



Product Documentation

Rapid SQL[®] XE and Rapid SQL[®] 7.7.1

New Features Guide

Corporate Headquarters
100 California Street, 12th Floor
San Francisco, California 94111

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

© 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.

August 16, 2010

TABLE OF CONTENTS

New features summary	4
Top new features.....	6
InterBase/Firebird support	6
InterBase/Firebird datasources	6
Object management	6
SQL editor support.....	7
Database object management.....	8
Microsoft SQL Server object management enhancements.....	8
IBM DB2 for Linux, Unix, and Windows object management enhancements.....	10
IBM DB2 for z/OS object management enhancements.....	11
Oracle object management enhancements.....	13
Revamped editors and wizards.....	14
ISQL Editor	14
Automated error detection and coding assistance	14
SQL Preprocessor (#define/#include) improvement.....	15
Query option enhancements.....	15
Data editor	16
Improved record creation for tables with all-default column values.....	16
Datasource Management updates	16
Data source filtering	16
Datasource categorization.....	18
Datasource storage options	18
Additional miscellaneous features	20
Additional resources	21
Licensing your Embarcadero Technologies product	21
Embarcadero Technologies product support.....	21
Embarcadero Technologies technical support.....	21
Embarcadero Technologies on the Web	21

NEW FEATURES SUMMARY

InterBase support

This release introduces support for InterBase® SMP, an Embarcadero offering.

Preliminary support includes:

- *Connectivity/Datasource support* - InterBase connectivity has been incorporated into the standard Rapid SQL datasource scheme. You can register datasources, edit datasource details, and use other common datasource management features.
- *Object management* - Editors/Wizards are available for a full set of object types and a preliminary set of object actions are available.
- *SQL editor support* - Includes execution options, syntax checking, and Code Assist features

Unicode support

Rapid SQL now includes Unicode support throughout the application, allowing you to work with data in different languages and display the text correctly on screen.

Datasource management

This release offers a set of new features for managing and working with datasources:

- Existing filtering capabilities have been revised or enhanced, and new, more powerful filtering capabilities are now available.
- Datasource categorization lets you visually distinguish servers used for different purposes in your enterprise.
- In addition to the current Windows registry-based datasource catalog storage, new local file-based and network-shared storage methods are available.

ISQL Editor updates

The SQL parsing components have been upgraded for this release. The new components provide the following benefits when creating or editing scripts:

- Increased performance – Parsing activity is much faster than in previous releases.
- Platform concurrency – The ISQL editor now supports syntax and features of the most recent DBMS versions supported by Rapid SQL.
- Accuracy – Syntax issues in previous releases of Rapid SQL have been corrected.

New editor features are also provided. The ISQL Editor has been enhanced with a set of automated validation and code assistance features:

- On-the-fly syntax-checking locates and diagnoses syntax errors as you type.
- Semantic validation eliminates typographical errors and helps avoid outdated references in specified object names.

- Code assist analyzes statements as you type and offers intelligent suggestions with regard to object names.

Query options customizing the execution environment can now be saved and loaded. The frequency with which the options are sent to the server has been optimized.

Microsoft SQL Server exploitation

Index support in Rapid SQL has been upgraded to account for newer Microsoft SQL Server index functionality. You can now rebuild, reorganize, and disable indexes, primary keys, and unique keys. When creating, dropping or rebuilding indexes, primary keys, or unique keys, online options are available with these operations. This release also offers sparse columns, MAXDOP support, and page and row locking.

IBM DB2 for Linux, Unix, and Windows exploitation

This release introduces support for the following recent upgrades to IBM DB2 for Linux, Unix, and Windows functionality:

- Encoding scheme options of ASCII, UNICODE, or NONE for character or graphic string parameters in procedures and functions.
- Object ownership transfer for object owners and users with SECADM authority.
- Inherit isolation level with or without lock request options on Structured Type methods.

IBM DB2 for z/OS exploitation

Rapid SQL now offers the following DB2 for z/OS upgrades:

- Extended Truncate table support.
- Table cloning actions.
- Updated calculation and display of catalog table and index statistics.
- Long object names.
- Column size and default value modification with simple ALTER TABLE statements.
- Compressed indexes, index key randomization, and expression-based indexes.

Oracle exploitation

When creating or modifying tablespace datafiles, you can now set the availability of datafiles to either ONLINE or OFFLINE.

User interface changes and related features

This release also continues reworking of the object creation/editing/action windows and dialogs. A number of DB2 LUW and SQL Server editors and wizards have been migrated to the new look.

Additional currency or miscellaneous features

In response to customer requests, new or enhanced features are now available:

- Export and import of datasource definitions.
- Removal of reliance on proprietary stored procedures for IBM DB2 for z/OS.

TOP NEW FEATURES

Top new features for this release fall into the following functional categories:


- [InterBase/Firebird support](#)
- [Database object management](#)
- [ISQL Editor](#)
- [Data editor](#)
- [Datasource Management updates](#)
- [Additional currency or miscellaneous features](#)

INTERBASE/FIREBIRD SUPPORT

This release introduces native support for the InterBase DBMS. This release features basic connectivity, object management, SQL editing and execution facilities, and project support.

INTERBASE/FIREBIRD DATASOURCES

InterBase connectivity has been incorporated into the standard Rapid SQL datasource scheme. This release lets you register datasources, edit datasource details, and offers other common datasource management features.



For connected InterBase datasources, the Datasource Explorer offers object type navigation, access to object management features, and common time-saving such as Bookmarks and Favorites.

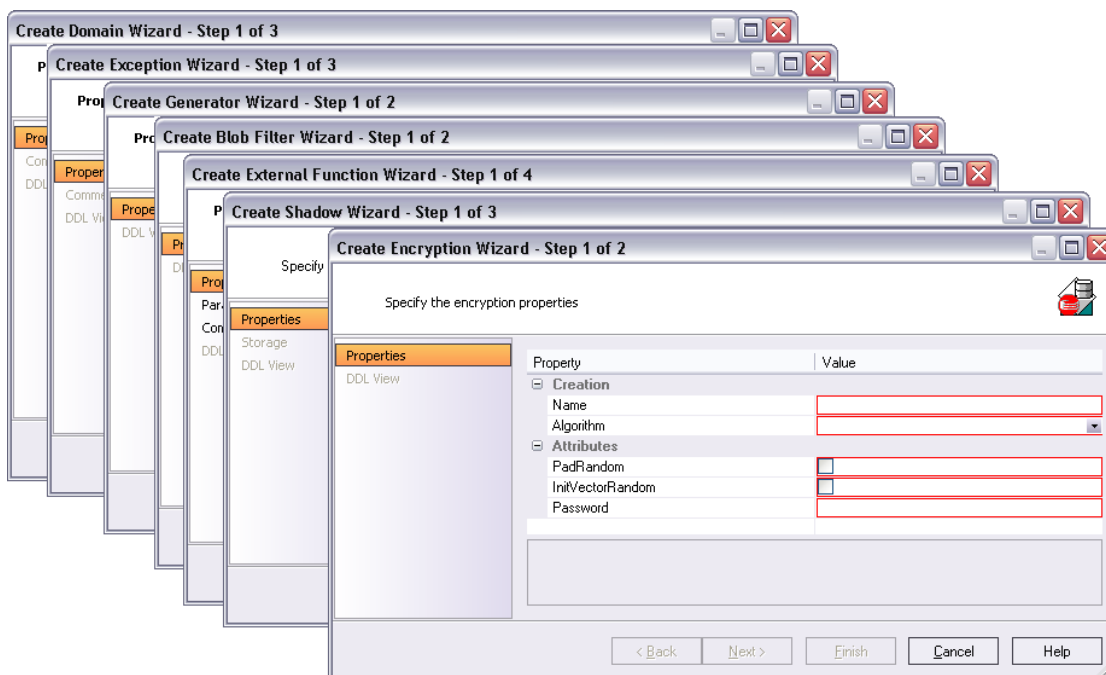
OBJECT MANAGEMENT

This release provides support for the following InterBase object types:

Blob Filters	Domains	Encryption Keys	Exceptions
External Functions	Foreign Keys	Generators	Indexes
Primary Keys	Procedures	Roles	Shadows
Tables	Triggers	Unique Keys	Users
Views			

Editors/Wizards

Similar to support for other DBMS platforms, Rapid SQL provides creation wizards and editors for all supported InterBase object types.



Object actions

While support differs by InterBase object type, support for DROP, EXTRACT, RENAME, and REPORT object actions is provided for most types. A more thorough set of object actions is available for indexes, procedures, tables, triggers, and views.

SQL EDITOR SUPPORT


Full SQL editor support is provided for InterBase. This includes:

- Execution options
- The ability to set query options
- Parser-based error-handling and coding assistance options such as syntactic and semantic validation, code complete, and code templates. For more information, see "Automated error detection and coding assistance" on page 14.

DATABASE OBJECT MANAGEMENT


MICROSOFT SQL SERVER OBJECT MANAGEMENT ENHANCEMENTS

Rebuilding, Reorganizing, and Disabling Indexes, Primary Keys, and Unique Keys
This release introduces the ability to rebuild, reorganize, and disable indexes. The same operations are now available for primary keys and unique keys as well.

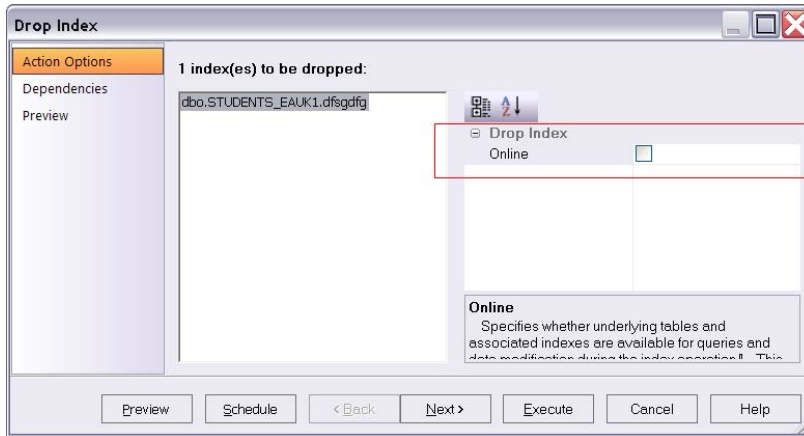


Online options when creating, dropping or rebuilding indexes, primary keys, or unique keys

Rapid SQL now offers ONLINE=ON clause options for operations against clustered indexes, primary keys, and unique keys. Create, drop, and rebuild operations can now be specified as online operations.



Similarly, when dropping a clustered index, primary key, or unique key from a table, you have the option to execute the drop as an online operation.



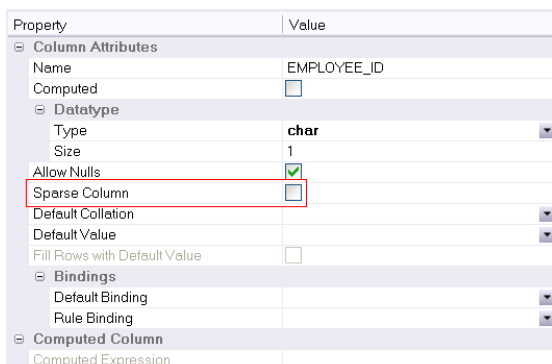
Page and row locking in SQL Server indexes

The wizards/editors for indexes, primary keys, and unique keys now let you enable locking granularity at the page and row level.




Sparse columns

This release introduces support for Sparse columns. Defining a column as sparse optimizes storage of columns that allow NULL values. This property does not apply to the following data types: text, ntext, image, timestamp, geometry, or geography types or to user-defined data types. 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.



MAXDOP support in SQL Server indexes

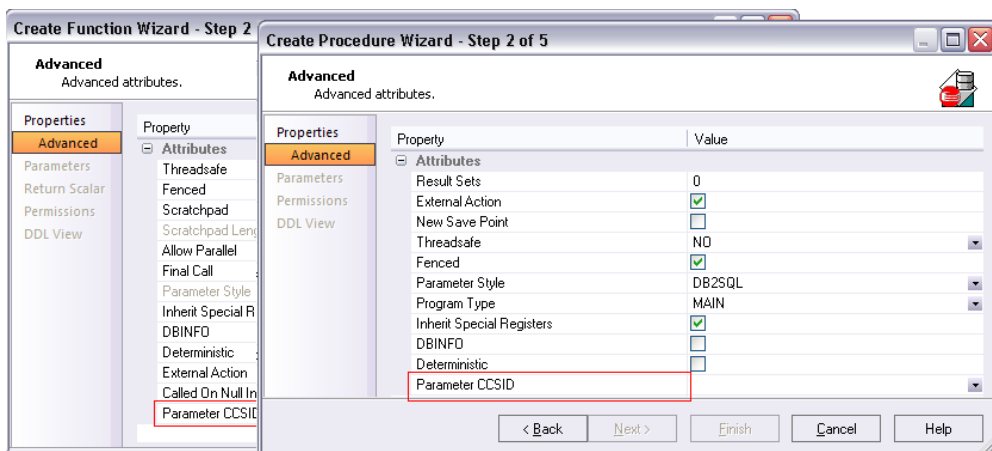
The wizards/editors for indexes, primary keys, and unique keys now support a MAXDOP value, limiting the number of processors used in parallel plan execution.



IBM DB2 FOR LINUX, UNIX, AND WINDOWS OBJECT MANAGEMENT ENHANCEMENTS

Encoding scheme options for procedures and functions

Parameter CCSID properties, available on the editors/wizards for procedures and functions, let you select an encoding scheme of ASCII, UNICODE, or NONE for character or graphic string parameters.



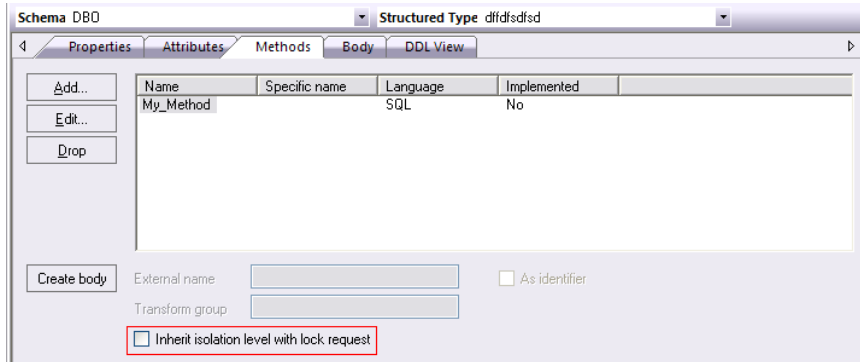
Ownership transfer

Rapid SQL now 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. You can transfer ownership of most DB2 object types.




Isolation Level Lock Request options on Structured Type methods

For structured type methods, you can now specify 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** property value of SQL.



Page and row locking in SQL Server indexes

The wizards/editors for indexes, primary keys, and unique keys now let you enable locking granularity at the page and row level.



IBM DB2 FOR Z/OS OBJECT MANAGEMENT ENHANCEMENTS

Extended Truncate table support

For this release, the truncate operation has been extended to control DROP STORAGE/REUSE STORAGE, RESTRICT WHEN DELETE TRIGGERS/IGNORE DELETE TRIGGERS, and IMMEDIATE clauses of the TRUNCATE TABLE statement.


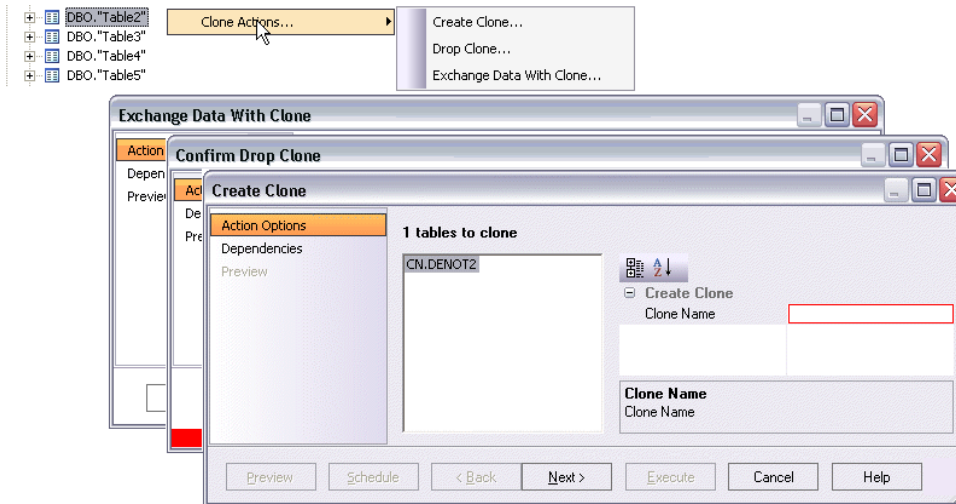


Table cloning actions

Rapid SQL now offers object actions corresponding to CREATE CLONE and DROP CLONE options for ALTER TABLE statements, functionality introduced in IBM DB2 for z/OS version 9. EXCHANGE DATA statement support is also provided, letting you quickly replace a table with its clone.



Catalog table statistics

Calculation and display of certain index and table statistics have been revised to reflect DB2 version 8.x updates to system catalogs. New statistics (and the relevant catalog table column) for this release impacted by this change include:

- Index statistics: **Data Blocks/Key** (AVGKEYLEN) and **DASD storage** (SPACEF)
- Table statistics: **Average Row Length** (AVGROWLEN) - **DASD storage** (SPACEF)

Existing statistics impacted by this change include:

- Index statistics: **Cluster Ratio** (CLUSTERRATIOF)
- Table statistics: **Number of rows** (CARDF) and **Number of pages** (NPAGESF)
- Column statistics: **Number of Distinct Values in the Column** (COLCARDF)

Long object names

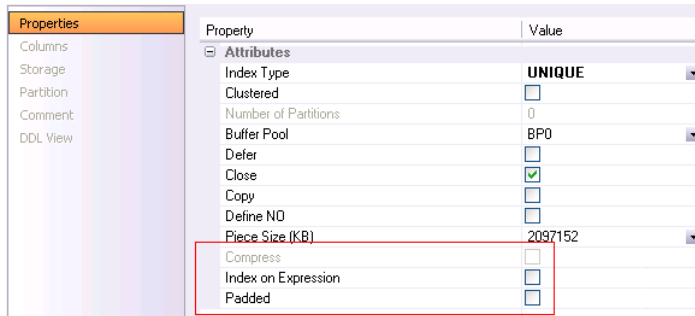
This release introduces support for DB2 Long Object Names, typically allowing identifiers of up to 128 bytes.

Modifying column sizes and default values with a simple ALTER TABLE statement


Changing column sizes and default values can now be handled through a simple ALTER when working against DB2 for z/OS version 8 and above. This eliminates the cost of using an extended ALTER.

Compressed indexes, index key randomization, and expression-based indexes


This release introduces support for these IBM DB2 for z/OS version 9 features.



- *Compressed Indexes* - Compressed indexes are useful when space is at a premium, insertions are sequential, and deletions are few.
- *Index Key Randomization* – Entries are stored in random tree locations to avoid page contention on insertion. Available for indexes created with the **Padded** property selected, the **RANDOM** sort setting is available on the **Columns** panel.



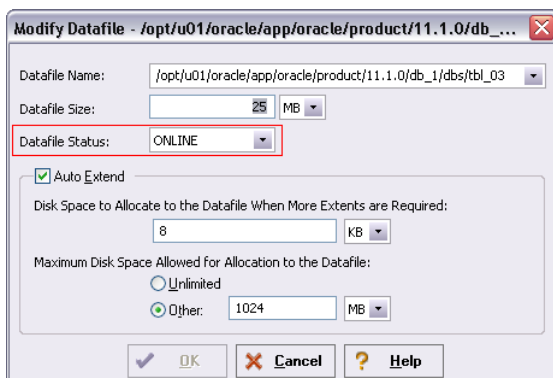
- *Expression-based Indexes* – This feature is activated using the **Index on Expression** property. The key-expression is provided on the **Columns** panel.



ORACLE OBJECT MANAGEMENT ENHANCEMENTS

ONLINE/OFFLINE status of tablespace datafiles

When creating or modifying tablespace datafiles, you can now set the availability of the datafile to either **ONLINE** or **OFFLINE**.



REVAMPED EDITORS AND WIZARDS

This release continues the reworking of object editors and wizards. Aimed at improving workflow and the user interface experience, the new editors and wizards are being reworked for consistency. The following object editors/wizards were upgraded for this release:

- **Microsoft SQL Server** – Roles, User Messages
- **DB2 for Linux, Unix, and Windows** – Structured Types

ISQL EDITOR

AUTOMATED ERROR DETECTION AND CODING ASSISTANCE

This release introduces a number of automated, as-you-type ISQL editor features. While the existing debug and execution options provide assistance at the code-completed or ready-for-testing stage, the new features optimize productivity as you create or edit SQL. The new features minimize typographical errors, let you quickly correct obsolete object name references, and save time and keystrokes.


On-the-fly syntax checking

The ISQL editor now provides real-time, continuous syntax checking. With errors persisting until corrected, this feature lets you quickly find and diagnose coding errors.

```
22 BEGIN
23     v_start_year := TO_NUMBER(TO_CHAR(p_start_date, 'YY'));
24     Encountered unexpected token "v_start_year"
25     p_years_elapsed := v_end_year - v_start_year;
26
27 END YEARS_ELAPSED;
```

Semantic validation


Semantic validation ensures that the names of tables, columns, and views, are properly specified. A semantic error is an indication that a name references an object not found in the database.



Code Complete


Within SELECT, INSERT, UPDATE, DELETE, CALL, and EXEC statements, the Code Complete feature offers intelligent suggestions at those points in the statement where object names are specified. A Code Complete suggestion list lets you choose from a

list of object names. This both eliminates keystrokes and minimizes errors. Suggestions are offered for tables, columns, views, packages, procedures, functions, and synonyms.



Code templates

Code templates are complete code blocks that can be easily added to open windows or scripts with a few keystrokes. When you type CTRL+SPACE, the Code Assist menu opens, letting you select a code template for insertion in the editor window.



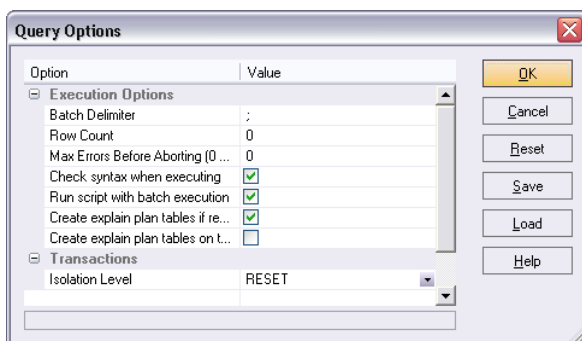
Code templates are defined and maintained using the Code Workbench.

SQL PREPROCESSOR (#DEFINE/#INCLUDE) IMPROVEMENT

Prior to this release, preprocessing a script made #define and #include substitutions in the same ISQL editor window, losing the current state of your opened or edited script. For this release, you now have the option of opening a preprocessed script in a new ISQL editor session, unexecuted.

QUERY OPTION ENHANCEMENTS

You can now save your Query Options settings for subsequent loading in another session. You can also revert your current settings to the Rapid SQL defaults.



In addition, the frequency with which query options are sent to the server has been optimized for locked ISQL sessions. Once locked, a session will only send query options on the first execution. While the session remains locked, only options you explicitly modify are sent to the server on the next execution.

DATA EDITOR

IMPROVED RECORD CREATION FOR TABLES WITH ALL-DEFAULT COLUMN VALUES

For tables with all columns defined as having a default value, there is no longer a requirement to provide a value in at least one column cell when adding records. You can now add multiple records by repeatedly clicking the Insert New Record and Save Current Row buttons.


DATASOURCE MANAGEMENT UPDATES

DATASOURCE FILTERING

Existing filtering capabilities have been revised or enhanced for the current release. In addition, new, more powerful filtering capabilities are now available.

Simple, name-based filtering

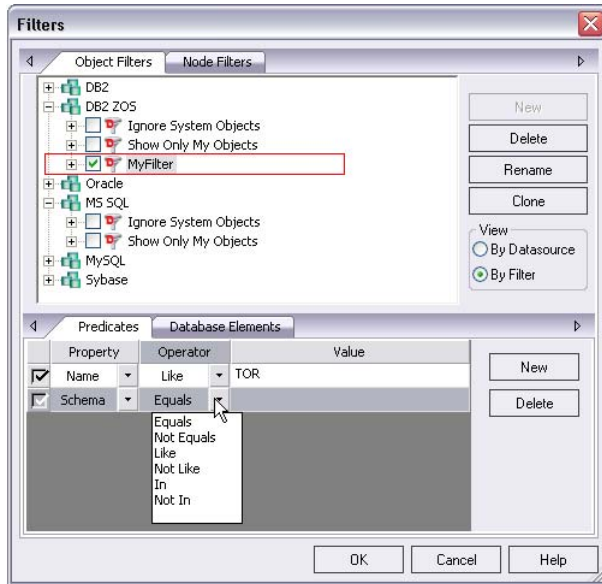
The Filter box at the top of the Datasource Explorer lets you hide datasources based on naming. When you type one or more characters, the tree is updated to show only those nodes of connected datasources whose name contains the characters you typed.



Complex, user-defined object name filtering

Filters that hide system objects or that display only objects owned by the current Rapid SQL user were available in previous releases. Now you can create your own custom filter definitions. Each named filter consists of one or more ANDed conditions, based


on schema or object names. You can also specify the object types to which your object filter applies.



Filters can be defined at the datasource level or at the DBMS level and are enabled and disabled at the datasource level. When a filter is enabled, the tree display for that datasource will include only those objects whose name, schema, and object type satisfy the filter conditions.


Node filtering

Node filtering, defined at the DBMS platform level, lets you hide and show objects based on object type.



Filters node

Each datasource now features a **Filters** node. In addition to any custom object filters you create, this node provides access to two default filters: **Ignore System Objects** and **Show Only My Objects**.




The defaults replace functionality previously available in the Options editor. They can be enabled and disabled at the datasource level but cannot be edited or deleted. Custom filters can also be enabled and disabled but the **Filters** node also provides access to creation and maintenance features.

DATASOURCE CATEGORIZATION

This release introduces the 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. When you assign a category to a datasource, the Datasource Explorer icons for the datasource show a distinctive color-coded category indicator. You can use the Options editor to customize your categorization scheme to include:

- 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



DATASOURCE STORAGE OPTIONS

Rapid SQL now lets you select one of three methods of datasource catalog storage:

- **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. 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. 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.

The Datasource panel of the Options editor lets you choose a method and specify the network location of a Network shared Datasource catalog definition file. Existing functionality such as manually registering datasources and automatic datasource discovery are available for the new storage methods.

In addition, Import/Export Datasource functionality, newly available from the Manage Datasources facility, lets you account for the new storage methods. For example, to set up a **Network shared datasource catalog** scheme, you can build a catalog first, using the **Windows Registry datasource catalog** or **File-based 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.

More generally, the import/export functions can be used to exchange datasource definitions between users employing the different storage methods.

ADDITIONAL MISCELLANEOUS FEATURES

In response to customer requests, a number of new or enhanced features are now available. General enhancements for this release include:

- As mentioned above, the Manage Datasources facility's export and import capabilities are helpful for users employing different datasource storage schemes. Even for users with the same storage scheme, though, these features can save time in organizations with large numbers of datasources and multiple users. For example, 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.

IBM DB2 for Z/OS improvements include the following:

- **Removal of reliance on proprietary stored procedures** - For the past several releases, Rapid SQL relied on a set of Embarcadero-provided stored procedures in supporting certain DB2 z/OS functionality. For this release, that functionality is now provided through native IBM DB2 for z/OS stored procedures. In addition to simplifying installation and setup, this better positions Rapid SQL for further DB2 for z/OS enhancements.

ADDITIONAL RESOURCES

LICENSING YOUR EMBARCADERO TECHNOLOGIES PRODUCT

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

EMBARCADERO TECHNOLOGIES PRODUCT SUPPORT

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

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

EMBARCADERO TECHNOLOGIES TECHNICAL SUPPORT

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

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

EMBARCADERO TECHNOLOGIES ON THE WEB

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