum entries 84
- Overflow Rows setting 187
- Overhead setting 196
- ownership (objects), transferring 577
- P
- Package Bodies 152
  - Compiling 487
- Package Size package setting 204
- Package Source package setting 204
- Packages 150
  - Compiling 487
  - Rebinding 550
  - Synonyms 503
- packages
  - alarms/thresholds, time consumption 77
  - associated plans 206
  - connections 205
  - creating, Oracle 400
  - editor, IBM DB2 Z/OS 204
  - editor, IBM LUW 188
  - editor, Oracle 254
- packet size 78
- Pad Index setting 233, 235, 239
- Page Fetch Pairs setting 186, 189, 204, 207, 214
- page settings
  - tablespaces 196
- Page Size tablespace setting 193
- Page Writes package setting 204
- pages
  - locking 233, 369
  - number of, table 210
- Pages Reserved setting 240
- parallelism, maximum degree of 368, 371, 378
- Parallel Degree setting 261
- Parallel Execution index settings 248
- Parallel Instances setting 261
- Parallel Query Option
  - materialized view logs 251, 253
- parallel query option 244
- Parameter CCSID setting 299
- Parameter Style setting 184, 299
- parameters
  - configuration DB2 LUW databases 183
  - external functions 348–349
  - functions, DB2 184
  - methods, DB2 LUW 192
  - procedures, IBM LUW 189, 223
  - style option 192
- Parsed Size setting 247, 256
  - packages 254
- Partition Type index setting 249, 255, 269, 277
- Partition Upper Bound 543
- range- 336
- Partitioned Columns 541
- range- 336
- Partitioning 414
  - Composite 414
  - Hash 414
  - List 414
  - Range 414
- partitions
  - editing in DB2 Z/OS tables 210
  - for materialized query tables 187
  - indexes, Oracle 249
  - primary keys, Oracle 255
  - tables, DB2 LUW 193
  - tables, Sybase 284
  - tablespace 336
  - unique keys, Oracle 269
- Password Expiration setting 281
- password, Developer Network 19
- Passwords
  - Changing 480
- passwords
  - database links, associated 246
  - expired 425
  - logins, associated with 278, 281
  - policy 233
  - profile limits 257
  - user, Oracle 270
- Paste Columns for Check Constraint 544
- paster end of line options 90
- paths
  - trigger functions 196, 213
- PDS Name package setting 204
- Percent Free primary key setting 207
- Percent Free setting 248, 262, 268
- percent free, extents 245, 251
- Percent Free,/Used setting 252
- Percent Increase setting 248, 252, 262–263, 268

## INDEX

- percent increase, extents 245
- percent increase, materialized views extents 251
- percent used, extents 245, 251
- Performance Center
  - client option 96
  - connectivity options 96
  - identifiers 96
  - options/preferences 75, 96
- Permissions
  - Cascading 661
  - Explicit 661
  - Using Roles to Grant 662
  - Using Roles to Revoke 662
- permissions
  - DBA/data dictionary 84
- Permissions Management 661
- Piece Size index setting 204, 214
- Piece Size primary key setting 207
- PL/SQL Profiler 898
  - Functionality 900
    - Clear Profile Table 904
    - Flush 901
    - Run Detail 902
    - Run Summary 902
    - Start 900
    - Stop 906
    - Unit Detail 905
      - Advanced View 906
    - Unit Summary 903
  - Requirements 899–900
  - Using 907
    - Sample Session 908
- Place 545
- placement, databases 274
- plans
  - editor, IBM DB2 Z/OS 205
- port numbers
  - Performance Center 96
- precision
  - function parameters 203
  - parameters, DB2 functions 184
- precision, function parameters 203
- Precompile Timestamp package setting 204
- preferences
  - code analyst 77
- preferences (application) 75
- pre-fetch options 82
- Prefetch Size setting 196
- Prefetch Strategy setting 277, 283, 432
- Preview 546
- Preview Create 546
- Primary Keys 153
  - Changing Constraint Status 481, 573
- primary keys
  - creating, DB2 z/OS 329
  - creating, InterBase 351
  - creating, Oracle 400
  - creating, Sybase 434
  - editor, IBM LUW 188
  - editor, InterBase 222
  - editor, Oracle 254
  - editor, SQL Server 234
  - editor, Sybase 279
  - editor, z/OS 207
  - table, DB2 LUW 194, 225
- Primary Space Allocation 207
- primary keys
  - creating, DB2 LUW 304
- private SGA 257
- Privileges
  - Deny Privileges From 664
  - Grant Privileges 662
  - Revoke Privileges To 663
- Procedures 153
  - Compiling 487
  - Recompiling 485
  - Synonyms 503
- procedures
  - alarms/thresholds, time consumption 77
  - creating for Oracle 403
  - creating, DB2 LUW 305
  - creating, DB2 z/OS 331
  - creating, InterBase 352
  - creating, Oracle 403
  - creating, SQL Server 371
  - creating, Sybase 434
  - editor, IBM LUW 189, 222
  - editor, Oracle 256
  - editor, SQL Server 235
  - editor, Sybase 280
- PROD category 127
- Product Design 43
- Profiler
  - Setting Up 899
- Profiles 154
- profiles
  - creating, Oracle 403
  - editor, Oracle 258
  - user-associated 270
- profiling level, code analyst 77
- profiler
  - options/preferences 82

## INDEX

### Projects

- Adding Database Objects 677
- Adding Files 681
- Build Order 676
- Closing 675
- Creating 665
- Creating From a Database 667
- Creating From a Version Control Project 672
- Creating From Existing Files 671
- Creating Without Initialization 673
- Executing Files 680
- File Properties 685
- Management 665
- New Project 666
- Opening a File 682
- Opening a Recent Project 674
- Opening an Existing Project 674
- Project Management 665
- Properties 684
- Setting the Build Order 676
- Subproject Properties 686
- Subprojects 682
  - Creating 682
  - Deleting 683
  - Renaming 684
  - Sorting 684
- Working with 673

- public, database device 246
- published 229, 363

### Q

- QA category 127

- Qualifier package setting 204

### queries

- builder options 97
- optimization for views 199
- optimization, DB2 LUW views 198
- optimization, MQTs 187
- options, detailed descriptions 636
- options, introduced 634
- options, setting 634
- options, when sent 641

### Query Builder

- Auto Join options 97
- code generation options 97
- display/style options 97
- execution options 97
- syntax check options 97

## INDEX

- Query Builder 478
  - Adding a Join 751
  - Auto Join Properties 730
  - Auto Joins 753
  - Auto Layout 753
  - Building a CREATE VIEW Statement 743
  - Building a DELETE Statement 742
  - Building a Query 738
  - Building a SELECT Statement 739
  - Building an INSERT Statement 740
  - Building an UPDATE Statement 741
  - Column Properties 732
  - Creating a Clause 754
  - Creating a GROUP BY Clause 758
  - Creating a HAVING Clause 759
  - Creating a WHERE Clause 754
  - Creating an AND Clause 755
  - Creating an ORDER BY Clause 757
  - Design 725
  - Display Options 730
  - DML Tab 726
  - Editing Joins 751
  - Execution Properties 730
  - Explorer 725
  - General Properties 730
  - Getting Started 735
  - Inserting an AND or OR Clause 756
  - Joins 749
    - Editing 751
    - Inner 749
    - Left Outer 750
    - Properties 733
    - Self 750
  - Keyboard Commands 728
  - Moving Additional Tables to the SQL Diagram Pane 746
  - Moving Columns in the SQL Statement Tree 761
  - Moving Tables 745
  - Moving Tables and Columns in the SQL Statement Tree 761
  - Moving Tables and Views 746
  - Moving Tables in the SQL Statement Tree 761
  - Opening Queries 763
  - options/preferences 75
  - Removing a Join 751
  - Results Tab 727
  - Saving Queries 763
  - Saving Queries as XML 763
  - Schema Change Detection 763
  - Selecting a Statement 736
  - Selecting ALL Columns 748
  - Selecting an Instance or Database 735
  - Selecting Columns 737
  - Selecting Columns in the SQL Diagram Pane 737
  - Selecting Columns in the SQL Statement Tree 737
  - Selecting DISTINCT Columns 748
  - Selecting Tables 745
  - Setting Column Properties 732
  - setting options/preferences 97
  - SQL Diagram Pane 726
  - SQL Statement Pane 726
  - SQL Tab 727
  - Statement Properties 729
  - Subqueries 761
  - Syntax Checker 762
  - Table Properties 731
  - Table Shortcut Menus 728
  - Tables in the SQL Diagram Pane 745
  - Tables/Views Tab 726
  - Toolbar 727
  - Using 734
  - Using Data Editor with Query Builder 764
  - View Properties 731
  - Working with Tables in the SQL Diagram Pane 745
  - Workspace Window Options 725
- Query Optimization package setting 188
- Query Plan Facility 632
- query plans
  - default display type 87
  - default layout 86
- query rewrite, materialized views 251
- Quiesce Instance 547
- Quiesce Tablespace 547
- quotas, user 263
- quoted identifier 229, 363
- quoted identifiers 78
- R
  - range partitions 284
  - range-hash/list composite 277
  - read only 229, 274, 363, 428
  - reads
    - logical 257
- READS SQL DATA function setting 189
- Reassign by Category 549
- Rebind Packages 550
- Rebind Plans 551
- rebuild
  - InterBase indexes 554
  - SQL Server objects 553
- Rebuild Outlines 555
- Recompile 555
- recompute statistics 233
- records
  - maximum length 210
- Recovery Status setting 196
- recursive triggers 229, 363
- Redistribute 555
- Redo Log Groups 156
- reducing a container 196
- REF CURSOR
  - Display in ISQL Window 589
- Refresh Method materialized view setting 251
- Refresh Summary Table 556
- refresh, materialized query tables 187
- registration code 19
- registration, product 19

## INDEX

- registry
  - datasource storage 80
- Release Bound package setting 204
- Rely primary key setting 255
- Rely unique key setting 268
- Remote Servers 156
- Rename 557
- rename action style 85
- Reoptvar package setting 204
- reordering parameters 184
- Reorganize 559
- reorganize
  - SQL Server objects 564
- Replacement Strategy setting 262
- replication
  - check constraints 229
- Report 565
- Reserve Page Gap setting 277, 279, 283, 432
- Reserved Psges setting 233, 235, 277, 279, 286
- Resize 653
- resizing a container 196
- Resource Release package setting 204
- Restart 566
- RESTRICT delete/update rule 183
- RESTRICT WHEN DELETE TRIGGERS clause 582
- RESTRICTdelete rule 202
- RestrictDrop setting 193
- RestrictDrop table setting 210
- Result Sets
  - Closing 647
  - Editing 648
  - Formatting 654
  - Mailing 646
  - Saving 647
  - Saving as XML 647
- results (ISQL) options 98
- Results Editor 643
  - Closing a Result Set 648
  - Closing a Result Window 647
  - Configuring 643
  - Editing Result Sets 648
  - Exporting Data 645
  - Mailing a Result Set 646
  - Multiple Result Sets 643
  - One Attached Tab Window 644
  - One Separate Unattached Window 644
  - Opening Result Files 608
  - Read-only Mode 646
  - Renaming a Result Window 647
  - Saving Results 647
  - Separate Unattached Windows 643
  - Single Result Sets 644
  - Status Bar 646
- Results editor
  - type support 659
- results window
  - grid options 85
  - options/preferences 75
- results, SQL
  - format options 99
  - window/file/set options 98
- returns
  - column, DB2 functions 185
  - scalar, DB2 functions 185
- returns, functions 202
- reuse max, password 257
- REUSE STORAGE clause 582
- reuse time, password 257
- Reverse index setting 248
- Reverse primary key setting 255
- reverse scans, indexes 186
- Reverse unique key setting 268
- Right Outer Joins 750
- Role Locked setting 281
- Roles 156
- roles
  - assigned to logins 278
  - creating, Oracle 405
  - Data Dictionary usage 85
  - editor, SQL Server 223
  - editor, Sybase 280
  - login-associated 233
- Rollback Segments 156
- rollback segments
  - creating, Oracle 406
  - editor, Oracle 258
- rollbacks, materialized views 251
- roundrobin partitions 284
- row
  - length, average 210
- Row Compression setting 193
- Row Movement setting 261
- Row Organization setting 261
- ROWPAGES 283
- rows
  - locking 233, 369
  - total in table 210
- rule binding 241, 286
- Rules 157
- rules
  - creating, SQL Server 372
  - creating, Sybase 435
  - editor, SQL Server 236
- runstats
  - last 210
- types
  - sequences, IBM LUW 190
- S
- save point function setting 189
- scalar, return 185
- scale
  - function parameter 203
  - parameters, DB2 functions 184
  - parameters, IBM LUW procedures 189, 223
- Scale setting 210

## INDEX

- Scheduling
  - ISQL Window Scripts 719
  - Microsoft Task Scheduler 719
- Schema 566
- Schema Path package setting 204
- Scratchpad Length setting 299
- scratchpad option 192
- Scratchpad setting 299
- Scratchpad settings 184
- screen tips 69
- Script Execution Facility
  - Using 716
- Script Files 604
  - Inserting Files 606
- scripts
  - query options 634
- Secondary Space Allocation 207
- Secure Socket Layer encryption 120
- Security Policy setting 193
- segment management, extents 263
- Segment Name setting 277, 279, 283, 432
- Segment setting 286
- segmented tablespaces 336
- Segments 158
- segments
  - creating, Sybase 436
  - LOB column usage 262
  - object creation warnings 100
- Select \* From 567
- select into/bulkcopy/pllsort 229, 274, 363, 428
- semantic validation
  - options/preferences 93
- Sequences 159
  - Synonyms 503
- sequences
  - creating, Oracle 407
  - editor, IBM LUW 190
  - editor, Oracle 259
- Sequential Pages setting 186, 189, 197, 204, 207, 214
- serial number 19
- Server Messages 54
- Servers
  - Configuring 125
  - Disconnecting 110
- servers
  - connection options/preferences 78
- Session Files
  - Executing 125
- Session Recording 125
- sessions
  - defining profile limits 257
- sessions Per User setting 257
- SET DEFAULT update/delete rule 231
- SET NULL delete rule 183
- SET NULL update/delete rule 231
- SET NULLdelete rule 202
- Set Online/Offline 568
- set statistics action 569
- Set Tablespaces Quota 570
- Set UNDO Tablespace 571
- set... options
  - defaults 88
- Setting Access Levels While Updating Statistics 586
- shadows
  - creating, InterBase 354
  - editor, InterBase 223
  - support for 159
- shortcut
  - code templates 780
- Shortcut Menu 46
- shortcuts 69
- Show Only My Objects
  - default settings 84
- Show System Objects
  - default settings 84
- Shrink 572
- Shrinks setting 259
- single user 229, 274, 363, 428
- size
  - user datatype 286
- Size setting 210
- Smart Docking 69
- SMTP
  - encoding 99
  - notification identifiers
    - encoding
      - SMTP email 99
  - options/preferences 76, 99
- Snapshot Logs 160
- Snapshots 160
  - Synonyms 503
- Sort in Tempdb setting 233
- sort order
  - indexes, InterBase 351
- sort, index setting 248
- sorting
  - index columns 186, 203
  - index option 248
- Source Size setting 247
  - packages 254
- Source Size setting 256
- sp\_rename action style 85
- space
  - clusters 244
- space utilization
  - clusters 245
- special registers 184, 189
- specific name, methods, DB2 LUW 192
- SQL
  - logging options 95
  - method type, DB2 LUW 192
  - parameter style 192
  - preprocessing/compiler directives 591
  - query options 634
  - script directory 83
- SQL \* Loader
  - directory 96
- SQL Access Level setting 184, 189

## INDEX

- SQL Editor 589
  - Closing a Query Window Tab 614
  - Closing a Script 617
  - Editing a Script 618
  - Executing Scripts 627
  - Goto 612
  - ISQL Windows 595
  - Navigating 612
  - Opening DDL Editors 604
  - Opening Script Files 604
  - Opening SQL Windows 596
  - Regular Expressions 611
  - Renaming a Query Window Tab 613
  - Saving a Script 616
  - Windows 607
    - Splitting 607
- SQL Error package setting 204
- SQL Scripting 589
  - DDL Editor 603
  - Results Editor 643
  - SQL Editor 589
- SQL Warn/Math Warn package setting 188
- SQL Windows 595
  - Find and Replace 608
  - Replacing 610
  - Searching 609
- SSLencryption 120
- Statistics Recompute setting 233, 235, 239
- statistics.computng, SQL Server indexes 233
- Status 573
- status
  - datafiles 265, 418
  - functions 247
  - libraries 250
  - packages 254
  - procedures 256
  - tablespaces 481
  - triggers 266, 482
- status bar
  - ISQL default 90
- Status package setting 188
- Stay Resident procedure setting 208
- Step Execution Facility 632
- Stogroups 161
- stogroups
  - creating, DB2 z/OS 332
  - editor, IBM DB2 Z/OS 209
  - specified in tablespace 336
- Storage Group database setting 201
- Storage Group primary key setting 207
- String Delimiter package setting 204
- Structured Types 161
- structured types
  - adding/editing methods 192
  - editor, IBM LUW 191
- style
  - aoto-indent 91
  - editor 91
  - ISQL window 91
  - syntax highlighting options 91
  - tabs to spaces 91
- Subpartition type index setting 249, 255, 269, 277
- subscribed 229, 363
- subtypes, blob filters 345
- Summary Definition 576
- Switch Online 576
- Synonyms 163
- synonyms
  - creating, DB2 z/OS 334
  - creating, Oracle 409
  - editor, DB2 Z/OS 209
  - editor, Oracle 260
- syntax
  - checking durng excute, ISQL 86
  - checking, Query Builder 97
- Syntax Check 595
- SysEntries package setting 204
- SYSTEM
  - tablespace, default 568
- system
  - segment 286
  - segment name 283
- System Required setting 186, 189, 197, 204, 207, 214
- SYSTEM tablespace
  - object/user creation warning 101
- T
- Tab Menu 69
- Table Status setting 210
- Table Tablespace setting 187
- Table Type setting 210
- Tables 163
  - Analyzing 465
  - DBCC 508
  - Deallocating Unused Space 514
  - Histograms 465
  - Moving Tables in the SQL Diagram Pane 745
  - Reorganizing 562
- tables
  - creating, DB2 LUW 310
  - creating, DB2 z/OS 334
  - creating, InterBase 355
  - creating, Oracle 410
  - creating, SQL Server 374
  - creating, Sybase 437
  - default segment warning 101
  - editor, DB2 LUW 193
  - editor, DB2 Z/OS 210
  - editor, InterBase 224, 226
  - editor, Sybase 282
  - partitioning for Oracle 414
- Tablespace Buffer Pool database setting 201

## INDEX

- Tablespaces 165
  - Changing Status 481
  - Quiescing 547
- tablespaces
  - and materialized views 251
  - changing status 481
  - coalescing 484
  - creating, DB2 LUW 313
  - creating, DB2 z/OS 336
  - creating, Oracle 415
  - database placement 183
  - datafiles, adding/modifying 418
  - default, setting 568
  - editor, DB2 LUW 195
  - editor, DB2 Z/OS 212
  - editor, Oracle 262
  - extent details 264
  - for materialized query tables 187
  - in indexes 248, 262
  - index creation warning 100
  - inunique keys 268
  - map 263
  - materialized view logs 252
  - object creation warnings 100
  - primary key settings 255
  - segment management 264
  - user default/temporary 270
  - user quotas 263
- tabs
  - to spaces 91
- Task Scheduler 719
- Technical Requirements 13
  - Database 16
  - Hardware 13
  - Operating System 13
- templates
  - code. See Code Template feature
- temporary offline mode 481
- TEXT
  - unicode support, Data editor 768
  - unicode support, Results editor 659
- Text Files
  - Importing 775
- Threadsafe function setting 184, 189
- Threadsafe setting 299
- thresholds/alarms
  - code analyst 77
- time formats
  - data editor 79
- time-based settings
  - profile limits 257
- timeouts
  - debugger 82
  - server connection/query 78
- Timestamp package setting 204
- timing, trigger 238
- timing, triggers 266
- timing, triggers 196
- Toggle 590
- Token Matching Capability 642
- toolbars
  - hiding/displaying 68
  - ISQL default 90
  - options/preferences 89–90
  - preferences 69
- Tools
  - Data Editor 765
  - Database Search 713
  - Embarcadero Products 778
  - File Execution Facility 715
  - Find in Files 712
  - Query Builder 724
  - Scheduling 719
  - Script Execution Facility 715
  - Visual Difference 720
- torn page detection 229, 363
- Total Keys setting 186, 189, 197, 204, 207, 214
- Total Pages Modified setting 233, 235, 240
- Total Rows Modified Since Last setting 240
- Total Sections package setting 188
- Transaction Log
  - Moving 539
- transaction log files 229
- transfer ownership 577
- Transfer Rate setting 196
- Triggers 166
- triggers
  - alarms/thresholds, time consumption 77
  - changing status 482
  - creating, DB2 LUW 314
  - creating, DB2 z/OS 337
  - creating, InterBase 358, 361
  - creating, Oracle 419
  - creating, SQL Server 376
  - creating, Sybase 439
  - editor, DB2 Z/OS 213
  - editor, IBM DB2 LUW 196
  - editor, Oracle 266
  - editor, SQL Server 238
  - editor, Sybase 285
  - events 196, 213
  - function path 196, 213
  - timing 196, 213
  - type 196, 213
- trunc log on chkpt 229, 274, 363, 428
- Truncate 578
- type
  - columns, DB2 z/OS 210
  - user datatype 286
- Type Bodies 168
  - Compiling 488
- type bodies
  - editor, Oracle 267
- Types 167
  - Compiling 488



## INDEX

### types

- alarms/thresholds, time consumption 77
- editor, Oracle 267
- parameters, DB2 functions 184
- parameters, IBM LUW procedures 189, 223
- sequences, IBM LUW 190
- triggers 266

## U

### UI

- Setting Appearance 69
- Setting Classic Microsoft Office 97 style 69
- Setting Microsoft Office 2003 style 69
- Setting Microsoft Office XP style 69

### UNICHAR

- unicode support, Data editor 768
- unicode support, Results editor 659

### Unicode support

- Data editor 767
- Results editor 659

### uniform allocation 263

### unique 186

- indexes, InterBase 351

### unique auto\_identity index setting 428

### unique index type 248

### Unique Keys 168

### unique keys

- creating, DB2 z/OSS 339
- creating, InterBase 354, 359
- creating, Oracle 422
- creating, SQL Server 377
- creating, Sybase 440
- editor, InterBase 226
- editor, Oracle 267
- editor, SQL Server 239
- editor, Sybase 285
- editor, z/OS 213
- table, DB2 LUW 194, 225

### unique setting

- indexes 186

### UNITEXT

- unicode support, Data editor 768
- unicode support, Results editor 659

### UNIVARCHAR

- unicode support, Data editor 768
- unicode support, Results editor 659

### range- 336

### Unlimited Max Size setting 230

### unload data

- default directory 79

### update

- trigger firing 285

### update rule 231

### Update Statistics 582

### UPDATES

- trigger firing on 238

### Updating Table Statistics 586

### usage

- user datatypes 286

### USE statements

- DDL extract options 80

### Used Pages setting 233, 235, 240, 277, 279, 286

### user allocation 263

### User Datatypes 169

### user datatypes

- creating, DB2 LUW 316
- creating, DB2 z/OS 341
- creating, SQL Server 380
- creating, Sybase 441
- editor, DB2 LUW 197
- editor, IBM DB2 Z/OS 214
- editor, SQL Server 241

### user ID, Developer Network 19

### user interface preferences 69

### user messages

- creating, SQL Server 378
- editor, SQL Server 241

### USER views 85

### Users 170

- SCHEMA 170

- SQLID 170

### users

- alias definition 273
- assigned to logins 278
- assigning to groups 276, 287
- creating, DB2 LUW 317, 342
- creating, Oracle 424
- creating, SQL Server 379
- creating, Sybase 442
- database links, associated 246
- editor, DB2 Z/OS 215
- editor, IBM DB2 LUW 198
- editor, Oracle, editors
  - users, Oracle 269
- editor, Sybase 287
- extracting logins with 80
- login-associated 233
- profile assignment 258
- role assignment 223
- tablespace quotas 263

## V

### Valid package setting 204

### Validate index constraint status 255

### Validate package setting 204

### Validate unique key setting 268

### ValidProc table setting 210

### VARCHAR

- unicode support, Data editor 768
- unicode support, Results editor 659

### varchar

- unicode support, Data editor 768
- unicode support, Results editor 659

### VARCHAR2

- unicode support, Data editor 768
- unicode support, Results editor 659

## INDEX

- VC Files Tab
  - Closing 698
  - Opening 698
- VC Files Tab Functionality 696
- VCAT 336
- VCAT Catalog primary key setting 207
- verify function, password 257
- Version Control 689
  - Adding Files to Version Control 693
  - Basics 694
  - Checking In Projects or Files 702
  - Checking Out Projects or Files 700
  - Collapse All 709
  - Delete 698
  - Expand All 708
  - Getting Latest Version 699
  - Integrating with Merant/Intersolv PVCS Version Manager 690
  - Integrating with Microsoft Visual SourceSafe 690
  - Integrating with MKS Source Integrity 691
  - MKS Source Integrity 691
  - Open 698
  - Options 100
  - Properties 708
  - Refresh 709
  - Removing a Project or File 707
  - Removing a Project or File from Version Control 707
  - Sort 698
  - Undoing a Check Out 703
  - VC Files Tab 695
  - Viewing File Differences 705
  - Viewing History 705
  - Visual SourceSafe 690
  - Working with Files 693
- view
  - creating, Sybase 443
- Views 172
- views
  - creating, DB2 LUW 317
  - creating, DB2 z/OS 343
  - creating, Oracle 425
  - creating, SQL Server 381
  - editor, DB2 LUW 198
  - editor, DB2 Z/OS 216
  - editor, SQL Server 242
  - editor, Sybase 288
  - query optimization 199
- Visual Difference
  - Comparing Files 720
  - Navigating 722
  - Opening 720
  - Options 723
  - Printing 722
  - Searching 723
- visual style 69
- Volatile setting 193
- Volatile table setting 210
- Volatile, materialized query tables 187
- W
- Waits setting 259
- Warnings
  - options/preferences 76
- warnings
  - options/preferences 100
  - Query Builder options 97
- Without Login setting 240
- wizards
  - blob filters 344
  - domains 345
  - exceptions 346–347
  - external functions 348
  - generators 350
- wizards, default directories 83
- WORKFILE database type 336
- Wraps setting 259
- Writes setting 259
- X
- XML 648
  - type support, Data editor 767
  - type support, Results editor 659
- xmltype
  - viewing as CLOB 87