DBArtisan® 8.7 User Guide

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

Contents

Welcome to DBArtisan

Welcome to DBArtisan, the database administration solution for managing IBM DB2 for Linux, Unix, and Windows, IBM DB2 for z/OS and OS/390, Microsoft SQL Server, MySQL, Oracle, and Sybase ASE databases.

The table below describes the major sections of Help.

Section	Description
DBArtisan Tutorial	Explores the features and benefits of DBArtisan. You learn how to make schema modifications, add users and grant database privileges. Finally, you explore some of the powerful database utilities.
Application Basics	Provides conceptual information about DBArtisan's product design and navigation, as well as datasource management and configuring the application.
Using DBArtisan	Provides detailed information about using DBArtisan. Using DBArtisan shows how to use each component of the application. Using breaks each process and procedure into simple step-by-step instructions.
Utilities	Provides detailed information about the extraction, migration and publication of your schema, the loading and unloading of data, database searching, the use of the Script Execution Facility and Query Builder, the scheduling of processes, the backing up and restoration of databases, the use of the Database Monitor, Database Manager, Resource Limits, and Logfile.
Tools	Provides detailed information about DBArtisan's powerful tools to help you create, edit and manage your development environment.
Embarcadero SQL Debugger	Provides detailed information about using the Embarcadero SQL Debugger modules. It includes step-by-step walkthroughs to familiarize you with the features of the Embarcadero SQL Debuggers.
DBArtisan PL/SQL Profiler	Provides detailed information about using the DBArtisan PL/SQL Profiler. Includes a step-by-step walk through to familiarize you with the features of the PL/SQL Profiler.
Capacity Analyst	Provides detailed information about using Capacity Analyst.
Performance Analyst	Provides detailed information about using Performance Analyst.
Space Analyst	Provides detailed information about using Space Analyst.
Backup Analyst	Provides detailed information about using Backup Analyst.

Additional Product Resources

The Embarcadero Web site is an excellent source for additional product information, including white papers, articles, FAQs, discussion groups, and the Embarcadero Knowledge Base.

Go to www.embarcadero.com/support, or click any of the links below, to find:

- Documentation
- Online Demos
- Technical Papers
- Developer Network

>

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.

To save you time, Embarcadero Technologies maintains a knowledge base of commonly-encountered issues and hosts discussion forums that allow users to discuss their experiences using our products and any quirks they may have discovered. To search the knowledge base or participate in forums see our Online Community.

To speak directly with Embarcadero Technical Support, see Contacting Embarcadero Technologies Technical Support, below.

NOTE: Evaluators receive free technical support for the term of their evaluation (14 days).

Contacting Embarcadero Technologies Technical Support

When contacting Embarcadero Technologies Technical Support please provide the following to ensure swift and accurate service:

Personal Information

- Name
- Company name and address
- Telephone number
- Fax number
- Email address

Product and System Information

- Embarcadero product name and version number. This information is found under Help, About.
- Your client operation system and version number.
- Your database and version number.

Problem Description

A succinct but complete description of the problem is required. If you are contacting us by telephone, please have the above information, including any error messages, available so that an Embarcadero Technical Support Engineer can reproduce the error and clearly understand the problem.

There are three ways to contact Embarcadero's Technical Support department:

- Via the Web
- Via Phone
- Via Email

Via the Web

Embarcadero Technical Support provides an online form that lets you open a Support case via the Web. To access this form, go to http://www.embarcadero.com/support/open_case.jsp.

We normally acknowledge the receipt of every case on the same day, depending on the time of submission.

Via Phone

United States

Embarcadero Technologies Technical Support phone number is (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 licensing issues, including Product Unlock Codes, 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.

The Embarcadero Technologies Technical Support fax number is (415) 495-4418.

EMEA

Embarcadero Technologies Technical Support phone number is +44 (0)1628 684 499. The hours are Monday to Friday, 9 A.M. to 5:30 P.M. U.K. time.

For licensing issues, including Product Unlock Codes, call +44 (0)1628-684 494. The hours are Monday to Friday, 9 A.M. to 5:30 P.M. U.K. time

The Embarcadero Technologies Technical Support fax number is +44 (0)1628 684 401.

Via Email

United States

Depending on your needs, send your email to one of the following:

- support@embarcadero.com - Get technical support for users and evaluators
- upgrade@embarcadero.com - Request upgrade information
- key@embarcadero.com - Request a product key
- wish@embarcadero.com - Make a suggestion about one of our products

EMEA

Depending on your needs, send your email to one of the following:

- uk.support@embarcadero.com- Get technical support for users and evaluators
- uk.upgrade@embarcadero.com - Request upgrade information
- uk.key@embarcadero.com - Request a product key
- uk.wish@embarcadero.com - Make a suggestion about one of our products

Technical Requirements

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

- Microsoft Internet Explorer 5.5 or later
- Hardware Requirements
- Operating System Requirements

- Database Support

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:

- 256 MB of memory
- 256 MB of disk space

Operating System Requirements

DBArtisan supports the following operating systems:

- Windows 7
- Windows XP
- Windows Vista

For more information, see XP support and Vista and Windows 7 support.

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 Windos 7 provide two user types:

- Standard user
- Administrator

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

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

Database Support

The table below describes the database platforms DBArtisan® 8.7 User Guide 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, DBArtisan® 8.7 User Guide only supports connecting to a v8 Database on Windows, Linux and Unix.
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. (DBArtisan® 8.7 User Guide 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

IBM DB2 for Z/OS stored procedure requirements

When working against an IBM DB2 for z/OS data source, DBArtisan 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 DBArtisan 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.

About Embarcadero DBArtisan

Embarcadero DBArtisan offers easy access to your license information. The About Embarcadero DBArtisan dialog box not only lets you see version and build information, but also allows you to see your license type and licensed modules (if any), and lets you manage and register licenses. You can also connect to numerous Embarcadero Technologies' links.

DBArtisan Tutorial

DBArtisan is an industry-leading database administration solution for managing Oracle, Microsoft SQL Server, Sybase Adaptive Server, IBM DB2 for Linux, Unix, and Windows, and IBM DB2 for z/OS and OS/390 databases. Its cross-platform capability allows users to efficiently manage heterogeneous database platforms easily using a single front-end tool. Using DBArtisan, users boost their productivity by utilizing a single tool for all their databases, regardless of database vendor.

Database Administrators

DBArtisan enables database administrators to accomplish more with the time they have available in their workday. It eliminates the tedious tasks associated with researching schema dependencies when making object changes. Also included are a host of utilities, which condense DBA tasks taking hours or days down to minutes.

Developers

DBArtisan provides additional administration functionality to database developers over standard development platforms. Using the powerful schema extraction, schema migration, and publication wizards, developers can quickly extract and move schema from development to other environments, as well as create objects much quicker than using old-fashioned hand coding techniques.

Proceed to About This DBArtisan Tutorial

About This DBArtisan Tutorial

This tutorial is intended to help you get started using Embarcadero's DBArtisan, the industry-leading solution for administering enterprise databases from a single point of control. While DBArtisan supports current versions of Oracle, Microsoft SQL Server, Sybase Adaptive Server, IBM DB2, and IBM DB2 for z/OS and OS/390, the examples in this tutorial are Oracle-centric. Unless otherwise noted, all features and functionality highlighted in this tutorial are applicable to all supported platforms.

After completion of this tutorial, you will have the foundation you need to explore the many features and benefits of DBArtisan. You'll have learned how to competently manage the major database administration disciplines using DBArtisan's standard cross-platform console. In addition, you will have a solid understanding of DBArtisan's more advanced Space, Performance and Capacity management capabilities.

This tutorial is divided into 10 sessions:

>

You can use this basic tutorial as a road map of product highlights and also to help you find your own path to explore DBArtisan.

Once you have started, you can select **Help** from the toolbar to find the many additional resources that complement and build on many of the activities shown in this brief tutorial.

Proceed to Session 1: Getting Started with DBArtisan

Session 1: Getting Started with DBArtisan

You can obtain the latest version of the DBArtisan software from the Embarcadero web site at https://downloads.embarcadero.com/free/dbartisan.

Start DBArtisan

1. On the **Start** menu, point to **Programs** > **Embarcadero DBArtisan 8.7.0 > DBArtisan 8.7.0**.

   The first time DBArtisan opens, it displays a message asking if you would like DBArtisan to automatically discover and register datasources.

2. If you have installed and used other Embarcadero tools that use registered datasources, click **Yes**.

   DBArtisan finds any datasources that are defined in the datasource catalog on the machine identified in the Data Source Tab of the Options Editor (see Setting Environment Options below).

   OR

   For the purpose of this tutorial, click **No**.

3. Click **OK** to continue.

DBArtisan provides a Discover Datasources feature that automatically searches the DBMS configuration files on your system for datasources that are not currently registered. The Discover Datasource feature is a dialog box that contains a list of all unregistered datasources found on your network or local machine, including the name of the server or instance and the type of DBMS for each datasource. Once discovered, you have the option to register datasources.

Proceed to Registering Cross-Platform Datasources

Registering Cross-Platform Datasources

You are presented with the Datasource Registration dialog box if you selected **OK** in the previous dialog box. The Datasource Registration dialog box also opens from the main menu when you point to **Datasource**, and then click **Register Datasource.**

Registering a datasource (Oracle Sample)

1. On the first panel of the wizard select **Oracle** and click **Next**.
2. On the next panel, use the **Host**, **Port**, **SID/Service**, **Type**, and **Instance Name** controls to provide connection information for an Oracle datasource.
3. Provide a name that will appear in your datasource list in the **Datasource** field and click **Next**.
4. Use the **User Id**, **Password**, and **Connect As** controls to provide login credentials.
5. To save and encrypt your password, select **Auto-Connect?**, optionally set the **Connect using Windows Authentication** control, and then click **Next**.

6 You can ignore the **Datasource Properties** panel for now as it provides support for an advanced feature. Click **Next**.

7 In the **Datasource Group** tree, place the datasource you are registering.

8 Click **Finish**.

DBArtisan offers the same easy-to-use Datasource Registration process for IBM DB2, Microsoft SQL Server, Oracle, MySQL, and Sybase connections. The connection information only needs to be set up one time for each datasource and can be saved locally or in a common datasource catalog for use by other Embarcadero products.

You can configure Embarcadero database applications to use a datasource catalog stored in the system registry of your machine (local).

All Embarcadero database administration products share the datasource catalog, which means that when you set up your datasource catalog using one product such as DBArtisan, the same list of datasources is available in other Embarcadero Technologies products. Any changes you make to the datasource catalog are reflected in all Embarcadero database management products.

Proceed to General Usability Features

General Usability Features

DBArtisan provides many “user in mind” features that make the product configurable to meet individual needs and preferences. These features are designed to shave time off the tasks that you perform many times on any given working day.

Retaining Datasource Explorer View Settings

1 At the top of the **Explorer** tree, click to expand the drop-down menu.

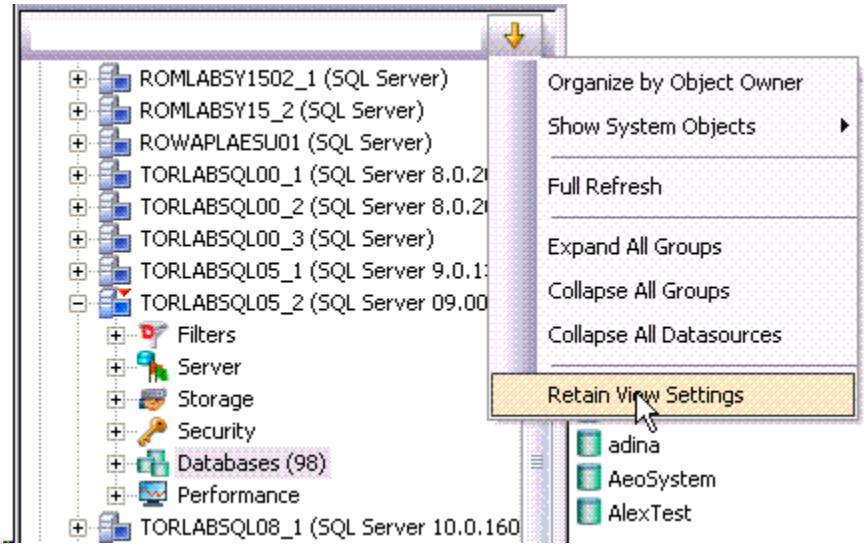

2 Select **Retain View Settings**.

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

Adding Datasource Explorer Bookmarks

1 In the **Explorer** tree, right-click any node.

2 Select **Add Bookmark**.

DBArtisan opens the Add Friendly Bookmark Name dialog box.

3 Click **OK**.

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

Setting Keyboard Shortcuts and Hotkeys

1 In any open space above the **Datasource Explorer**, right-click.

DBArtisan opens a shortcut menu.

2 From the shortcut menu, select **Customize**.

The Customize dialog box opens.

3 In the **Customize** dialog box, open the **Keyboard Tab**.

TIP: The Keyboard Tab can be used to set Keyboard shortcut hot keys for all areas of DBArtisan functionality.

4 Click **Close**.

Opening Most Recently Used Datasources

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

DBArtisan opens the datasource in the Datasource Explorer, ready to work with an active connection.

Proceed to Session 2: Object and Schema Management

Session 2: Object and Schema Management

Advanced Object Creation and Management

DBArtisan provides unparalleled database object management capabilities. Its database platform- and version-specific graphical object editors and wizards enable you to easily create, drop, or alter any of your accessible database objects. The following example walks you through creating and then altering a standard Oracle table object. This concept carries across all supported object types and across all supported platforms.

Creating a Table

1 On the **Datasource Explorer**, expand an Oracle datasource.

2 On the Oracle datasource, right-click the **Tables** node, and then select **Create**.

DBArtisan opens the Table wizard and leads you through the process of creating a table object.

3 Complete the wizard panels.

4 Click **Finish**.

DBArtisan lets you preview any and all generated scripts before you submit them to the database. This is standard for all object related scripts.

>

Changing a Table

Changes to database tables, such as modifying column lengths, inserting new columns, or deleting unneeded ones, can require dropping of a table. This requires knowledge of the underlying object dependencies so that these dependent objects are rebuilt after the table has been re-created.

DBArtisan provides the ability to perform "extended" table alterations by constructing a SQL script with the steps necessary to save off the original data, create the new table, and populate it with the original data. Once these steps are complete, all dependent objects are then rebuilt and permissions re-applied. Following is a sample table change:

1. From the **Datasource Explorer**, expand the **Tables** node, select the table you created in the previous example.
2. Double-click the table.

   OR

   From the **Command** menu, click **Open**.

   DBArtisan opens the Table Editor at the **Columns** tab. In addition to columns, Table Editor tabs let you work with other table properties, indexes, constraints, partitions, and other aspects of table usage/definitions.
3. Select a column from the column list. In the right panel, column attributes are updated to display those of the selected column.
4. Select a different **Type**. The change is reflected in the column list.
5. On the Table Editor toolbar, select the **Alter** button.

   DBArtisan lets you preview the SQL script before you submit it to the database.
6. Click **Cancel** to dismiss the dialog and the close the Table editor window without saving changes.

Working with Object DDL

DBArtisan allows you to easily extract DDL for single or multiple objects using several methods. The most straightforward is described here:

1. On the **Datasource Explorer**, expand an Oracle datasource.
2. On the Oracle datasource, click the **Tables** node.
3. In the right pane of the **Explorer window,** right-click any table or group of tables (SHIFT+CLICK), and then select **Extract**.

   The DDL for all highlighted objects is extracted directly in a DDL Editor where it can be altered, executed, and saved to the database. No intermediary steps are required.
4. Close the Extract window.

Proceed to Advanced Schema Management

Advanced Schema Management

In addition to standardized object support, DBArtisan provides you with advanced Schema management features. These features include full or object-level schema extraction, migration (with or without data), and publication. This example walks you through a simple cross-platform migration between Oracle and SQL Server datasources. Because DBArtisan automatically resolves differences between the disparate DBMS platforms, you can concentrate on what you want to do, rather than how to actually do it. The Schema Migration Wizard sequences the creation of objects in the proper order to eliminate dependency errors. It also has the intelligence to use the fastest method available for copying table data.

Migrating Schema

NOTE: While this example focuses on schema migration, the same wizard principle applies to schema extraction and publication.

1 From the **Utilities** menu, select **Schema Migration**.

DBArtisan opens the Schema Migration Wizard.

2 On the first panel of the wizard, select **Perform New Migration** and click **Next**.

DBArtisan opens the second panel of the Schema Migration Wizard.

3 Select a **Source Datasource** and **Target Datasource** and click **Next**. Provide credentials if prompted.

NOTE: If working with a DBMS that supports multiple-database datasources, you must additionally select a database.

DBArtisan opens the next panel of the Schema Migration Wizard.

4 In the **Object Types** box, select the object types you want to migrate to the target datasource. Use the **Objects** list to select the specific objects you want to migrate. When ready, click **Next**.

5 Choose a set of dependency, owner, script, migration, and action options to use for this migration job. As you make your selections, the changes are reflected in the **Example Preview** pane. Click **Next**.

DBArtisan opens the next panel of the Schema Migration Wizard.

6 You can elect to save the current selecections to be played back in a subsequent migration job. Click **Next**.

7 Click **Finish.**

DBArtisan executes the migration operation.

Proceed to Session 3: Database Security Management in DBArtisan

Session 3: Database Security Management in DBArtisan

DBArtisan can help you efficiently establish and maintain database security and related objects. Whether you are managing an existing production database or setting up a new environment, you will find consistent support across all supported platforms.

Creating a User

NOTE: While this example focuses on creating a new Oracle user, the same wizard-driven principle applies to all security objects (groups, roles, etc.).

1 On the **Datasource Explorer**, expand an Oracle datasource, and then the **Security** node.

2 On the **Security** node, right-click **Users**, and then click **Create**.

DBArtisan opens the User Wizard and leads you through the process of adding a user.

3 Complete the **Create User Wizard**. and when finished click **Finish**.

DBArtisan allows you to preview any and all generated scripts before they are submitted to the database. This is standard for all object related scripts.

4 Click **Execute** to create the new user.

DBArtisan opens the User Editor for the new user. The standard User Editor can be used to manage existing database users as shown below.

Granting and Editing User Privileges

Privileges can be easily granted, revoked, or viewed from within either of two editors within DBArtisan; the User Editor, or the individual object editor (table, procedure, etc.). The User Editor provides a tabbed interface, which can be used to view and modify individual attributes of the user.

1 Double click the user you just created to open the User editor and then open the **Object Permissions** tab.

By default, the tab shows table permissions but you can use the **Object Type** dropdown to work with other permissions.

2 In the permissions grid, right-click a cell coresponding to a permission for a particular table and select **Grant**.

3 Use the same method to **Revoke** the permission you just granted.

4 You could subsequently make additional changes for this user but for now close the User editor window without saving changes.

Proceed to Session 4: DBArtisan Space Management

Session 4: DBArtisan Space Management

Managing space is vital to ensuring the availability and performance of your databases. DBArtisan incorporates many built-in space features that enable you to smartly manage and exploit all aspects of your database's storage. The following example walks you through a review of DBArtisan's built-in support for reporting Oracle tablespace storage and space data.

Managing Object Space

NOTE: While this example is specific to Oracle tablespaces the same concept applies to all supported platforms.

1 On the **Datasource Explorer**, expand any Oracle datasource.

2 On the Oracle datasource, expand the **Storage** node, and then right-click **Tablespaces**.

3 Right-click any tablespace listed in the right pane of the Explorer window, and then click **Open**.

Embarcadero DBArtisan opens the Tablespaces Editor.

4 On the **Tablespaces Editor**, click the **Storage** tab.

The Storage Tab displays and lets you edit the tablespace extent limits.

NOTE: Changes to this tab enable the Alter function on the Tablespace Editor toolbar.

5 On the **Tablespace Editor**, click the **Space** tab.

The Space Tab displays a graphical view of the Free space and Fragmentation Index for the target tablespace.

6 Finally, on the **Tablespace Editor**, click the **Map Tab**.

The Map Tab displays a color-coded map of the objects contained on the tablespace.

NOTE: The map segments are proportional to the actual size of the objects on the tablespace.

>

7 Close the Tablespace editor.

Proceed to Advanced Space Management

Advanced Space Management

For advanced space analysis and management, DBArtisan's optional Space Analyst component contains sophisticated diagnostics to help you pinpoint all space-related problems in your database, as well as an intelligent Reorganization Wizard that can reorganize all or selected parts of your database.

Embarcadero Space Analyst

1 On the Analyst toolbar click the **Space Analyst** button.

The Space Analyst launches in the DBArtisan workspace.

Embarcadero's Space Analyst provides sophisticated diagnostic capabilities to troubleshoot bottlenecks and performance inefficiencies that result in poor space management.

2 Close the Space Analyst window.

Proceed to Session 5: DBArtisan SQL Management

Session 5: DBArtisan SQL Management

DBArtisan provides powerful visual tools for creating and analyzing complex SQL statements and server-side code objects. The following examples walk you through DBArtisan's Visual Query Builder, feature-rich ISQL facility, and some of the advanced analysis and debugging capabilities provided by the Embarcadero SQL Debugger and SQL Profiler.

Building Queries

1 In the right pane of the Datasource Explorer, right-click a table, and then select **Build Query**..

DBArtisan opens the Query Builder.

2 Use the **Locations** window, select the columns to return in the result.

Query Builder generates the query text in the lower SQL window.

TIP: You can build advanced queries using the options supplied in the DML Tab.

3 After the query is built, click the **Execute** button on the Query Builder toolbar.

Query Builder displays results in the lower SQL window.

4 Close the Tablespace editor.

Using the ISQL Window

1 On the **File** menu, click **New**, and then **ISQL**.

DBArtisan opens the ISQL Editor window.

2 Add SQL code via your method of choice (free-form typing, retrieve from a file, paste copied code, etc.).

The ISQL Editor window includes the following features and options:

- The ISQL window highlights all platform and general keywords and provides the options for SQL code formatting, syntax checking and analysis.
- Once code is executed you have control over whether your transaction is committed or rolled back from the database.
- For all open ISQL windows, there are also options for connection locking, scheduling, executing your code across multiple datasources, explain plan generation, and SQL Tuning.

3 Press **F8** prior to SQL execution.

DBArtisan opens the Query Options dialog box that lets you set platform specific Query Options to immediately determine if your code is optimized.

4 Close the ISQL window.

SQL Debugging, Analysis, and Tuning

To analyze and debug your SQL code, DBArtisan provides cross-platform SQL code debuggers. For your Oracle databases DBArtisan offers a robust PL/SQL code profiler that helps you to pinpoint and eliminate “hot spots” within poorly running server-side code. To ensure code efficiency, the ISQL window provides tight integration with Embarcadero’s SQL Tuner, so you can perform multiple “test then tune” iterations without having to leave an open ISQL window.

Debugging SQL

NOTE: While this example is specific to Oracle PL/SQL Debugging the same interface and functionality applies to all of the supported platforms.

1 On the **Datasource Explorer**, expand any Oracle datasource node.

2 On the Oracle datasource, expand the **Procedures** node.

3 In the right pane of the **Explorer**, right-click any stored procedure, and then select **Debug**.

4 If prompted, enter any input parameters the Procedure Execution input window.

5 Click **Continue**.

After the SQL Debugger interface is displayed you can step through code, step into dependencies, set and watch variables, and even perform basic code profiling for each line of executed code.

Profiling SQL

NOTE: The Embarcadero SQL Profiler is available only for Oracle version 8.0 and later.

1 From the **Tools** menu, select **SQL Profiler** and then click **Start**.

2 When prompted add a Profile comment that will help you identify the session later.

3 On the **Datasource Explorer**, expand any Oracle datasource.

4 On the Oracle datasource, expand the **Procedures** node.

5 In the right pane of the **Explorer** window, right-click any stored procedure and select **Execute**. If applicable, enter any input parameter and press **Execute**.

6 After execution is complete, from the **Tools** menu, select **SQL Profiler**, then click **Stop**.

>

7 When prompted select **Stop & Analyze**.

DBArtisan displays the Run Detail view of the profiling data for the current session.

8 To locate the most expensive code for individual units within the run, double-click any of the units and the Unit Detail view is presented. The Unit Detail view allows you to pinpoint the code that may need to be tuned for efficiency.

9 To find previously captured Profiling sessions, on the **Datasource Explorer**, select the **Performance** node, and then **PL/SQL Code Profiling**.

Proceed to Session 6: Job Management

Session 6: Job Management

DBArtisan freely integrates with the Microsoft Windows Task Scheduler, which allows you to schedule virtually any task to run on your computer whenever and however often you want. While this example is specific to an Oracle table redefinition the same concept applies to any job or script that can be scheduled. For this example, you will schedule a table redefinition.

Scheduling Jobs

To schedule a job:

1 On the **Datasource Explorer**, expand any Oracle datasource.

2 On the Oracle datasource, expand the **Tables** node, and then right-click any table.

3 Select **Extract**.

4 From the ISQL window toolbar, click **Schedule**.

The Scheduler Action dialog box opens where you can provide a name, set notifications, and specify an output directory for the new job.

5 After you have completed the dialog box, click **OK**.

This opens the Windows Job Scheduler dialog. For the purposes of this exercise, you can either finish scheduling the task and inspect the results when it completes, or you can click Cancel to proceed to the next session.

Proceed to Session 7: Data Management

Session 7: Data Management

DBArtisan provides comprehensive facilities to help you manage the data in all of your databases. A visual data editor helps you add, change, and delete data from your tables with all referential integrity enforced. You can create insert statements for tables using current data and also extract data as XML documents for certain databases. Rounding out its rich schema management capabilities, DBArtisan also allows you to migrate schema objects and associated table data from one database server to another, across the same or different platforms.

Editing Visual Data

To start the Visual Data Editor, do the following:

1 In the **Datasource Explorer**, right-click any table or tables, and select **Edit Data**.

DBArtisan opens the Data Editor Filter.

2 In **Columns**, select the columns to include in the Edit.

3 You can also filter the editable rows by including your own Select statement.

4 Click **OK**.

CAUTION: In Live mode all changes are applied to the database when you move off of an updated or inserted row. Deleted rows are immediately removed from the database.

NOTE: Batch mode allows you to make changes and then save all when all are complete.

5 Experiment with the editing and execution options and when finished close the editor window..

Creating Insert Statements

1 On the **Datasource Explorer**, select any Oracle datasource.

2 On the Oracle datasource, expand the **Tables** node.

3 In the right pane of the Explorer window, right-click any table, and then select **Create Insert Statements**.

DBArtisan opens the **Create Insert Statements** dialog box.

4 In **Columns**, select the columns you want to include in the Insert statement.

5 You can also filter what rows are included by adding your own Select statement.

6 **OPTIONAL:** Select Owner information and row limits

7 Click **OK**.

The resulting insert statements are created and presented in an active ISQL window. At this point they can be executed immediately, scheduled to later or saved. Note that all extracted insert statements can be run against the same or different databases containing a similar schema.

Extracting Data as XML

NOTE: This feature is available for Oracle 9i, SQL Server 8.0, and Sybase 12.5.1 and later (for **Sybase only**, you must purchase and install a special service from Sybase). The following example is specific to Oracle 9i, but the concept applies to SQL Server 8.0 as well.

1 On the **Datasource Explorer**, expand any Oracle datasource.

2 On the Oracle datasource, 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.

Proceed to Session 8: Performance Management

Session 8: Performance Management

DBArtisan offers a number of options to help you manage the performance of your databases. First, DBArtisan ships with a built-in process monitor that helps you understand who is connected to your database along with each user's current activity and session-related data. For more robust performance details DBArtisan's Performance Analyst add-on is a powerful client-side database monitor that runs fully contained in the DBArtisan console.

Proceed to Monitoring Sessions

Monitoring Sessions

NOTE: While this example is specific to Oracle, the Process Monitor is available for all of the supported platforms.

To start the DBArtisan Process Monitor:

1. On the **Datasource Explorer**, select any Oracle datasource.
2. From the **Utilities** menu, select **Database Monitor**.

The Database Monitor includes the following options and features:

- Highlight any session and any currently running SQL is displayed in the lower pane.
- You can drill-down into a specific session to display session-level statistical details, historical and current wait events along with a working copy of the currently running SQL that can be copied to an ISQL for explain plan generation.
- By using the Monitor drop down options you can display more advanced database-level monitoring data such as locks, blocking locks, hit ratio by user, Top 20 SQL etc.

Proceed to Performance Monitoring

Performance Monitoring

For advanced performance monitoring and management, DBArtisan's optional Performance Analyst provides intelligent diagnostic information and strong drill-down details to help you get to the heart of any episode of performance degradation. Performance Analyst integrates completely with DBArtisan so you can fix any performance problems with a few clicks of the mouse.

Embarcadero Performance Analyst

1. On the **Datasource Explorer**, select any Oracle datasource.
2. From the **Analyst** toolbar, click the **Performance Analyst** button.

   The Performance Analyst opens in the DBArtisan workspace for the target Oracle datasource.

Please see the DBArtisan online help for a detailed walkthrough of all available features and functionality.

Proceed to Session 9: Capacity Management

>

Session 9: Capacity Management

Planning for the future of your critical databases used to be a difficult task. However, DBArtisan's optional Capacity Analyst tool makes it easy to understand where your databases are today and where they are headed in the future. Capacity Analyst lets you track key database metadata and performance metrics over time so you can perform trend analysis on key areas like growth, object fragmentation, I/O, and session load. Like all of the Analyst Series products, Capacity Analyst runs fully contained within DBArtisan so you have access to smart, built-in forecasting mechanisms that allow you to predict when your databases will run out of space and give you the ability to pro actively manage your storage assets, all from the same console.

Advanced Capacity Planning

1. On the **Datasource Explorer**, select any Oracle datasource.
2. From the **Analyst** toolbar, click the **Capacity Analyst** button.

   The Capacity Analyst opens in the DBArtisan workspace for the target Oracle datasource.

Proceed to Session 10: General Utilities and Tools

Session 10: General Utilities and Tools

No evaluation of DBArtisan would be complete without a mention of the general Utilities and Tools that are available across all supported platforms.

Utilities Menu

The main toolbar Utilities menu contains the more advanced DBArtisan features. The available menu items are context-sensitive and version-specific for the selected datasource DBMS platform.

Tools Menu

The main toolbar Tools menu contains those features that are common across all DBMS platforms. All DBArtisan utilities and tools provide a common interface that walks you through all input and execution requirements. All results are consistently presented so you can easily move between features without effort or confusion.

Proceed to Additional Evaluation Resources

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.

Electronic Documentation

Detailed reference documentation is available on the DBArtisan Evaluation CD or online at www.embarcadero.com/support.

>

Online FAQ

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

E-mail Support

You can contact DBArtisan support engineers, consultants and engineers directly by sending inquiries to 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 DBArtisan offers in the way of design, navigation, and application features. The information presented here is high-level and conceptual.

Application Basics is divided into two sections. The table below describes each section:

Section	Description
DBArtisan Product Design	This section describes key interface elements and concepts, including how to use the Datasource Explorer, parallel query execution, and how to navigate and make the most use of DBArtisan's Multiple Document Interface (MDI) environment.
Specifying DBArtisan application and feature options	This section describes how you can configure the application to meet your specific needs.

DBArtisan Product Design

The Datasource Explorer has its own context-sensitive toolbar. As you switch object types in the Datasource Explorer, this toolbar and shortcut menus adapt to expose functionality pertaining to the selected object type. This section describes key interface elements and concepts, such as how to use the Datasource Explorer, parallel query execution, and how to navigate and make the most use of DBArtisan's Multiple Document Interface (MDI) environment.

> **TIP:** DBArtisan lets you use a mouse with a wheel to navigate in active windows.

For more information, see:

Datasource Explorer

DBArtisan Toolbars and Menus

DBArtisan Dialog Boxes

Customizing the DBArtisan Display

Datasource Explorer

The Datasource Explorer provides a visual method for browsing, accessing, and manipulating objects. The Explorer is designed to give you an overall picture of the contents and configuration of your datasources. From the Explorer you can connect, disconnect, filter, and browse the objects in all of the supported datasources on your network.

The Datasource Explorer is a separate tabular window, with a tree in the left pane. When you select an object on the Datasource Explorer, you can view corresponding data in the right pane.

> **TIP:** You can open any number of Explorer tabs so that you can administer and focus on a particular group of servers or objects in one window and accomplish separate tasks in another tab.

The following topics provide basic information on understanding the Datasource Explorer:

- Basic viewing options in the Datasource Explorer
- Datasource Explorer nodes by DBMS platform
- Filtering in the Datasource Explorer

>

- Opening and closing Explorer tabs

For detailed information on datasources, see Datasource Management.

Basic viewing options in the Datasource Explorer

The Datasource Explorer dropdown menu provides basic organization and display options.

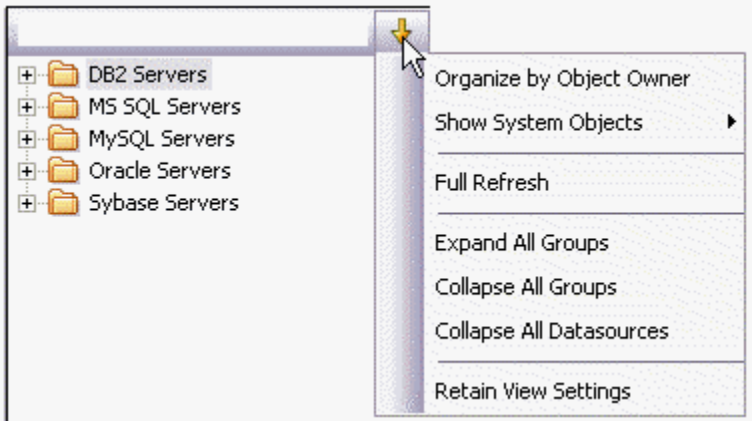

To change basic organization and viewing options:

1 From the Datasource Explorer dropdown menu, select an option. Use the following table as a guide to your selection.

Option	Description
Organize by Object Owner or **Organize by Object Type**	Toggling these options has the following effect: **Owner** - the Datasource Explorer displays objects for each user, categorized by object type. This display mode is most efficient if you are working with databases containing a high number of objects. **Type** - the Datasource Explorer displays all objects, categorized by object type, for all users in the same list. This display mode speeds or slows performance in databases that contain many objects.
Show System Objects	Lets you control display of system objects: NEVER, ALWAYS, or USE DATASOURCE FILTER. For information on filter options, see Filtering in the Datasource Explorer.
Full Refresh	Select to refresh.
Expand All Groups	Select to expand all groups.
Collapse All Groups	Select to collapse all groups.
Collapse All Datasources	Select to collapse all datasources.
Retain View Settings	Select to retain the current state of the Datasource Explorer so that the it opens the same way the next time you start DBArtisan.

TIP: You can also set these option on the Explorer options tab of the Options editor.

Datasource Explorer nodes by DBMS platform

The Datasource Explorer tree organizes and nests subjects as branches, or *nodes*. By expanding or collapsing the tree, you can efficiently browse multiple servers and databases and penetrate to the appropriate level of detail.

The Datasource Explorer tree nodes give you quick access to instance and server configuration, as well as security and performance functionality. The node names depend on the platform. For example, the Instance node in IBM DB2 is very similar to the Server node in Sybase ASE.

When you click the Datasource node of the Explorer, DBArtisan lists all the databases available for that datasource. Similarly, by clicking the Databases Node you can view all the database object types available for that particular datasource. Finally, by drilling down to the object type, DBArtisan displays available objects on in the right pane.

> **TIP:** For IBM DB2 for Linux, Unix, and Windows, DBArtisan offers enhanced navigation of the system, instance, and database nodes of the Datasource Explorer. DBArtisan lets you view, manage, and navigate DB2 for Windows/Unix datasources at the instance node level rather than at the database level.

The Datasource Explorer displays the following nodes:

IBM DB2 for Windows, Unix, and Linux	IBM DB2 for OS390 and Z/OS	Microsoft SQL Server	MySQL	Oracle	Sybase
	Application Processes				
	Databases Node	Databases Node	Databases Node		Databases Node
Filters Node	Filters Node	Filters Node	Filters Node	Filters Node	Filters Node
Instance Node				Instance Node	
Performance Node	Performance Node	Performance Node	Performance Node	Performance Node	Performance Node
Schema Node	Schema Node			Schema Node	
Security Node	Security Node	Security Node	Security Node	Security Node	Security Node
		Server Node	Server Node		Server Node
Storage Node	Storage Node	Storage Node		Storage Node	Storage Node
	Subsystem				

Databases Node

> **NOTE:** The Databases node is available for IBM DB2 for OS390, Microsoft SQL Server, MySQL, and Sybase ASE.

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. Information management centers on the deployment and administration of databases and database applications.

>

The Datasource Explorer List

When you select the Databases node of the Datasource Explorer, DBArtisan lists the databases and displays pertinent information about each of them in the right pane. The tables below describes the information displayed for those DBMS that have a Databases node:

IBM DB2 for OS390 and Z/OS	Microsoft SQL Server	MySQL	Sybase	Node description
AlterEdts				Time when the most recent alter edit was applied
Bufferpool				Name of default bufferpool for the database, blank if it is a system database
CreateEdts				Time when the create statement was created for the database
	Create Date/Time		**Create Date/Time**	When the database was created
Creator				Name of creator
	Database Status		**Database Status**	Status of the database, such as online, offline, read-only or single-user
DB ID				Internal identification number for database
Group Member				The DB2 sharing member name
Index Bufferpool				Indicates name of the default bufferpool for indexes
			Last Dump Tran	The last time that the transaction log for the database was dumped
MIXED_CCSID				Indicates graphic characters
Name	**Database Name**	**Name**	**Database Name**	Name of database
			Online	Yes or no.
			Size	Size of the database in MB
Stogroup				Name of stogroup for the database, blank if it is a system database
			Suspended	Yes or no.
Timestamp				Indicates timestamp
Type				Indicates type of database; W indicates a work file

For more information, see Datasource Explorer nodes by DBMS platform.

NOTE: MySQL does not support the concept of owner and databases for MySQL are registered manually. When you select the database node for MySQL, you see the name of the database.

Filters Node

The Filters node contains default Datasource Explorer tree filters and any custom filters you create for that datasource or create across that DBMS type. Controls let you enable or disable the default and custom filters. Options for creating and maintaining cusom filters are provided. For details, see Filtering in the Datasource Explorer.

Instance Node

NOTE: The Instance node is available for IBM DB2 and Oracle.

DBArtisan places Instance as the first level of information under the Datasource node in the Datasource Explorer. Instance includes:

- DB Manager Configuration
- Datasources
- Version
- Configuration
- Oracle Job Queues

Accessing an Instance

Connect a datasource to access an instance. After you have connected to your datasource:

1. On the **Datasource Explorer**, expand the target datasource node.
2. On the **Datasource Explorer**, expand the **Instance** node.

Editing Node Connection

1. On the **Datasource Explorer**, expand the target datasource node.
2. On the **Datasource Explorer**, right-click the **Instance** node.
3. Select **Edit Node Connection**.

   DBArtisan opens the Connect Login for Node dialog box.

For more information, see Datasource Explorer nodes by DBMS platform.

Database Configuration

The Database Configuration node displays under the Instance node of the Datasource Explorer. When you select the Configuration node, DBArtisan displays configuration information about your datasource in the right pane of the Explorer window. You can use the Datasource Explorer or shortcut menu to:

- Open the Instance Parameters dialog box.
- Open the Edit Configuration dialog box.
- Extract SQL.
- Execute an HTML report.

Opening the Instance Parameters dialog box

Connect to a datasource to access an instance.

1. On the **Datasource Explorer**, expand the target datasource node.
2. On the **Datasource Explorer**, expand the **Instance** node.
3. Click **Configuration**.

   DBArtisan displays the configuration information in the right pane of the Explorer window.

>

4 On the **Datasource Explorer** toolbar, click **Open**.

OR

In the right pane of the **Explorer** window, right-click the target configuration parameter, and then click **Open**.

OR

In the right pane of the **Explorer** window, double-click the target configuration parameter, and then click **Open**.

DBArtisan opens Instance Parameters dialog box.

Version Information

NOTE: Version is available for MySQL, Microsoft SQL Server and Oracle.

DBArtisan places Version under either the Server or Instance node of the Datasource Explorer. When you select Version, DBArtisan shows version information about your datasource in the right pane of the Explorer window. Version information varies according to the database platform's manufacturer.

Displaying Version Information

Connect to a datasource to access version information. After you have connected to a datasource, do the following:

1 On the **Datasource Explorer**, expand the **Datasource** node.

2 On the **Datasource Explorer**, expand the **Server** or **Instance** node.

3 Click **Version**.

DBArtisan displays version information in the right pane of the Explorer window**.**

Performance Node

NOTE: The Performance node is available for all DBMSs.

The Performance node on the Datasource Explorer lets you quickly access to the following features of the Database Monitor:

IBM DB2 for Windows, Unix, and Linux	IBM DB2 for OS390 and Z/OS	Microsoft SQL Server	MySQL	Oracle	Sybase
Blocking Locks		Blocking Locks		Blocking Locks	Blocking Locks
	Buffer Pools				
	Image Copies				
Locks	Locks	Locks	Locks	Locks	Locks
Processes		Processes	Processes	Processes	Processes
	Threads				
	Traces				
	Utilities Node				

You can select the Performance node and view the corresponding statistics in the right pane of the Explorer window. To utilize all Database Monitor functionality, on the Utilities menu, click Database Monitor.

>

IBM DB2 for Linux, Unix, and Windows and IBM DB2 OS390

Performance is a result of the DB2 API calls. The DB2 API calls are subject to change depending on the version and DB2 SDK call interface of the DB2 server.

NOTE: Performances do not appear in the system catalog tables.

For more information, see Datasource Explorer nodes by DBMS platform.

Image Copies

NOTE: Image Copies are available for IBM DB2 for z/OS and OS/390 only.

The Image Copies node shows the image copies on the subsystem. Image Copies are a full image copy of all pages in a tablespace, partition, data set, or index space.

For more information, see Datasource Explorer nodes by DBMS platform.

Processes

NOTE: Processes is available for IBM DB2, Microsoft SQL Server, MySQL, Oracle, and Sybase ASE.

The Processes node on the Datasource Explorer displays processes that are currently running on your DBMS.

For more information, see Datasource Explorer nodes by DBMS platform.

Locks

NOTE: Locks are available for all DBMSs.

The Locks node on the Datasource Explorer displays the lock objects held on the table and row IDs for your database. Locks are a way of controlling concurrent events and access to data. DBArtisan lets you view all locks associated with a database. To utilize all Database Monitor functionality, on the Utilities menu, click Database Monitor.

For more information, see Datasource Explorer nodes by DBMS platform.

Refreshing the Locks Tab

1. On the **Datasource Explorer**, select the **Locks** node.

   DBArtisan displays the locked sessions in the right pane of the Explorer window.

2. On the **Datasource Explorer** toolbar, click **Refresh**.

Blocking Locks

The Blocking Locks node on the Datasource Explorer displays only the locks that are blocked. To utilize all Database Monitor functionality, on the Utilities menu, click Database Monitor.

For more information, see Datasource Explorer nodes by DBMS platform.

Refreshing the Blocking Locks Tab

1. On the **Datasource Explorer** select the **Blocking Locks** node.

   DBArtisan displays the blocked sessions in the right pane of the Explorer window.

2. On **Datasource Explorer** toolbar, click **Refresh**.

Threads

NOTE: Threads are available for IBM DB2 for OS390.

Threads are control structures in IBM DB2 for z/OS and OS/390, which:

- Describe application connection.
- Trace application progress.
- Process resource functions.
- Communicate SQL statement status after SQL execution.

There are three types of threads:

1 Active Threads

2 Indoubt Threads

3 Inactive Threads

The table below describes the information DBArtisan displays in the right pane when you select the Active, Indoubt or Inactive Thread nodes on the Datasource Explorer:

Column Name	Description
Token	Token number designation
Name	Location name of the system
Status	Indicates status of the conversation or socket
A	Indicates event control
Req	Indicates Req number value
ID	Program currently running
Auth ID	User authorization ID
Plan	Bound application plan
ASID	Address space ID
Location	IP address of client program
Session ID	Session Instance Identifier

For more information, see Datasource Explorer nodes by DBMS platform.

Active Threads: Active threads are the processes running on your database. When you select the Active Thread node on the Datasource Explorer, DBArtisan displays the active threads that are currently running on the IBM DB2 for z/OS and OS/390 database. To utilize all Database Monitor functionality, on the Utilities menu, click Database Monitor.

Refreshing Active Threads

1 On the **Datasource Explorer**, select the **Active Threads** node.

DBArtisan displays the active sessions in the right pane of the Explorer window.

2 In the right pane of the **Explorer** window, click the target active thread(s).

3 On the **Datasource Explorer** toolbar, click the **Refresh** button.

OR

In the right pane of the **Explorer** window, right-click the active thread(s), and then click **Refresh**.

DBArtisan refreshes the active thread.

Killing Active Threads

1 On the **Datasource Explorer**, select the **Active Threads** node.

DBArtisan displays the active sessions in the right pane of the Explorer window.

2 In the right pane of the **Explorer** window, click the target active thread(s).

3 On the **Datasource Explorer** toolbar, click the **Kill** button.

OR

In the right pane of the **Explorer** window, right-click the active thread(s), and then click **Kill**.

DBArtisan refreshes the active thread(s).

For more information, see Datasource Explorer nodes by DBMS platform.

Indoubt Threads: Indoubt threads are processes waiting for resources to become active. When you select the Indoubt Thread node on the Datasource Explorer, DBArtisan displays the indoubt threads that are currently running on the IBM DB2 for z/OS and OS/390 database. To utilize all Database Monitor functionality, on the Utilities menu, click Database Monitor.

Refreshing Indoubt Threads

1 On the **Datasource Explorer**, select the **Indoubt Threads** node.

DBArtisan displays the indoubt sessions in the right pane of the Explorer window.

2 In the right pane of the **Explorer** window, click the target indoubt thread(s).

3 On the **Datasource Explorer** toolbar, click the **Refresh** button.

OR

In the right pane of the **Explorer** window, right-click the indoubt thread, and then click **Refresh**.

DBArtisan refreshes the indoubt thread.

Killing Indoubt Threads

1 On the **Datasource Explorer**, select the **Indoubt Threads** node.

DBArtisan displays the indoubt sessions in the right pane of the Explorer window.

2 In the right pane of the **Explorer** window, click the target indoubt thread(s).

3 On the **Datasource Explorer** toolbar, click the **Kill** button.

OR

In the right pane of the **Explorer** window, right-click the indoubt thread, and then click **Kill**.

DBArtisan kills the indoubt thread.

For more information, see Datasource Explorer nodes by DBMS platform.

>

Inactive Threads: Inactive threads are processes that are idle. When you select the Inactive Thread node on the Datasource Explorer, DBArtisan displays the inactive threads that are currently running on the IBM DB2 for z/OS and OS/390 database. To utilize all Database Monitor functionality, on the Utilities menu, click Database Monitor.

Refreshing Inactive Threads

1 On the **Datasource Explorer**, select the **Inactive Threads** node.

DBArtisan displays the inactive threads in the right pane of the Explorer window.

2 In the right pane of the **Explorer** window, click the target inactive thread(s).

3 On the **Datasource Explorer** toolbar, click the **Refresh** button.

OR

In the right pane of the **Explorer** window, right-click the inactive thread(s), and then click **Refresh**.

DBArtisan refreshes the inactive thread.

Killing Inactive Threads

1 On the **Datasource Explorer**, select the **Inactive Threads** node.

DBArtisan displays the inactive threads in the right pane of the Explorer window.

2 In the right pane of the **Explorer** window, click the target inactive thread(s).

3 On the **Datasource Explorer** toolbar, click the **Kill** button.

OR

In the right pane of the **Explorer** window, right-click the inactive thread(s), and then click **Kill**.

DBArtisan kills the inactive thread.

For more information, see Datasource Explorer nodes by DBMS platform.

Traces

NOTE: Traces are available for IBM DB2 for z/OS and OS/390 only.

The Traces node displays a list of active traces. Traces let you monitor and collect DB2 for OS/390 monitoring, auditing, performance, accounting, statistics, and serviceability (global) data.

NOTE: This information is read-only.

For more information, see Datasource Explorer nodes by DBMS platform.

Buffer Pools

The Buffer Pools node displays a list of buffer pools, the main storage that is reserved to satisfy the buffering requirements for one or more table spaces or indexes.

For more information, seeDatasource Explorer nodes by DBMS platform.

Utilities Node

NOTE: The Utilities node is available for IBM DB2 for OS390 only.

>

The Utilities node displays wizards that are typically run from the command line and do not apply to a particular type of schema object. Some of these wizards are also available via the command menus for objects. The wizards include:

- Check Data
- Check Index
- Copy
- Copy to Copy
- Load
- Merge Copy
- Modify Recovery
- Modify Statistics
- Quiesce Tablespace
- Rebuild Index
- Recover
- Reorg Index
- Reorg Tablespace
- Runstats Index
- Runstats Tablespace
- Unload
- Maintain Template

TIP: To run a wizard, select the command in the right pane of the application and click the Execute button on the Datasource Explorer toolbar.

For more information, see Datasource Explorer nodes by DBMS platform.

Schema Node

NOTE: Schema node is available for IBM DB2 for Linux, Unix, and Windows, IBM DB2 for OS390, and Oracle.

DBArtisan offers an environment for managing database objects across different database platforms. DBArtisan stores information about server and object types in object editors. Each editor is a tabbed dialog box that groups related information together. Each editor includes a context-sensitive Commands menu with pertinent functionality for the desired object. Many of the object editors contain a DDL tab.

The Schema node displays the underlying target object's SQL. The object editor toolbar has a refresh button, which lets you refresh object editor contents, clear the object editors, and log SQL if SQL Logging is turned on. This can be useful when you have completed an object and want to see the changes immediately. For objects that have dependent objects, such as tables, triggers, procedures, and views, you can view and access their dependent objects directly from the object tab, instead of having to go back to the correct node of the Datasource Explorer.

For more information, see Datasource Explorer nodes by DBMS platform.

>

Security Node

NOTE: The Security node is available for all DBMSs.

The Security node gives you quick access to related security functions and lets you administer and control security on your database from a single vantage point. Rather than having to sift through a large number of objects, you can simply create and alter security settings to be applied across the server from this top-level.

Each DBMS offers different security options:

IBM DB2 for Windows, Unix, and Linux	IBM DB2 for OS390 and Z/OS	Microsoft SQL Server	MySQL	Oracle	Sybase
Groups					
		Logins			Logins
		Linked Servers			
				Profiles	
		Roles		Roles	Roles
Users	Users		Users	Users	

Application Processes

NOTE: The Application Processes node is available for IBM DB2 for z/OS and OS/390.

DBArtisan places Application Processes as the first level of information under the Datasources node in the Datasource Explorer. The Application Processes node of the Datasource Explorer includes portions of the mainframe DB2 subsystem that relate to application programming. The Application Processes node also includes two components that are industry innovations: a search tool for statements, and a means to expose previously established access paths.

Application Processes includes:

- Plan
- Packages
- DBRM Node
- Procedures
- Triggers

Plan

A Plan is an executable application created in the bind process. It can include one or more packages or debris.

>

The Datasource Explorer List

You can select the Plan node in the Datasource Explorer and view the corresponding statistics. The table below describes the information DBArtisan derives and displays about plans from the system catalog:

Column Name	Description
Plan Name	Name of plan
Owner/Creator	User name of plan owner or creator
Qualifier	Name of qualifier
Valid	Indicates if plan is valid or invalid
Operative	Indicates if plan is operative
Isolation	Level of isolation
Reoptvar	
Packlists	Number of packlists
Environments	Number of environments
Bind Date/Timestamp	Time stamp of bind

TIP: Double-click a specific plan open the Plan editor. For details, see Plans Editor (IBM DB2 Z/OS).

For more information, see Application Processes.

DBRM Node

NOTE: DBRM is for IBM DB2 for z/OS and OS/390 only.

A 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.

The Datasource Explorer List

When you select the DBRM node of the Datasource Explorer, DBArtisan lists pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays DBRMs from the system catalog:

Column Name	Description
DBRM Name	Displays the name of the DBRM.
Version	Displays DBRM version.
Contoken	The contoken (consistency token) in the object compares to the contoken in the plan/package. If the contokens match, then the program will execute. If they do not match and autorebind = y, and the rebind is successful, then the program will execute. If they do not match and autorebind = n or fails, then an -818 is returned to the program - contoken mismatch, and the job ends. DB2 displays the mismatched contokens. The analysis associated with a -818 requires finding the DBRM or package contoken. This pane lets you easily expose the contokens and find the version mismatch.
Language	Displays the Host language.
Precompile Timestamp	Precompilation time of the DBRM.

For more information, see Application Processes.

Server Node

NOTE: The Server node is available for MySQL, Microsoft SQL Server, and Sybase ASE.

DBArtisan places Server as the first level of information under the Datasource node in the Datasource Explorer. Server includes:

MySQL

- Version Information
- Server Configuration
- Scheduling

Microsoft SQL Server

- Version Information
- Database Configuration
- Microsoft SQL Server Agent
- Error Logs

Sybase ASE

- Version Information
- Database Configuration

Accessing a Server

Connect to a datasource to access a server. After you have connected to your datasource, do the following:

1. On the **Datasource Explorer**, expand the datasource node.
2. On the **Datasource Explorer**, expand the **Server** node.

Server Version Information

DBArtisan places Version under the Server node of the Datasource Explorer. When you select Version, DBArtisan shows version information about your datasource in the right pane of the Explorer window.

Displaying Version Information

Connect to a datasource to view version information. After you have connected to a datasource, do the following:

1. On the **Datasource Explorer**, expand the **Datasource** node.
2. On the **Datasource Explorer**, expand the **Server** node.
3. Click **Version**.

   DBArtisan displays version information in the right pane of the Explorer window.

For more information, see Server Node.

Configuration

DBArtisan places Configuration under the Server node of the Datasource Explorer. When you select Configuration, DBArtisan shows configuration information about your datasource in the right pane of the Explorer window. You can use the Datasource Explorer or shortcut menu to:

- Open the Server Configuration dialog box.
- Open the Edit Configuration dialog box.
- Extract SQL.
- Execute an HTML report.

Opening the Server Configuration Dialog Box

Connect to a datasource to access a server. After you have connected to your datasource, do the following:

1. On the **Datasource Explorer**, expand the datasource node.
2. On the **Datasource Explorer**, expand the **Server** node.
3. Click **Configuration**.

   DBArtisan displays configuration information in the right pane of the Explorer window.
4. Select the target configuration parameter.
5. On the **Datasource Explorer** toolbar click **Open**.

   OR

   In the right pane of the **Explorer** window, right-click the configuration parameter, and then click **Open**.

   DBArtisan opens the Server Configuration dialog box.

For more information, see Server Node.

Microsoft SQL Server Agent

DBArtisan places SQL Server Agent under the Server node of the Datasource Explorer. DBArtisan assists you in using the Microsoft SQL Server Agent service to automate administration of a Microsoft SQL server. Microsoft SQL Server Agent manages, and schedules, allowing programmed responses to server events or the execution of scheduled administrative tasks. Automating recurring administrative tasks and responses to server events can save considerable time for other administrative tasks.

DBArtisan lets you define the three main components of automatic administration managed by the SQL Server Agent service. These components are:

- Alerts
- Operators
- Jobs

To use the Microsoft SQL Serve Agent features in DBArtisan, you need to install Microsoft SQL Enterprise Manager, Query Analyzer, or Profiler tool.

NOTE: The SQL Server Agent must be installed if you have purchased and plan to use the Backup Analyst.

NOTE: Be aware that when you work with the Alerts, Operators, or Jobs features and click Help, you connect directly to SQL Server Books online. Embarcadero Technologies is not responsible for the contents of the SQL Server Books materials.

Alerts

An alert is a definition that matches one or more Microsoft SQL Server events with a response. Errors and messages, or events, are generated by Microsoft SQL Server and entered into the Microsoft Windows application log. SQL Server Agent reads the application log and compares events to alerts that you defined. When SQL Server Agent finds a match, it fires an alert. DBArtisan lets you define alerts to respond to SQL Server events by notifying operators, forwarding the event to another server, or executing a job. These administrative tasks runs automatically if the SQL Server Agent service has been started.

When you select the Alerts node on the Datasource Explorer, DBArtisan displays all alerts for the server in the right pane of the Explorer window. The column headings in the right pane of the Explorer window provide important information about alerts. These headings are:

- Name
- Enabled
- Error
- Severity
- Last Occurred
- E-mail
- Panel
- Net Send
- Count

You can use the shortcut menu to access functionality associated with alerts. Shortcut menus let you:

- View Properties
- Create a New Alert
- Drop an Alert
- Find Object
- Select All
- Unselect All

NOTE: On the shortcut menu, when you click New, DBArtisan opens the Microsoft SQL Server Agent Alert Properties dialog box where you can enter the new alert information.

You can also open the *Microsoft SQL Server Agent* Alert Properties dialog box by double-clicking a particular alert, or by clicking the Properties button on the Datasource Explorer.

Operators

Operators are individuals responsible for the maintenance of one or more computers running Microsoft SQL Server. E-mail, panel, and Net Send can notify these operators.

When you select the Operators node on the Datasource Explorer, DBArtisan displays all operators defined for the server in the right pane. The column headings in the right pane of the Explorer window provide important information about operators. These headings are:

- Name
- Enabled
- Last Email

- Last Panel
- Last Net Send

You can use the shortcut menu to access functionality associated with operators. The shortcut functionality lets you:

- View Properties
- Create a New Operator
- Drop an Operator
- Find Object
- Select All
- Unselect All

NOTE: On the shortcut menu, when you click New, DBArtisan opens the Microsoft SQL Server Agent Operator Properties dialog box where you can enter the new operator information.

You can also open the *Microsoft SQL Server Agent* Operator Properties dialog box by double clicking a particular operator, or by clicking the Properties button on the Datasource Explorer.

Jobs

Jobs define administrative tasks that can execute on a local server according to schedules or in response to alerts. A job is comprised of job steps that can be:

- Executable Programs
- Windows Commands
- Transact-SQL Statements
- Microsoft ActiveX Script
- Replication Agents

When you select the Jobs node on the Datasource Explorer, DBArtisan displays all jobs for the target server in the right pane of the Explorer window. The column headings in the right pane of the Explorer window provide important information about jobs. These headings are:

- Name
- Category
- Scheduled
- Status
- Next Run Status (Start Date)
- Next Run Date

You can use the shortcut menu to access functionality associated with jobs. The shortcut menu functionality lets you:

- View Properties
- Start Job
- Stop Job
- Enable Job
- Disable Job

- View Job History
- Create a New Job
- Drop a Job
- Find Object
- Select All
- Unselect All

On the Shortcut menu, when you click New, DBArtisan opens the Microsoft SQL Server Agent Job Properties dialog box where you can enter the new job information.

You can also open the *Microsoft SQL Server Agent* Job Properties dialog box by double clicking on a particular job, or by clicking the Properties button on the Datasource Explorer.

Error Logs

DBArtisan places Error Logs under the Server node of the Datasource Explorer. Error logs are comprised of information recorded by Microsoft SQL Server. When you expand the Error Logs node, you can see the current and archived error logs. When you select the Error Logs node, DBArtisan displays the current and archived error logs in the right pane of the Explorer window. To note information, DBArtisan uses a blue information icon, and to note errors, DBArtisan uses a red error icon.

The column headings in the right pane of the Explorer window provide important information about error logs. These headings are:

- Name
- Source
- Message

You can use the Datasource Explorer toolbar or the shortcut menu to access functionality associated with error logs. The shortcut menu functionality lets you:

- Find
- Toggle All Messages
- Generate an HTML Report
- Refresh the Window
- Find Object
- Select All
- Unselect All

TIP: Double-click any error log to display the message.

For more information, see Server Node.

Storage Node

NOTE: The Storage node is available for all DBMSs.

>

DBArtisan places all storage mechanisms under the Storage node of the Datasource Explorer. Storage objects include:

IBM DB2 for Windows, Unix, and Linux	IBM DB2 for OS390 and Z/OS	Microsoft SQL Server	Oracle	Sybase
Bufferpools	Bufferpools			
			Control Files	
			Directories	
				Data Caches
				Database Devices
		Dump and Backup Devices		Dump and Backup Devices
Nodegroups				
			Redo Log Groups	
			Rollback Segments	
	Stogroups			
Tablespaces	Tablespaces		Tablespaces	

Subsystem

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

DDF

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

NOTE: The information is read-only.

For more information, see Subsystem.

Parameters

Parameters displays the DB2 subsystem parameters.

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

For more information, see Subsystem.

Filtering in the Datasource Explorer

For ease of navigation, DBArtisan lets you filter the Datasource Explorer tree to display only selected nodes and objects. DBArtisan 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 datasources, 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.

Simple, on-the-fly, character-based tree filtering

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

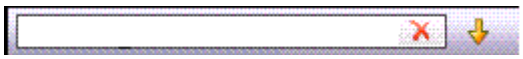

To restrict the Datasource Explorer tree display to only user-defined nodes whose name contains a particular character string:

- 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 Datasource Explorer tree display for a datasource by name, owning schema, or name/schema combinations. These filters are defined at the DBMS platform level and enabled or disabled on a datasource by datasource level. Multiple object filters can be created for a DBMS platform, letting you enable combinations of filters for a single datasource, 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 Datasource Explorer tree for a datasource to display only objects belonging to the Schema **QA** and whose name starts with the string **Test**.

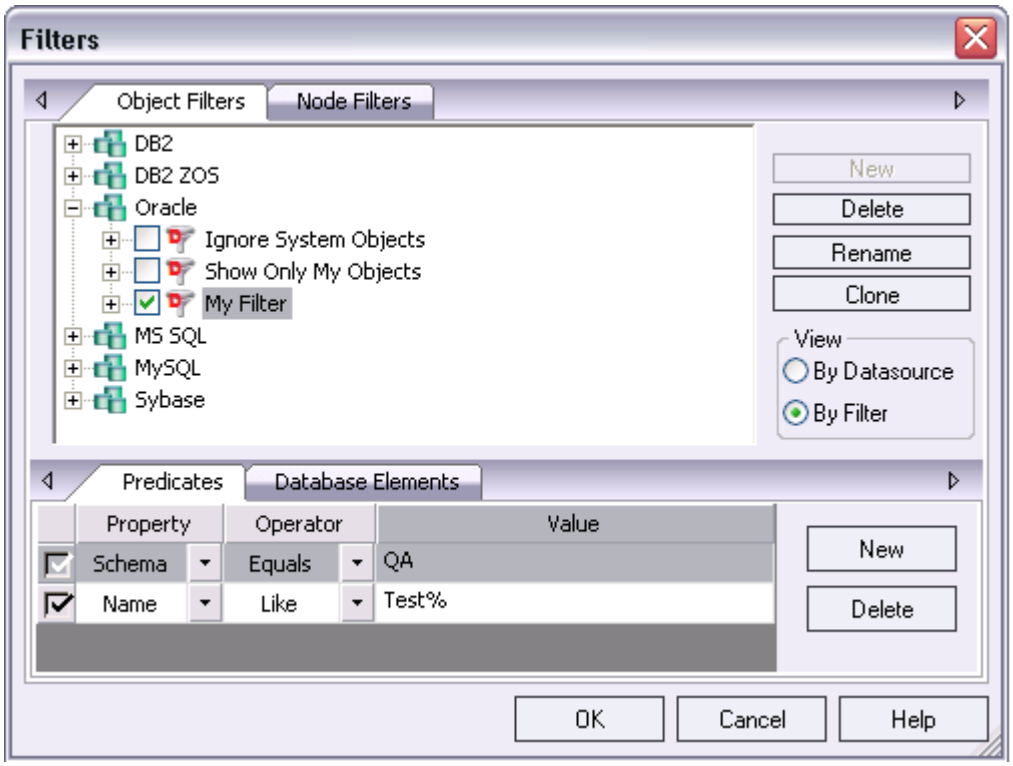

Each filter condition is a simple expression that tests object names for equality (**Equals**, **Not Equals**), inclusion in a specified list (**In, Not In**), or using pattern matching (**Like, Not Like)**. Optionally, an object filter can also be used to restrict display of specified object types on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

To create an object filter that can be enabled for all datasources 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. The **Object Filters** tab opens by default, letting you create an object filter definition. The **Node Filters** tab, on the other hand, lets you perform real-time filtering of object types for all datasources ona DBMS platform. For details, see Node, or object type, filtering.

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 datasource, select the **By Datasource** view, select a specific datasource, 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 **Value**s. DBArtisan supports the wildcard characters and conventions supported on the target DBMS.
- 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 datasource of that DBMS type.

To disable or enable an object filter for a datasource:

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 Datasource**.

3 Expand the icon corresponding to the target DBMS and then expand the icon corresponding to the specific datasource 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 datasource in the Datasource 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, groups, profiles or logins, you may wish to filter those out of the Datasource explorer tree. You could either disable display of the individual object types or if practical, disable display of the **Security** node for the DBMS.

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 Datasource Explorer tree and select **Filter** from the context menu. The **Filters** dialog opens.

The **Filters** dialog opens. The **NodeFilters** tab opens by default, letting perform real-time filtering of object types for all datasources ona DBMS platform. The **Object Filters** tab, on the other hand, lets you create permanent, name-based or schema-based filters. For details, see Object name/schema filtering.

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

Opening and closing Explorer tabs

DBArtisan provides flexibility in the configuration of its display. DBArtisan lets you open or close multiple Explorer Tab windows. This feature lets you configure the DBArtisan desktop to suit the task at hand. The default setting for the Datasource Explorer is a single tab docked on the left-hand side of the application.

To open a new Datasource Explorer tab:

1 On the **Datasource** menu, click **New Explorer**.

To close a Datasource Explorer tab:

1 Right-click the tab label at the bottom of the target **Datasource Explorer** tab and select **Close.**

Enterprise Management Console

The Enterprise Management Console interface is an enterprise-wide view of managed databases. It is available when the application is focused on the top-level node of a datasource group. You can use this feature to do multiple administrative operations with one or two clicks.

The Enterprise Management Console gives you quick access to the following functionality:

- Connecting to your datasources
- Modifying datasource registration
- Database Monitor
- Data Load
- Data Unload
- Database Schema Extraction
- Database Schema Migration

Parallel Query Execution/Multithreading

DBArtisan takes full advantage of 32-bit architecture by supporting parallel query execution or multithreading.

DBArtisan accelerates the development process by executing multiple queries simultaneously using separate operating system threads and improving response time because most operations can start immediately and return control to the primary application thread.

DBArtisan lets you run multiple, concurrent queries against the same datasource.

DBArtisan employs a database connection pooling technique that dynamically spawns connections as needed and disconnects them when the associated query has completed.

DBArtisan increases throughput by automatically assigns query threads to separate processors if your system has multiple CPUs. When you execute parallel queries against different database servers using a multi-processor computer, you experience near linear gains in throughput.

Describe Window

The Describe Window is a floating window that lets you view procedures, tables, views, and functions (Oracle and IBM DB2 for Linux, Unix, and Windows 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

DBArtisan offers three ways to open the Describe window:

1. In an editor, right-click an object and then click **Describe**.
2. On the **View** menu, select an object and then click **Describe**.
3. On the right pane of the **Explorer** 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.

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).

DBArtisan Toolbars and Menus

DBArtisan offers multiple ways to navigate through the application:

- Main Menu
- Toolbars
- Shortcut Menus

For more information on customizing these items, see Customizing general user interface appearance.

Main Menu

You can access DBArtisan's features from the Main Menu by clicking the menu name, and then selecting the target command on the sub-menu. The menus are context-sensitive, changing as you perform tasks. DBArtisan menu items include:

Menu Item	Description
File	Create, open, close, print and save script files and result sets. Set application options and defaults. **Note**: When you open New, you have easy access to a new ISQL window, the DDL Editor, a Java file, or the Active Script.
Datasource	Create, modify, select, manage, connect to, and disconnect from datasources. Access database object editors. Display or hide the Datasource Explorer. Access session recording for Tivoli integration.
Utilities	Access utilities to extract schema, migrate schema, export and import data, execute scripts, search databases, monitor process activity and locks, schedule tasks, and restore, manage, and backup databases.
Logfile	Activate/deactivate, open, flush, and annotate the DBArtisan application log.
View	Configure DBArtisan to hide or show toolbars, status bars, and windows. Access full screen mode.

Menu Item	Description
Tools	Access customizable tools menu. Access installed DBMS connectivity utilities and server documentation.
Analyst	Opens the Space Analyst, Capacity Analyst, Backup Analyst, and Performance Analyst.
Bookmarks	Access and manage bookmarks.
Window	Cascade and tile open windows. Arrange icons. Toggle between open windows.
Help	Lets you access context-sensitive Help, link to Embarcadero on the Web, and see license information when you click About...
Edit	Edit and manipulate the text in your SQL scripts. Available only when an ISQL Window is selected.
Query	Execute and set query options for a SQL script. Available only when an ISQL Window is selected.
Format	Format the contents of result sets. Available only when a Result Window is selected.

For more information, see DBArtisan Toolbars and Menus.

Explorer Bookmarks

DBArtisan supports easy datasource navigation with the Explorer bookmarks. You can right-click any node of the Explorer tree and add a bookmark. After a bookmark is created, you can simply select it from the Bookmark menu, and be instantly positioned and connected to the resource in the Datasource Explorer tree.

Creating Explorer Bookmarks

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

   DBArtisan opens the Add Friendly Bookmark Name dialog box.

2. Type the Explorer bookmark name.
3. Click **OK**.

   DBArtisan 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**.

   DBArtisan opens Bookmark Manager.

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

   DBArtisan 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 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.

Toolbars

DBArtisan supports two classes of toolbars:

- Feature-specific toolbars appearing at the top of the window or panel of tools such as the Datasource Explorer, Object editors, the ISQL editor, and the Analyst tools:

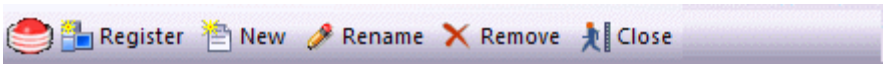

  For information on the functions provided by feature-specific toolbars, see the documentation for the associated feature.

- Optional-display toolbars, appearing in one or more rows below the menu bar.

You can hide or display the individual optional-display toolbars and they change to reflect the element of the application you are using. These toolbars contain buttons that are the fastest way to access commonly used features of DBArtisan.

The following topics provide information you need to work with optional-display toolbars:

- Using toolbar viewing options
- Understanding the short cuts available on specific toolbars

  NOTE: You can set preferences as to the default toolbars displayed and their appearance. For details, see Customizing general user interface appearance.

Using toolbar viewing options

DBArtisan 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 and the commands available on each, see Understanding the short cuts available on specific toolbars.

Understanding the short cuts available on specific toolbars

The following optional-display toolbars are available using the **Toolbars** command on the **View** menu.

Analyst Series Toolbar

This toolbar offer shortcuts tthat open the Space Analyst, Capacity Analyst, Performance Analyst, and Backup Analyst tools.

Datasource Toolbar

This toolbar offers shortcuts that let you connect to or disconnect from a datasource, select a datasource and for Sybase and SQL Server, select a specific database.

Main Toolbar

This toolbar has common application functions.

It offers the following shortcuts:

New	Open	Save	Save All
Print	Send Mail	Options	New Datasource Explorer
Output	Describe Window	Full Screen	Log SQL
Session Recording	Find Object		

Registration Toolbar

This toolbar offers shortcuts to registering a datasource, editing a datasource registration, and unregistering a datasource.

SQL Edit Toolbar

This toolbar offers functions used in editing SQL:

It offers the following shortcuts:

Undo	Redo	Cut	Copy
Paste	Insert File	Paste SQL Syntax	Paste SQL Statement
Comment	Undo Comment	Upper Case	Lower Case
Match Token	Find	Repeat	Replace
Toggle Bookmark	Previous Bookmark	Next Bookmark	Clear Bookmark
Increase Indent	Decrease Indent		

SQL Profiler Toolbar

This toolbar offer functions used when working with the Oracle Profiler.

It offers the following shortcuts:

Start Profiler	Flush Profile	Run Summary	Run Detail
Unit Summary	Unit Detail	Clear Profile Tables	Stop Profiling

Tools Toolbar

This toolbar offers functions otherwise available on the **Tools** menu.

It offers the following shortcuts:

Customize Tools Menu	Find in Files	Database Search	Script Execution Facility
File Execution Facility	Scheduler	Visual Difference	Edit Data
Build Query	Table/Index Size Estimator	SQL* PLus	Performance Center
Code Workbench	Import Data From File		

Utilities Toolbar

This toolbar offers functions otherwise available on the **Utilities** menu.

It offers the following shortcuts:

Schema Extraction	Schema Migration	Shema Publication	Data Unload
Data Load	Database Monitor	Server Configuration	

Windows Toolbar

This toolbar offers functionality for working with windows.

It offers the following shortcuts:

Cascade	Tile Windows Horizontally	Tile Windows Vertically	Next
Previous	Windows		

Customizing general user interface appearance

DBArtisan 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 DBArtisan:

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	
Toolbars	Select the toolbars you want to display in the application. For information on the toolbars available, see Understanding the short cuts available on specific toolbars.	
Application Visual Style	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: **Use default WIndows XP colors**, **OneNote style tabs**, **Docking Tab Colors**, **Allow MDI Tab Swapping**, **Enable Smart Docking**, **Enable Tab Menu**, and **3D Rounded Docking Tabs**.	
	Menu animations	Lets you specify a menu animation style of UNFOLD, SLIDE, or FADE.
	Menu Shadows	Displays shadowed 3D effects.
Tools	Lets you define external applications to run from the **Tools** menu of DBArtisan: the text displayed on the Tools menu command (**Menu contents**), the path and file name of the executable (**Command**), optional **Arguments**, and an optional **Initial DIrectory**.	
Keyboard	**Category**	Select a general category for a hot key for the command.
	Commands	Select a hot key command, based on the general category.
	Description	Displays the command description.
	Set Accelerator for	Select application area where you want new hot key to be active.
	Current Keys	Displays current hot key.
	Press New Shortcut Key	Press keyboard key or an unassigned F key.
Options	**Show ScreenTips on toolbars**	Select to display a ScreenTip when you hover your mouse over a button. For example, when you hover your mouse over the New button, DBArtisan displays the ScreenTip “New.”
	Show shortcut keys in ScreenTips	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, DBArtisan displays the ScreenTip “New (CTRL+N).”

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

Shortcut Menus

DBArtisan incorporates shortcut 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.

Accessing Shortcut Menus

1 On **Desktop**, right-click **DBArtisan**.

DBArtisan opens the shortcut menu.

2 Click the target shortcut menu command.

For more information, see DBArtisan Toolbars and Menus.

Keyboard Commands

DBArtisan provides a number of Keyboard commands that expedite IISQL Editor development. The table below describes the available keyboard commands:

Category	Action	Keyboard shortcut
General editing	Delete one character to the left	BACKSPACE
	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
Extending selections	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
Moving the insertion point	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
	Up one line	UP ARROW
	Down one line	DOWN ARROW
	To the end of a line	END
	To the beginning of a line	HOME
	Up one screen (scrolling)	PAGE UP
	Down one screen (scrolling)	PAGE DOWN
	To the end of a document	CTRL+END
	To the beginning of a document	CTRL+HOME
Bookmarks	Toggle bookmark on/off	CTRL+F2
	Go to next bookmark	F2
	Go to previous bookmark	SHIFT+F2
Splitter windows	Go to next pane	F6

Category	Action	Keyboard shortcut
	Go to previous pane	SHIFT+F6
Debugger operations	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	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
SQL Window	Go to the Result Tab	CTRL+ALT+R
	Go to the Query Tab	CTRL+ALT+Q

For more information, see DBArtisan Toolbars and Menus

Customizing the Tools Menu

The Customize Tools Menu dialog box lets you add up to sixteen commands to the Tools menu. You can associate any program that running on your local machine. You can use this feature to associate frequently used programs and commands, thereby adding functionality directly to the DBArtisan application. In addition, you can specify arguments for any command that you add to the Tools menu. You can use these arguments to decipher how you want to execute commands. DBArtisan includes many pre-defined arguments that let you customize the application.

During the initial installation, DBArtisan automatically adds the following items (if installed locally) to the Tools menu.

- ODBC Administrator
- Oracle Service Name Manager
- Sybase Open Client/Server Config
- Sybase Directory Service Edit (DSEDIT)
- SQL Server Network Library Config
- SQL Server Books Online
- DB2 Client Configuration Assistant

For more information, see DBArtisan Toolbars and Menus.

Customize User Tools Menu Dialog Box: DBArtisan lets you customize the Tools menu.

The table below describes the options and functionality on the Customize User Tools Menu dialog box:

Option	Description
Menu Contents	In the box, click the target command.
Menu Text	In the box, type the name of the tool as you want to appear on the Tools menu.
Command	In the box, type the path and name of the program for the new tool or click Browse to locate the program. For example, C:\WINDOWS\NOTEPAD.EXE
Arguments	To pass an argument to the command, in the box, type any argument or click the Arrow, and then click the target argument. **NOTE:** To display a menu of pre-defined arguments, click the drop-down arrow next to the Argument box, and then select an argument from the list to insert the syntax into the box.
Initial Directory	To specify an initial directory, in the box, type the path and name of the file directory or click Browse to locate the file directory.
Add Button	Click to add a new command to Tools menu.
Remove Button	In the Menu Contents box, click the command you want to remove, and then click the button.
Move Up Button	Click to move the command up one position in the menu.
Move Down Button	Click to move the command down one position in the menu.

Completing the Customize User Tools Menu Dialog Box

To complete the Customize User Tools Menu dialog box, do the following:

1. On the **Tools** Menu, select **Customize**.

   DBArtisan opens the Customize User Tools dialog box.

2. Select options.
3. Click **OK**.

For more information, see DBArtisan Toolbars and Menus.

>

Specifying Arguments

DBArtisan provides a number of pre-defined arguments that you can pass to programs that you have added to the Tools menu. You can specify an argument to be passed to a program for newly added commands by choosing one of DBArtisan's pre-defined arguments or typing in your own arguments. The table below describes the pre-defined 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 nonsource 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 nonsource window is active.
$$CurLine$$	The current pointer line position within the active window.
$$CurCol$$	The current pointer column position within the active window.
$$CurText$$	The current text (the word under the current pointer 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 DBArtisan.
$$CurUserID$$	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.

The table below describes examples for using the pre-defined arguments and command-line arguments:

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

TIP: DBArtisan lets you specify an argument in the DBArtisan opens the Customizing general user interface appearance.

For more information, see DBArtisan Toolbars and Menus.

DBArtisan Dialog Boxes

DBArtisan offers the following common dialog boxes. Additional dialog boxes are discussed in the Using section.

- Preview
- Report

DBArtisan Wizards

DBArtisan offers several wizards that let you create objects without knowing the underlying SQL commands. DBArtisan also offers wizards to perform common tasks.

DBArtisan Editors

DBArtisan 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.

Customizing the DBArtisan Display

DBArtisan lets you customize your display to suit your work habits.

For more information, see:

Full Screen Mode

Active Windows Dialog Box

Navigating DBArtisan Windows

Output Window

Set Sort Columns Dialog Box

Full Screen Mode

DBArtisan is equipped with Full Screen mode capabilities that let you conceal the application framework and take advantage of the entire monitor area. Full screen mode hides any other applications running on the computer and uses every available pixel for DBArtisan. Main menu functionality is accessible through keyboard commands when you use Full Screen Mode.

Activating Full Screen Mode

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

DBArtisan expands the application to fit the entire monitor area, and displays the Full Screen Mode icon as a stand-alone floating toolbar.

Dismissing Full Screen Mode

1 Click **Full Screen**.

DBArtisan reduces the size of the application.

TIP: If you closed the Full Screen Mode toolbar, right-click the top of the DBArtisan desktop to display the full screen.

For more information, see Customizing the DBArtisan Display.

Active Windows Dialog Box

DBArtisan's multi-tasking environment lets you open multiple active windows. DBArtisan's Active Windows dialog box lets you see all the open windows in the application workspace. DBArtisan displays active windows as tabs in the Datasource Explorer. Each tab represents a separate window that lets you execute, create or view the objects or schema in your database.

TIP: DBArtisan lets you use a mouse with a wheel to navigate in active windows.

Opening the Active Windows Dialog Box

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

DBArtisan opens the Active Windows dialog box and any open windows in the application workspace.

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

Option	Description
Activate	Sets the focus on the selected window and closes the Active Windows dialog box.
OK	Closes the Active Windows dialog box and makes any changes you have made to the windows in the workspace.
Save	Saves the contents of the selected window. If you are saving a file, DBArtisan prompts you to provide a name and location for the file.
Close Window	Closes the selected window. If you have not saved the contents of the window, DBArtisan prompts you to do so.
Help	Initiates and displays this Help topic in the DBArtisan Help.

For more information, see Customizing the DBArtisan Display.

Navigating DBArtisan Windows

DBArtisan offers functionality through the Windows toolbar that lets you:

- Tile windows vertically or horizontally.
- Cascade open windows.
- Minimize all windows and move to the next active window.

TIP: DBArtisan lets you use a mouse with a wheel to navigate in active windows.

>

The table below describes the commands available on the Windows menu:

Command	Description
Window	Activates the Active Windows dialog box.
Cascade	Layers all active windows on the desktop so that you can view all active windows simultaneously.
Tile Horizontal	Tiles all active windows on the desktop horizontally.
Tile Vertical	Tiles all active windows on the desktop vertically.
Next	Moves the focus of the application to the next active window.
Previous	Moves the focus of the application to the previous active window.

For more information, see Customizing the DBArtisan Display.

Output Window

DBArtisan has a separate window that captures all messages returned by the server and to log entries about the progress of operations started by the application.

The Output Window includes the following tabs:

- General
- Find In Files 1
- Find In Files 2

The Output Window lets you do the following to server messages:

- View
- Save
- Print
- Copy
- Flush

For more information, see Customizing the DBArtisan Display.

CONFIGURING THE OUTPUT WINDOW

Basic Output window options you can configure include:

- Floating the Output Window (with Smart Docking)
- Docking the Output Window
- Undocking the Output Window
- Hiding the Output Window
- Displaying the Output Window

Floating the Output Window (with Smart Docking)

This is the default view when you first open DBArtisan.

1. Place your pointer on the **Output Window** and drag it to another location.

   Arrows appear and you can move the window to the top, bottom, or sides of the view within DBArtisan or to another location on your desktop.

2. To change the location of the **Output Window**, repeat the first step.
3. To disable the feature, right-click and choose **Docking** or **Hide**.

Docking the Output Window

1. Place your pointer on the **Output Window**.
2. Right-click, and then click **Docking View**.

   DBArtisan docks the Output Window to the bottom of the application frame.

Undocking the Output Window

1. Place your pointer on the **Output Window**.
2. Right-click, and then click **Docking View**.

   DBArtisan **floats** the Output Window.

Hiding the Output Window

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

   DBArtisan hides the Output Window.

Displaying the Output Window

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

   DBArtisan displays the Output Window.

For more information, see Customizing the DBArtisan Display.

SAVING OUTPUT WINDOW MESSAGES

DBArtisan lets you save messages in the Output window to a file.

1. Right-click a message in the **Output Window**, and then click **Save**.

   DBArtisan opens the Save As dialog box.

2. In the **File Name** box, type the location and name of the file.

   DBArtisan defaults the file extension to.msg.

3. Click **Save**.

   DBArtisan saves the file.

For more information, see Customizing the DBArtisan Display.

PRINTING OUTPUT WINDOW MESSAGES

DBArtisan lets you print the messages you see in the Output window.

1 Right-click a message in the **Output Window**, and then click **Print**.

DBArtisan opens the Print dialog box.

2 Configure your print job.

3 Click **Print**.

DBArtisan prints the file.

For more information, seeCustomizing the DBArtisan Display.

COPYING OUTPUT WINDOW MESSAGES

DBArtisan lets you copy all or part of the contents of the Output window to the clipboard.

1 Select the target **Server Messages.**

2 Right-click, and then click **Copy**.

DBArtisan copies the selected messages to the Windows clipboard.

3 Paste the text into other Windows applications.

For more information, seeCustomizing the DBArtisan Display.

CLEARING OUTPUT WINDOW MESSAGES

DBArtisan lets you clear server messages from the Output Window.

1 Right-click a message in the **Output Window**, and then click **Clear**.

DBArtisan clears the server messages.

For more information, see Customizing the DBArtisan Display.

Set Sort Columns Dialog Box

The Set Sort Columns dialog box lets you sort multiple columns, and sort column identification.

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**.

DBArtisan 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.

Specifying DBArtisan application and feature options

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

To specify options for the DBArtisan application or a particular feature:

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

DBArtisan opens the Options Editor.

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

Connection options	Specifies the timeout parameters, packet size for a connection, and ANSI to OEM settings.
Data Editor options	Specifies settings for Data Editor.
Data transfer options	Sets default directories when performing data unload and load and data export and import operations.
Explorer options	Sets defaults for the organization of objects in the Datasource Explorer.
Datasource options	Specifies how to store the datasource registry on the local machine.
DDL extract options	Specifies whether or not DBArtisan should include DROP statements when extracting the schema of different database object types.
Debug options	Sets the duration of your debug initialization and debug session, enable or disable DBMS_OUTPUT, and enable the refresh option.
Directories options	Sets the default directories for placing the output of different operations such as HTML reports or schema extractions.
General options	Sets defaults for automatic login, restoring the last session, and other general application options.
Grid properties (Results window) options	Dictates the physical appearance of the results window grid.
ISQL options	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.
Java options	Specifies load Java files and drop Java files.
JDBC options	Lets you specify Java Virtual Machine options in addition to the JDBC driver location for each platform.
Logging options	Sets defaults for SQL Logging.
Migration options	Sets the default options for schema migration operations.
MySQL Utilities Options	Specifies paths to MySQL utilities.
Oracle Utilities options	Specifies the location of the Oracle Utilities.
Perf Center options	Specifies Performance Center's integration with DBArtisan.
Query Builder options	Specifies global settings for Query Builder.
Results (ISQL) options	Specifies auto format result sets, sets the display and format of Result Windows, and the mail file type and default fonts.
SMTP mail options	Lets you specify defaults for your outgoing mail notifications.
Space Analyst options	Sets Space Analyst options.
Warnings options	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, DBArtisan 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.

Connection options

After opening the Options editor (see Specifying DBArtisan 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
Login Timeout	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
Query Timeout	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
Packet Size	Specifies the network packet size to be used when communicating with the server, in bytes.	512
Max Connections (ctlib)	The maximum number of connections allowed by the ctlib setting on the client.	0
Client Character Set	Character set of client computer.	Local character set
Host Name	Name of the client computer.	Local name
Use Quoted Identifiers	If you plan to use delimited identifiers, select this option.	Not selected

Data Editor options

After opening the Options editor (see Specifying DBArtisan 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
Default Execution Mode	**Live Mode** lets you execute changes one row at a time. **Batch Mode** lets you make unlimited changes prior to execution. Batch Mode offers three sub-options: **Ignore errors-continue processing**, **Prompt on Error**, and **Stop after error(s)** (number of errors allowed before stopping execution)		Live Mode
Data Editor File	**Mail File Type**	Sets the default mail output style as Results, Tab Delimited, Comma Separated, or HTML.	Results
	Include column titles when saving	Includes column titles when saving.	Not selected
Grid Font	Customizes font style, and size for the Data Editor and the Results Grid.		Available
Printer Font	Sets font style, and size for printing output.		Available
Auto Format (Best Fit)	Fits formatting to match your desktop.		Selected
Begin and End Transaction Statements	Adds a beginning and ending transaction on each statement.		Selected
Default Date/Time Format	Displays the current date/time format and lets you customize the date/time display.		Results
	Use Calendar Control as default	If selected, DBArtisan uses the Calendar Control window.	Not selected
	2 digit year system setting warning	If selected, DBArtisan sends a warning when you use a two digit year system setting.	Selected
Confirmation Dialog Options	Enabling **Show Delete Confirmation Dialog**, **Show Update Confirmation Dialog**, or **Show Update LOB Confirmation Dialog** lets you display a confirmation dialog box when you use a delete command, update a row, or update a LOB, repectively..		Selected

Data transfer options

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

You can configure DBArtisan 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
Data Unload	Specifies the name and location of the default directory in the Data Unload box.
Data Load	Specifies the name and location of the default directory in the Data Load box.
Oracle Export (for Oracle)	Specifies the name and location of the default directory in the Oracle Export box.
Oracle Import (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 DBArtisan 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
Recent Datasource List Contains	Lets you specify number of **Datasources** to display in the datasource list.	8
Default to Alias Usage When Defining New Datasources		Unselected
Re-sort the Datasource Combo based on Connection state	When this box is checked, list of datasources on the Datasource toolbar identifies the currently connected datasources above all others. When this box is unchecked, datasources are listed in strict alphabetical order.	Selected

DDL extract options

After opening the Options editor (see Specifying DBArtisan 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
Extract to multiple windows	Select to extract the schema for each object into separate DDL windows. **NOTE:** This option only works when you extract DDL for multiple objects.
Extract in dependency order	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.
Script Use Statement	This option is for Sybase. Optimizes extraction through generating Use statements in the generated script.
Grouping Extracted DDL	There are two option as to how the DDL is generated for an extract operation: **Group the Drop and Create DDL for each object together** - 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. **Group DDL statements by type - all Drop statements, then all Create statements** - 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 DBMS-specific DDL Extract options.

DBMS-specific DDL Extract options

The DB2, DB2 (OS/390), Oracle, SQL Server, and Sybase 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.

The table below shows the object types you can include a DROP statement with when performing an ad hoc DDL extraction:

Option	DB2 for OS/390 and z/OS	DB2 for Windows, Unix, and Linux	Oracle	SQL Server	Sybase
Aliases	Yes (default)	Yes (default)			
Constraints	Yes	Yes	Yes	Yes	Yes
Database	Yes				
Defaults				Yes (default)	Yes (default)
Functions	Yes (default)	Yes (default)		Yes	
Indexes	Yes (default)	Yes (default)	Yes (default)	Yes (default)	Yes (default)
Packages		Yes (default)			
Procedures	Yes (default)	Yes (default)		Yes (default)	Yes (default)
Rules				Yes (default)	Yes (default)
Sequences		Yes (default)	Yes (default)		
Structure Types		Yes (default)			
Synonyms	Yes (default)		Yes (default)		
Tables	Yes (default)	Yes	Yes	Yes	Yes
Triggers	Yes (default)	Yes (default)		Yes (default)	Yes (default)
Users			Yes	Yes	Yes
User Datatypes				Yes (default)	Yes (default)
Views	Yes (default)	Yes (default)		Yes (default)	Yes (default)

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 DBArtisan 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
General	**Dependency Tree Option**	Lets you specify pre-fetch levels.	Pre-Fetch All Dependencies
Profiler	**Profiler Time Unit**	Lets you select a unit of milliseconds, seconds or minutes.	Milliseconds
	Save Profiler Reports	Lets you save profiler reports and type or browse for the report path.	Not Selected
Oracle	**Initialization Timeout (seconds)**	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
	Debug Session Timeout (seconds)	Specifies the point at which the application terminates your debug session due to idle time.	7200
	Enable DBMS Output	Enables the Oracle built-in package, DBMS_OUTPUT, letting you send messages from stored procedures, packages, and triggers.	Selected
	Refresh Dependencies for each run	Refreshes the dependencies each time you run the debugger.	Not selected
	Compile with Debug Option	When debugging a procedure against an Oracle datasource, DBArtisan 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 DBArtisan, will result in procedures A2 and A3 being marked as invalidated. The **Compile with Debug Options** settings let you control compilation of dependent objects while debugging.	Compile dependent options
DB2	**Debug Session Timeout (seconds)**	Specifies the point at which the application terminates your debug session due to idle time.	300
	Compile with Debug Option before Debug Session	Lets you specify options.	Prompt Always

Directories options

After opening the Options editor (see Specifying DBArtisan application and feature options), you can make changes to the **Directory** tab. DBArtisan 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\DBArtisan\DefFiles**

Option	Description	Default
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\DBArtisan\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\DBArtisan\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\DBA870\HtmlTpl**
User SQL Scripts	Specifies the name and location of the default directory for SQL Scripts.	**C:\Documents and Settings*user name*\Application Data\Embarcadero\DBArtisan\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\DBArtisan\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 Registering or Editing a Datasource.	**C:\Documents and Settings\davidt\My Documents\Embarcadero\DBArtisan\Directories\IncludeFiles**

Explorer options

After opening the Options editor (see Specifying DBArtisan 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 **Organize by Object Owner** (groups objects, by object type, for each user**)** or **Organize by Object Type** (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 DBArtisan session. For details, see Basic viewing options in the Datasource Explorer.	Owner
Sort Type	Lets you select object sorting: **Alphanumeric** or **Lexicographic**.	Alphanumeric

Option	Description	Default
Show System Objects and **Show Only My Objects** (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: **Never** - the relevant objects are not displayed by default. **Always** - the relevant objects are always displayed by default. **Use datasource filter** - The relevant default datasource filter is enabled. For more information, see Filters Node. You can change the **Show System Objects** and **Show Only My Objects** settings manually during a DBArtisan session. For details, see Using the default filters.	N/A
Options	Lets you anble or dispble the following settings: **Refresh after Object Commands** - Refreshes the Explorer automatically after an object has been modified or created. **Retain View Setting on Close** - Select to retain the current state of the Datasource Explorer so that it opens the same way the next time you start DBArtisan.	N/A
Landing Page View	Enables landing page views on a DBMS-by-DBMS basis. Landing pages display informative metadata that gives you a quick overview of your server/database.	Selected

General options

After opening the Options editor (see Specifying DBArtisan 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
Main tab	**Confirm on Exit**	Instructs DBArtisan to issue a message confirming the operation before exiting the application.	Selected
	Max Editors in a Single Open Operation	Specifies the maximum number of editors allowable from a single Open operation.	5
	Max Entries in Output Window	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
Oracle	**Data Dictionary View Usage**	Lets you select the roles/privileges that can view Data Dictionary Usage.	DBA Role
	Preserve Case in Object Identifiers	Preserves case of the database object.	Not selected
Sybase and **SQL Server**	**Rename action style**	Lets you specify whether user-defined objects or datatypes are renamed using an extended ALTER or using the `sp_rename` procedure.	extended alter

Grid properties (Results window) options

After opening the Options editor (see Specifying DBArtisan 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
Titles and Gridlines	**3D-Buttons**	Enables or disables a 3-D appearance on row and column headings in the grid.	Set
	Horizontal Gridlines	Enables or disables ruling between rows of the grid.	Set
	Vertical Gridlines	Enables or disables ruling between columns of the grid.	Set
	Mark Headings	Enables or disables highlighted row and column headings.	Set
Preview	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 - Editor tab
- ISQL options - Auto format tab
- ISQL options - Index Manager tab
- ISQL options - Code Assist tab

ISQL options - Main tab

After opening the Options editor (see Specifying DBArtisan 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:

Max Errors Before Aborting Execution	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. DBArtisan refers to this value when step executing SQL scripts.
Execute Selected Text	Executes a portion of a selected SQL script.
Check Syntax When Executing	For DB2, required to execute DB2 call statements in the ISQL Window.

Automatically lock connection	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.
Prompt to lock database connection	Locks the database connection on execution.
Tabs	Sets the appearance of the ISQL Window tabs to either the top or bottom of the ISQL Window.
File Association	Specifies whether the application should open an unknown file type automatically into an ISQL Window or prompt you with a message that DBArtisan does not recognize the file type.
Query Plan Layout	Sets the default orientation of a graphical query plan. For more information, see Query options.

ISQL options - Oracle tab

After opening the Options editor (see Specifying DBArtisan 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
Enable DBMS Output	Lets you specify **Buffer** size. 0 is the default.
Auto-Commit changes	Applies auto commit status changes to all open windows.
View xmltype as clob	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 **OciTypeBinder** conversion error. With the option selected, the SELECT submitted by DBArtisan is modified to include a getclobval() method call.
Default Query Plan	Lets you select the default display, tree-based or graphical, when you generate a query plan. For more information, see Viewing a tree-based or graphical Query Plan (Oracle).
Load Query Options	If this option is not enabled, DBArtisan 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 DBArtisan to load a previously saved query options file each time an ISQL windows opens against this DBMS platform. This lets you override DBArtisan's default query options. 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 Query options.

ISQL options - DB2 tab

After opening the Options editor (see Specifying DBArtisan 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
Set Isolation Level	Sets the default for the value of the **Isolation Level** option on the Query Options dialog. For details, see Query options.
Auto-Commit changes	Applies auto commit status changes to all open windows.
Create Explain plan tables if required	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.

Option	Description
Create explain plan tables on the SYSTOOLS schema	If set to TRUE, Explain Plan tables are created on the SYSTOOLS schema. If the tables already exist in the user's default schema, DBArtisan 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.
Load Query Options	If this option is not enabled, DBArtisan 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 DBArtisan to load a previously saved query options file each time an ISQL windows opens against this DBMS platform. This lets you override DBArtisan's default query options. 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 Query options.

ISQL options - Sybase tab

After opening the Options editor (see Specifying DBArtisan 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
Auto-Commit changes	Applies auto commit status changes to all open windows.
Load Query Options	If this option is not enabled, DBArtisan 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 DBArtisan to load a previously saved query options file each time an ISQL windows opens against this DBMS platform. This lets you override DBArtisan's default query options. 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 Query options.

ISQL options - SQL Server tab

After opening the Options editor (see Specifying DBArtisan 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
Enable Set Query Options	Sets the default value of the **Send Set Options** setting on the Query Options dialog. For details, see Query options. Changing this value does not affect ISQL sessions currently open.
Auto-Commit changes	Applies auto commit status changes to all open windows.
Load Query Options	If this option is not enabled, DBArtisan 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 DBArtisan to load a previously saved query options file each time an ISQL windows opens against this DBMS platform. This lets you override DBArtisan's default query options. 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 Query options.

ISQL options - MySQL tab

After opening the Options editor (see Specifying DBArtisan 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
Load Query Options	If this option is not enabled, DBArtisan 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 DBArtisan to load a previously saved query options file each time an ISQL windows opens against this DBMS platform. This lets you override DBArtisan's default query options. 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 Query options.

ISQL options - DB2 OS390 tab

After opening the Options editor (see Specifying DBArtisan 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
Load Query Options	If this option is not enabled, DBArtisan 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 DBArtisan to load a previously saved query options file each time an ISQL windows opens against this DBMS platform. This lets you override DBArtisan's default query options. 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 Query options.

ISQL options - Editor tab

After opening the Options editor (see Specifying DBArtisan 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
Window	**Show Toolbar** and **Show Status Bar**	Enables/disables these ISQL Window user interface elements.	Selected
	Maximize on new or open	Indicates that DBArtisan 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
	File Tracking	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
	Auto-Reload File	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, DBArtisan prompts you before reloading your file if external changes have been saved.	Not selected

Group	Option	Description	Default
	Auto-Save File	Indicates that files in the ISQL Editor should automatically be saved at the indicated time interval. Backup files are saved to **C:\Documents and Settings***username***\Application Data\Embarcadero\DBArtisan\Backup**. Backup files use a naming convention of the form **~ET***xxxx***.tmp**., See the registry under **HKEY_CURRENT_USER\Software\Embarcadero\DBArtisan***version***\Backup** for a listing of backup file names and the path of the file to which they correspond.	Selected (5 minutes)
Command History	**Save File Before Overwriting**	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
	Save Most Recent	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
Printing	Lets you select common printer options.		
Appearance	**Enable Syntax Highlighting**	Sets syntax highlighting on so that all keywords and comments are colored for easier reading and debugging.	Selected
	Show Line Numbers	Places line numbers in the left column of an ISQL Window.	Selected
	Enable Outlining	Enables and disables outlining.	Selected
	Enable Text Wrapping	Enables and disables a typical text wrap feature.	Selected
	Editor Font	Sets the font face, style, and size displayed in the editor.	Available
	Syntax Colors	Sets syntax coloring for keywords, comments, quotes, and default text for various file types and scripts from the Syntax Coloring dialog box.	Available
Formatting	**Auto Indent**	Sets automatic indentation for each carriage return, new line in your SQL script.	Selected
	Expand Tabs and **Tab Size**	Sets tabs as the specified number of spaces in result sets.	Selected (4)

ISQL options - Auto format tab

After opening the Options editor (see Specifying DBArtisan 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
Keywords	Lets you select a character case (UPPERCASE, LOWERCASE, INITIALCAPS, or NOCHANGE) formatting treatment for SQL keywords.

Setting	Description
Right Margin	Specifies the maximum number of characters per line.
New line before keyword	If enabled, a new line is forced before every SQL keyword. If disabled, lines are only forced for statement type and clause type keywords.
BEGIN..END block	If enabled, empty BEGIN..END blocks occupy a single line.
Stack lists	If enabled, a new line and indenting is forced for each item in comma-separated list such as argument lists.
List indent size	If **Stack lists** is enabled, this control lets you specify the number of spaces each list item will be indented from the current offset.
Parenthesis indent size	Lets you specify the number of indenting spaces for forced lines following an open parenthesis.
Conditions format style	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.
Conditions stack threshold	For **Conditions format style** selections that specify stacking, this value specifies the minimum number of conditions that must appear in a clause before conditions are stacked.
THEN statements	If enabled, simple THEN clauses are kept to a single line.

When finished, click **OK**.

ISQL options - Index Manager tab

After opening the Options editor (see Specifying DBArtisan application and feature options), you can make changes to the **ISQL - Index Manager** tab.

The datasource index stores information required to resolve references to the tables, table columns, views, functions, procedures, packages, and synonyms for registered datasources. Required to support object name-based features such as Semantic Validation and Code Complete, the datasource index provides rapid lookup and assists in verification of object names.

> **NOTE:** For information on the features requiring datasource index support, see Code Assist.

When indexing is enabled, DBArtisan automatically builds and incrementally maintains the datasource index as follows:

- When you connect to a datasource that is currently unindexed, DBArtisan will update the index with entries for the new datasource
- The index is automatically updated with any DBArtisan operation that results in an ALTER, CREATE, DROP, RENAME, or EXECUTE action
- The index is automatically updated on a DBArtisan Refresh operation against object types that support that operation

In addition, the datasource index stores a timestamp at the object type level for each datasource. When you connect to a datasource for which a time stamp for any object type is older than the specified **Index expiration time**, new index entries for all objects of that type on the datasource are generated.

While the datasource index is largely self-maintaining, the Options editor's **ISQL - Index Manager** tab provides controls for manual maintenance of the index. You can clear the existing index, start creating a new index after clearing, and stop the indexing process. Typically, you only use these options in problem-solving and for manual intervention.

The table below describes the options and functionality on the **ISQL - Index Manager** tab.

Option	Description
Enable Indexing	Enables the datasource indexing feature. Manual and automatic indexing updates cannot be executed unless this feature is enabled.
Log SQL	This setting lets you enable or disable writing of the the SQL issued to create and maintain indexes to the SQL log.
Start Indexing	Updates the datasource index entries for the current datasource if entries exist for the connected datasource in the index. If no entries are present in the index, new entries for the datasource are created in the index.
Stop Indexing	Stops any automatic indexing updates currently in progress.
Objects to index	For objects not indexed, Code Assist features will not be available for those object types. Semantic validation applies to tables, views, columns, and synonyms. Code Complete applies to functions, procedures, packages, tables, views, columns, and synonyms.
Clear Index	Deletes the current datasource index. Automatic index updates, such as when connecting to a new datasource or making a subsequent change to the connected datasource, will commence immediately. You can also start indexing the entire datasource to which you are connected, using the **Start Indexing** button.
Index expiration time	When an object type timestamp older than this value is detected on a datasource, DBArtisan updates the index for all objects of that type on the datasource.
Index location	Specifies the location of the files used to store the datasource index. A change to the location results in a new index being created next time that the application is started.

ISQL options - Code Assist tab

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

Group	Option	Description
SQL Assist	**Enable auto-activation** and **Auto-activation delay**	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 **Auto-activation delay**. For more information, see Code Assist.
	Insert single proposals automatically	Specifies that if a Code Complete suggestion list would contain only a single suggestion, that suggestion is inserted automatically.
	Fully qualify completions automatically	Specifies that code completion results are returned fully qualified, rather than the minimum required to identify the object.
Semantic Validation	**Severity levels for semantic validation problems**	Lets you select the severity level (ERROR, WARNING, or IGNORE) associated with detected semantic errors. The default, WARNING, 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 About semantic validation..
SQL Parsing	**Enable Parsing**	SQL Parsing must be enabled if you are using features such as on-the-fly syntax checking, semantic validation, and Code Complete.

Group	Option	Description
	Parse always	If **SQL Parsing** is enabled, enabling this option ensures that SQL Parsing is always active.
	Parse based on file size and **Maximum file size**	If **SQL Parsing** is enabled, these options place an upper limit on file size. SQL Parsing for a file will be disabled when a file exceeds the specified size.

Java options

After opening the Options editor (see Specifying DBArtisan 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
Load Java files	**Use the DBMS_JAVA package**	Lets you schedule your SQL.	Selected
	Use batch file (oracle_home/bin/loadJava.bat)	Uses batch file (oracle_home/bin/loadJava.bat)	Not selected
	Default Encoding option	Leave blank to use the default.	Default
Drop Java files	**Use the DBMS_JAVA package**	Lets you schedule your SQL.	Selected
	Use batch file (oracle_home/bin/dropJava.bat)	Uses batch file (oracle_home/bin/dropJava.bat)	Not selected

JDBC options

After opening the Options editor (see Specifying DBArtisan application and feature options), you can make changes to the **JDBC** tab.

DBArtisan 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:		
Initial Heap Size	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
Maximum Heap Size	Set the upper limit for your heap size.	64 MB
Additional Options	Add options here ONLY in consultation with Embarcadero Technical Support	N/A
Default JDBC Driver		
DBMS Type	Select the DBMS appropriate to the driver you want to review/adjust from the drop-down list.	
Search	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.	

Option	Description	Default
	Sybase	
JConnect metadata procedures are installed on all servers	**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 DBArtisan 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
Log all SQL Statements to a File	Indicates that the application should log all of the SQL that it executes to a file. SQL logging provides an audit trail for DBArtisan. You can examine this log to determine the SQL executed to complete a task.	Not selected
Logfile Path	If you choose to log the SQL generated by the application, specify the drive, directory, and file name.	None
Max File Size	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
Truncate	Empties the entire contents of the logfile.	Not available
Log all Output Messages to a File	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
Logfile Path	If you choose to log the server messages generated in the Output window, specify the drive, directory, and file name.	None
Max File Size	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
Truncate	Empties the entire contents of the output logfile.	Not available

Migration options

After opening the Options editor (see Specifying DBArtisan application and feature options), you can make changes to the **Migration** tab. The **Migration** tab sets the default values effective when you open the Migration Wizard, a utility that lets you migrate schema and data. For informtion on this utility, see Using the schema migration wizard.

MySQL Utilities Options

After opening the Options editor (see Specifying DBArtisan application and feature options), you can make changes to the **MySQL Utilities** tab. DBArtisan integrates with MySQL utilities. For DBArtisan 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 DBArtisan application and feature options), you can make changes to the **Oracle Utilities** tab.

DBArtisan integrates with multiple Oracle utilities. For DBArtisan 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 DBArtisan application and feature options), you can make changes to the **Performance Center** tab. DBArtisan lets you customize Performance Center's integration with DBArtisan. The table below describes the options and functionality on the Perf Center Tab:

Option	Description	Default
Computer Name or IP Address	Specify where the Performance Center Server is installed.	localhost
Port Number	Specify the port for the Apache server or IIS web server.	80
Test	Verifies the settings.	
Connect to the server using	Specify if you want DBArtisan to open the **Web Client** or the Performance Center file (**PerfCntr.exe**) within DBArtisan.	Web Client

Query Builder options

After opening the Options editor (see Specifying DBArtisan 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
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.	Not selected
	Generate SQL/92 if supported by DBMS	SQL/92 is a standard for relational database management systems.	Not selected
Execution	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
General	**Show Column Data types in Query Diagram**	Reveals the data type in each column for tables in the SQL Diagram pane.	Not selected
	Confirm on Item delete	Opens a Confirm Delete dialog box when an item is deleted.	Selected
	Auto Populate Views	Checks syntax every time an execute statement, refresh or copy statement begins.	Not Selected
	Auto Format	Automatically sets style and spacing of display.	Selected
Auto Join	**Run Automatically**	Automatically detects names and data types, and create joins for multiple tables.	Selected
	Require Indexes	Joins only indexed columns. Requires an indexed column for joins.	Selected
	Require same data type	Automatically joins columns with the same data type.	Selected
Syntax Checker	**Automatic Syntax Check**	Automatically checks SELECT and CREATE VIEW statements for errors.	Selected
	Warn on non index join	Returns a warning when it detects a join against a non-indexed column, or a column not participating in a primary key	Not selected
Display	Lets you sets the style, size, and color of **Column Font** and **Title Font**. Also lets you set the background **Table Color** for the SQL Diagram Pane.		Available

Reports options

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

- Append a date and time time stamp to report file names. This prevents existing reports from being owverwritten.
- Save reports in date-specic 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 DBArtisan 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	
Results Option pane	**Result Window**	**Single Window** - Displays all results in one tabbed result window. Multiple result sets together in the window. Single Window and Multiple Windows options are mutually exclusive. **Multiple Windows** - Displays multiple result sets one result set per window. **Attached to Editor** - 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. **Unattached** - Sets results appear separate from the ISQL Window. Used in conjunction with Single Window option or Multiple Windows option. **Reuse Window** - Sets new result sets to overwrite any current result sets in an open Result Window. Only valid for Single and Attached to Editor combination.
	Results File	**Mail File Type** - 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. **Schedule File Type** - Selects the schedule file type. Valid formats include Tab delimited, Comma separated, and HTML. **Include column titles when saving** - Includes column titles when saving a result set. If this option is turned off, column titles DBArtisan does not save result sets.
	Result Set Options	**Default Rowcount** - Lets you limit the number of rows returned to the result window of the ISQL window (default 0). **Sybase and SQL Server: Text Size (bytes)** - Lets you specify the text size (default 8192). **Oracle:LONG Size (bytes)** - Lets you specify the LONG size (default 8192). **LOB Preview: Text Size (bytes)** - Specifies the length of the preview of LOB column data (default 4096).
	Sort Type	Lets you select a result sorting: **Alphanueric** or **Lexicographic**.
Format pane	**Column Formatting**	**Auto Format (Best Fit)** - 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. **Use pre-defined column** - Lets you select column type and character length
	Enable Date/Time Format	Lets you select the date/time format.

Interface Element	Options and descriptions	
	Format	**Standard Grid** - Displays all result sets in a standard grid format. DBArtisan displays result sets in grid format in IISQL Editors that are opened after you have selected this option. DBArtisan does not display IISQL Editors that are already open. **HTML** - Displays all result sets as HTML tables. DBArtisan displays result sets in HTML format in ISQL Editors that are opened after you have selected this option. DBArtisan does not display ISQL Editors that are already open. **ASCII Text** - Displays all result sets as ASCII Text. DBArtisan displays result sets in ASCII Text format in ISQL Editors that are opened after you have selected this option. DBArtisan does not display ISQL Editors that are already open. **Grid Font** and **Printer Font** buttons - Opens a Font dialog box, letting you select the font, style, and size for the result sets grid or printed result sets.

SMTP mail options

After opening the Options editor (see Specifying DBArtisan 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
Send messages through SMTP	Enables SMTP messaging and makes the other controls on this tab available.	Unselected
Name	Name that appear as the e-mail sender.	Name you specified during installation.
E-mail Address	Address to send e-mails from DBArtisan.	E-mail address you specified during installation.
Authentication	Lets you specify authentication options.	None
User Name	User name for authentication.	Not available
Password	Password for authentication.	Not available
Host Name	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.
Port Number	Port number you connect to on your outgoing SMTP server.	25
Test	Opens an SMTP Configuration Test e-mail addressed to your e-mail address. Click Send Mail to send the e-mail.	Available
Bind to	Your IP address the message is bound to.	ANY_IP_ADDRESS
Encoding	E-mail encoding.	Western Europe (ISO)
Send messages Mime encoded	Messages encoded using Multipurpose Internet Mail Extensions (MIME) support enriched content and attachments.	Selected
Send all messages as HTML	Messages include text formatting.	Selected
Auto Connect to the Internet	DBArtisan connects to internet at launch.	Selected

Space Analyst options

After opening the Options editor (see Specifying DBArtisan application and feature options), you can make changes to the **Space Analyst** tab. The table below describes the options and functionality on the **Space Analyst** Tab:

Tab	Option	Description	Default
Oracle	**Exclude objects owned by SYS in Reorganization Wizard.**	The Space Analyst Reorganization Wizard excludes SYS objects.	Selected
	Auto populate Table Reorganization Details Grid	Space Analyst auto-populates table reorganization details.	Not Selected
	Auto populate Index Reorganization Details Grid	Space Analyst auto-populates index reorganization details.	Not Selected
	Populate Index Reorganization Details Grid with Deleted Leaf Statistics	Space Analyst gathers deleted leaf row statistics for an index and displays details.	Not Selected
SQL Server, Sybase, and IBM DB2 for Windows, Unix, and Linux	Offers the option to update the Table Details Grid and Index Details Grid.		Not Selected

Warnings options

After opening the Options editor (see Specifying DBArtisan application and feature options), you can make changes to the **Warnings** tab. DBArtisan 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
Create an index on the same tablespace as the associated table	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
Create an index on the same tablespace as the associated table	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
Create an object in the SYSTEM tablespace	Issues a warning message whenever a user tries to create or place an object on the SYSTEM tablespace.	Oracle
Create a user with default or temp tablespace as the SYSTEM tablespace	Issues a warning message when a user is created with a default or temp tablespace on the SYSTEM tablespace.	Oracle
Create an object in the master database	Issues a warning message when an object is created in the master database.	Microsoft SQL Server and Sybase ASE
Create a table or index on the default segment	Issues a warning message when a table or index is created on the default segment.	Microsoft SQL Server and Sybase ASE
Create a non-clustered index on same segment as the associated table	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 DBArtisan

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
Datasource Management	This section describes the purpose of datasources, the process of establishing datasource connections, and managing your datasources and datasource groups.
Supported Objects	This section describes the different platform objects the application supports.
Modifying objects using editors	This section describes Object editors. Editors let you modify existing objects.
Common functions performed against objects	This section describes the different platform object functionalities.
SQL Scripting	This section describes the SQL scripting environment, that lets you write, debug, test and deploy solid SQL code for your database applications.
Permissions Management	This section describes the Permissions Management features.

Datasource Management

After installing DBArtisan, 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. DBArtisan stores information about the datasource specification in the system registry and provides a visual interface for maintaining it.

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

Registering or Editing a Datasource

A datasource is the connection profile that lets you have access to a database, and, therefore, to your data. So, when you register a datasource, you're telling DBArtisan how to get at the database(s) you want to operate on. In fact, all database activities are performed through datasources, so you need a datasource profile for each database instance (or database server) in your enterprise. Each connection profile (i.e., datasource) includes a unique name, the connection string, the default user ID, and optionally, a password.

DBArtisan requires a JDBC connection. If you did not already have a JDBC driver, one was included as a part of the installation process (to see what's available, go to the \Java\drivers\ subfolder of the main DBArtisan installation folder. The DBArtisan installation provides a JDBC Type-4 thin driver for each of the platforms except IBM DB2 (see note below) if one was not previously installed.

NOTE: JDBC Type-4 Drivers require TCP/IP. If you are upgrading from a prior version of DBArtisan and want to connect to a datasource that is not using TCP/IP, you will have to edit your current datasource registration information.

NOTE: IBM DB2 does not use a Type-4 "thin" JDBC driver because it won't connect to DB2 version 7 or the OS/390. Instead, IBM DB2 packages a Type-2 driver by default.

Datasource Registration

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

To register a new datasource:

- From the **Datasource** menu, select **Register Datasource**.

To register a new datasource based on an existing definition:

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

To edit registration information:

- On the **Explorer**, right-click a datasource, and select **Edit** from context menu.

The table below describes the fields and controls you will encounter when creating or editing a datasource registration:

Panel/Tab	Option	Description
Database Type	Select the radio box corresponding to the DBMS type for the datasource. If there is currently a DBMS node selected in the Explorer, that database type is automatically selected.	

Panel/Tab	Option	Description
Connection Information	For detailed descriptions of the options on this tab, see Datasource connection information on page 90.	
Advanced (only available on clicking the **Advanced** button on the **Connection Information** tab)	**JDBC Driver**	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 JDBC Driver Editor. **NOTE:** If you select a driver other than the default, you'll be queried about whether you want that driver to be used for datasources you register going forward. Datasources already registered will not be affected.
	Connection URL	The connection URL cannot be edited--it's for reference only.
	URL Check	Indicates whether the connection URL is valid or not.
Security Parameters	**User ID**	The User ID that DBArtisan will use to connect to the datasource.
	Password	The password associated with the User ID.
	Connect As: (Oracle)	When relevant, choose the appropriate user/administrator level.
	Domain:	For MS SQL only. Identify the domain if the user has restricted access.
	Auto connect?	Spares the user from reentering the password every time you connect.
	Connect using Windows Authentication or **Connect using OS Authentication** (IBM DB2 for Windows, Unix, and Linux, MySQL, Oracle, SQL Server)	Login to the server is verified using Windows/OS authentication
	Connect using Kerberos Authentication (Sybase)	If this option is selected, login to the server is verified using Kerberos authentication. **NOTE**: This is only available of the **Use alias information in the SQL.INI file/Sybase Server** option on the **Connection Information** tab is selected. For details, see Datasource connection information.
	Test Connection	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.
Datasource Properties	**#Include Search Directories**	Enter one or more paths on this datasource, which will be searched for files in conjunction with use of the **#include** 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 #incliude directiive, see SQL preprocessing: #define and #include. Note that if there are no entries specified here, the directory specified on the **Directories** tab of the Options editor will be searched. For more informtion, see Directories options.
	Database device default path	The database device default path is used when creating a new database device.
Database Group	Select the datasource folder under which you want this datasource to appear.	

>

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

NOTE: On initial connection to a Sybase Cluster edition datasource, 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 datasource. You can subsequently change the connection scope for the datasource by editing the datasource.

If you have completed the datasource registration and still encounter problems, please refer to the Unable to Connect? discussion.

NOTE: After you successfully connect to a datasource, you can see the name of the host and the connection string information displayed at the bottom of the DBArtisan window. For Oracle and Sybase ASE, should a question arise, you can see if perhaps you set up the datasource 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.

Datasource connection information

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

DBMS	Option	Description
IBM DB2	**Server**	The name of the server.
	Schema ID	This is an optional field. If you want to include the current schema as part of the datasource connection properties, you can set the implicit schema for unqualified object references to a schema different from the user's login id.
	Function Path	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.
	Datasource Name:	This field is automatically populated with the server name, but you can rename it to whatever you want.
SQL Server	**Use Network Library Configuration**	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.
	Alias	This box is enabled only when you opt to use the Network Library Configuration
	Protocol: TCP/IP or Named Pipes	You need to select one or the other to register a SQL Server database.
	Host	Enabled only when you have not selected an Alias from the Network Library.
	Port (optional)/Pipe Name	Depending on your section, you can optionally indicate the port for TCP/IP or the pipe name depending on your means of connection.
	Default DB (optional)	Optionally give the default name for the database.
	Datasource Name	The field is automatically populated with the host name, but you can rename to whatever you want.
MySQL	**Server**	Enter the name of the host, for example doctest01.
	Port (optional)	The default port is 3306. You do not need to enter this information.
	Default Database	You must enter the name of the default database.
	Datasource Name	This field is automatically populated with the server name, but you can rename it to whatever you want.

>

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

For context information, see Registering or Editing a Datasource.

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 datasource.

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.
Driver supports the use of native aliases	DBArtisan cannot detect whether the driver you are using can support native aliases. You need to know by checking the driver's documentation.

Required Information	Description
Required jars	DBArtisan 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.

For context information, see Registering or Editing a Datasource.

Unable to Connect?

If you are having trouble connecting to a datasource, as a first step please check the connection information to make sure it has been captured correctly.

1 Verify that the user/password combination is correct.

2 Check that the correct JDBC Driver is being called. You can see all platform-specific drivers when you browse from the Advanced Settings panel of the registration wizard. You can see ALL drivers by going to the \Java\drivers\ subfolder of the main DBArtisan installation folder. A list of supported drivers follows.

Drivers supported by DBArtisan:

Oracle:

JDBC OCI client-side driver--Type 2, Versions 8, 9, and 10. This is shipped by Oracle.

Oracle 9.0.2 Thin Driver--Type 4, shipped with DBArtisan.

Sybase:

jConnect v5.5--Type 4, shipped with DBArtisan.

Microsoft SQL Server:

jTDS Driver--Type 4, shipped with DBArtisan.

NOTE: This is the only driver for MS SQL Server that we support.

IBM DB2--Type 2 driver comes from IBM.

MySQL:

MySQL driver--shipped with DBArtisan.

3 Verify that the connection URL is correct. The URL always begins with JDBC and includes a permutation of the driver.

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.

Datasource Catalog

The Datasource Catalog is a collection of defined datasources. It is stored in the system registry of your computer. All Embarcadero database management products share the datasource catalog, which means that when you set up your datasource catalog using one product such as DBArtisan, the same list of datasources is available in other Embarcadero Technologies products. Any changes you make to the datasource catalog are reflected in all Embarcadero database management products.

Sharing a datasource Catalog

You can configure Embarcadero database applications to use a datasource catalog stored in the system registry of your machine (local).

You can determine the location of your datasource catalog by inspecting the application status bar. The 'Catalog' indicator displays 'Local' if the catalog is stored in your local system registry.

When you configure the location of the datasource catalog, you are affecting all Embarcadero datasource applications as they all read and maintain the same datasource catalog.

NOTE: You need the proper permissions to read/write registry entries on another computer. See your network administrator if you have problems.

Automatically Discovering Datasources

The first time you run DBArtisan a dialog box displays, giving you the option to Auto-Discover all configured datasources. If you click Yes, the DBArtisan 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 system call to your network. Provide login information (user name and password) the first time you connect to a datasource. For more information, see Connect.

NOTE: IBM DB2 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 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 Datasource Management.

Changing Datasource Groups

DBArtisan 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.**

DBArtisan opens the **Change Datasource Group** dialog box.

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

3 Click **OK**.

DBArtisan changes groups.

Connect

DBArtisan lets you set datasources to automatically connect each time you open the application.The first time you start the application, DBArtisan 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. For details, see Automatically Discovering Datasources.

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.

If you later want to automatically connect your datasources, you can edit the datasource to make that change.

TIP: To configure your datasource to login automatically, refer to Registering or Editing a Datasource.

To manually connect to a datasource:

1 In the left pane of the Datasource Explorer, right-click a datasource and select **Connect** from the context menu.

DBArtisan opens a **Login to...** dialog box.

2 Type a **Login ID** and **Password**.

3 Optionally, and only if available, provide **Login As**, **Default SChema ID**, and **Default Function Path** values.

4 Select **Auto Connect** to automatically connect to the datasource in the future.

5 Click **OK**.

DBArtisan connects to the datasource and expands the datasource display to show the top level of nodes available on that datasource. For more information, see Datasource Explorer nodes by DBMS platform.

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, DBArtisan attempts to reconnect to a registered datasource, if available.

Completing the Disconnect Dialog Box

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

1 On the left pane of the Datasource Explorer, right-click a datasource and select **Disconnect** from the context menu.

NOTE: If necessary, DBArtisan 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.

Discover Datasource

DBArtisan discovers datasources residing on your system that are not currently registered datasources through a Windows system call to your network. DBArtisan can retrieve a list of discovered datasources, 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**.

DBArtisan opens the Discover Datasources dialog box and when the search is complete, displays all discovered datasources.

2 Select the check box next to any datasources you want to register.

3 Click **Register**.

DBArtisan registers the datasource or datasources selected.

Extracting configuration statements

DBArtisan lets you extract the sp_configure statements for setting and displaying your server configuration parameters into an Interactive SQL window.

1 On the left pane of the Datasource Explorer, expand the **Instance** or **Server** node for a datasource and then select the **Configuration** node.

DBArtisan displays the server configuration parameters in the right pane of the Datasource Explorer.

2 In the right pane of the **Datasource Explorer**, select one or more paramerters, right-click tand select **Extract**.

DBArtisan opens the DDL Editor. For more information, see ISQL Editor.

For more information, see Servers.

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**.

DBArtisan opens the Manage Datasources dialog box.

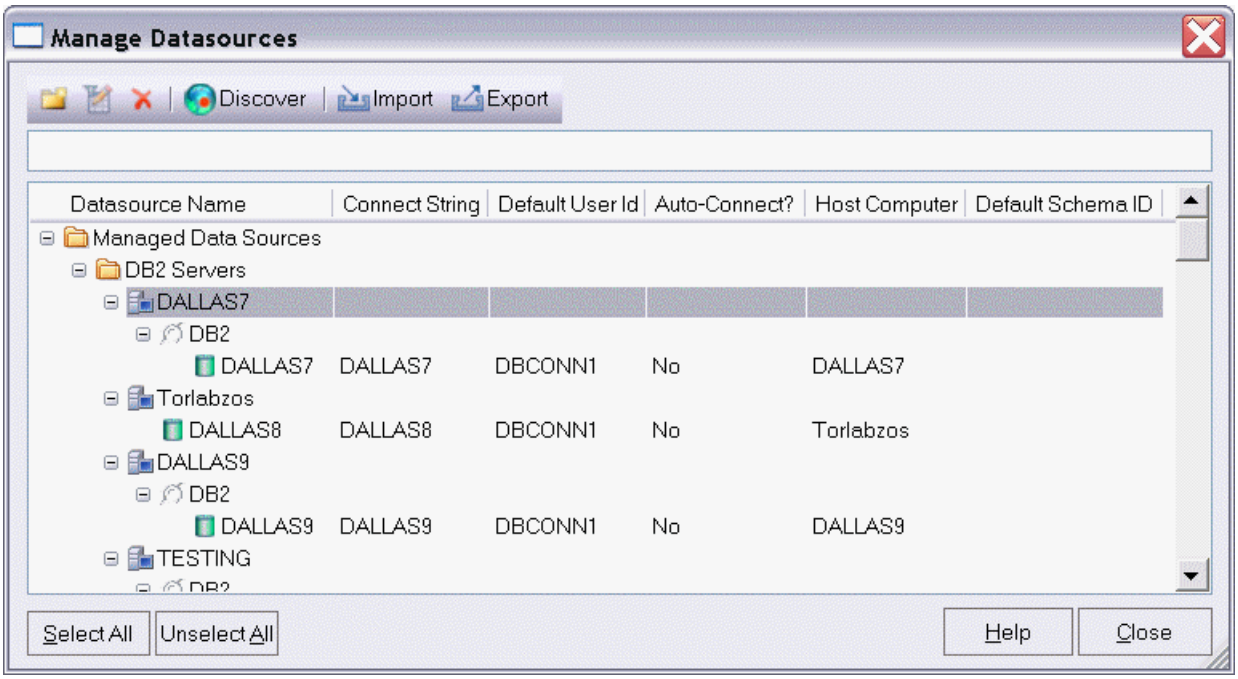

The grid format lets you access and view datasource specifications. datasources are grouped according to the default, DBMS-based folder structure and then within any custom datasource groups you have set up. The table below describes information available for each datasource entry:

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.
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 a Datasource.

- Select a datasource and click the Edit Datasource button to edit that datasource definition. For more information, see Registering or Editing a Datasource.
- 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

The ability to export and import datasource definitions can save time in organizations with large numbers of datasources and multiple users. For example, 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 (`.dsd` file suffix) files.

To export one or more datasource definitions:

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

   DBArtisan 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**. A **Browse For Folder** dialog opens.
4. Use that dialog to locate and select a target folder and then click **OK**.

Each datasource definition is stored in a separate file. User ID and password credentials are not exported to the datasource definition file. Other users can then import the datasource definition file.

To import one or more datasource definition files:

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

   DBArtisan opens the **Manage Datasources** dialog box.

2. Select the datasource group folder where the files or files are to be imported.
3. Click **Import**. A **Select datasource definitions to import** dialog opens.
4. Use that dialog to locate and select one or more datasource definition files (`.dsd`) to be imported and then 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**.

   DBArtisan opens the **Datasource Properties** dialog box.

For more information, see Datasource Management:

Selecting Datasources

The DBArtisan **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. Use **Connect** to connect to the datasource.
4. If available, click the **Database** list, and then click the target database.
5. Click **OK** to close the **Select Datasource** dialog box.

For more information, see Datasource Management.

Unregistering Datasource

DBArtisan lets you unregister datasources when you no longer need them.

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

1. On the left pane of the Datasource Explorer menu, right-click the target datasource and select **Unregister** from the context menu.
2. When prompted to confirm, click **Yes**.

   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 Datasource Management.

Datasource Groups

DBArtisan 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

DBArtisan offers the following functionality for Datasource Groups:

- Removing a datasource group
- Creating a new datasource group
- Renaming a datasource group

Removing a datasource group

DBArtisan 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 containining group.
- The default datasource groups correspond to the DBMS platforms supported by DBArtisan. 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**. DBArtisan prompts you to verify.
2. Click **Yes**.

For more information, see Datasource Groups

Creating a new datasource group

DBArtisan 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.

> **TIP:** SQL Server does not provide true clustering because separate instances of a database are actually running on each server in a cluster. To simulate a consolidated view, you can name and group clustered datasources within the Explorer by setting up a datasource group, named, for example, Cluster3. After you name the new group, simply drag-and-drop the datasources that make up the cluster into Cluster3.

Completing the New Datasource Group Dialog Box

1. In the Datasource Explorer tree take one of the following actions:
   - To create a new folder at the top level, right-click in empty space and select **New** from the context menu.
   - To create a folder within an existing folder, right-click the target folder and select **New** from the context menu.

   DBArtisan 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**.

Renaming a datasource group

DBArtisan lets you change the name of an existing datasource group.

To rename a datasource group:

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

   DBArtisan opens the **Rename Datasource Group** dialog.

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

3. Click **OK**.

Servers

DBArtisan lets you manage your servers from the Datasources menu. The Datasources menu provides the following options:

- Starting, pausing, and resuming a server (SQL Server)
- Shutting down a Server (SQL Server, Sybase)
- Configuring the Server
- Extracting configuration statements

Starting, pausing, and resuming a server (SQL Server)

DBArtisan lets you start a server that has been shut down. As well you can also pause and resume the SQL Server service.

> **NOTE:** For information how to shut down a server, see Shutting down a Server (SQL Server, Sybase).

To start a server, pause a server or resume after pausing:

- Use the **Datasource** menu's **Start Server**, **Pause Server**, and **Continue Ater Pause** commands.

Shutting down a Server (SQL Server, Sybase)

DBArtisan lets you shutdown a server using the **Shutdown Server** dialog box.

Completing the Shutdown Server Dialog Box

1. In the **Explorer**, select the target datasource.

2. On the **Datasource** menu, click **Shutdown**.

   DBArtisan opens the Shutdown Server dialog box.

The table below describes the options and functionality on the Shutdown Sever dialog box.

Option	Description
Normal	Shuts down the server in an orderly fashion. DBArtisan disables logins, and waits for currently executing Transact-SQL statements and stored procedures to finish. **Wait Time** (Sybase ASE only) - Lets you enable a grace period, an upper limit on how long to wait for normal operations to complete before forcing a shutdown. Wait time is specified in hours, minutes, and seconds, and if this option is not enabled, the Sybase ASE server will wait indefinitely before shutting down.
Immediate	Shuts down the server immediately. DBArtisan does not perform checkpoints in every database. The server terminates all user processes and rolls back any active transactions.

For more information, see Servers.

Configuring the Server

To configure and tune your sever, DBArtisan offers the **Server Configuration** dialog box that lets you view and edit the configuration parameters for each server. These configuration parameters are stored in the system catalog and used by DBArtisan when restarting the server. For detailed information on the meaning and purpose of each server configuration parameter, see the System Administration Guide and Performance and Tuning Guide.

NOTE: Many changes to server configuration parameters require you to shutdown and to restart DBArtisan before the changes take effect.

Completing the Server Configuration Dialog Box

NOTE: For Oracle the following dialog box is called Instance Parameters, and for SQL Server it is Server Configuration.

1 On the **Datasource** menu, click **Configure**.

The table below describes the options and functionality on the Server Configuration dialog box.

Option	Description
Edit Button	Click to modify the target parameter. For more information, see Edit Configuration Dialog Box.
New Value	Lets you type the value for the parameter.

Oracle Managed Files (OMF)

For Oracle 91 only, DBArtisan supports Oracle Managed Files (OMF.) OMF automatically manages files for tablespace, redo logs, and control files. Oracle automatically creates and removes operating system files as you create and drop storage objects.

You can enter the following query to set the db_create_file_dest parameter at the system level. You can set the parameter value for the session level by using Alter session clause:

```
ALTER SYSTEM SET DB_CREATE_FILE_DEST = '/u01/oradata/sample2'
```

TIP: You can also set the parameter at the system level in the Edit Configuration dialog box.

For more information, see Servers.

Edit Configuration Dialog Box

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

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

NOTE: To modify a parameter, the parameter must have a Yes in the Session Modifiable or System Modifiable column of the Instance Parameter grid.

For more information, see Servers.

Session Recording

DBArtisan includes a session recording feature that records changes to the server. The Session Recording feature lets you create a script file that you can execute as a batch at another time.

The session recording feature automates repetitive tasks that you routinely perform by issuing the commands in the collected script with a single command. Session recording lets you:

- Streamline time consuming, repetitive tasks.
- Execute multiple actions with from a single file.
- Automate a complex series of tasks by creating a batch file.
- Distribute automated tasks to multiple users with the session recording file.

The Session Recording feature is not available for operations that are destructive to your server. These operations include:

- Schema Migration
- Data Load
- ISQL and DDL Editors

When you activate session recording, these operations are grayed out and not available.

You can also use the session recording feature for Tivoli Integration. DBArtisan records schema change management sessions to instruction files, which are read by Tivoli Management Systems. Tivoli subsequently transports these to multiple managed endpoints for local execution.

For more information, see Recording a Session and Executing a Session File.

Recording a Session

You can create a session file that includes a series of commands or actions using DBArtisan **Session Recording Mode** dialog box.

>

Completing the Session Recording Mode Dialog Box

1 On the **Datasource** menu, click **Session Recording**.

The table below describes the options and functionality on the Enter Session Recording Mode dialog box.

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. DBArtisan uses *.ses file extension for session files.

NOTE: DBArtisan appends each task that you record in the session file. In place of an Execute button, DBArtisan displays a Record button. To record a certain task, click the Record Button to open the Task Recording dialog box.

2 When you want to exit session recording mode, open an ISQL Window that lets you review the session file.

3 On the **Datasources** menu, click **Session Recording**.

DBArtisan opens an SQL window.

For more information, see Session Recording and Executing a Session File.

Executing a Session File

You can execute a session file which includes a series of collected commands in a single batch from an ISQL Window.

1 On the **File** menu, click **Open**.

DBArtisan opens the Open Files dialog box.

2 Select the session file(s) (*.ses) from the file list, or in the **File Name** box, type a session file name.

3 Click **OK**.

DBArtisan opens the session file into a SQL window.

4 On the **SQL** tool bar, click the **Execute** button.

5 On the **SQL** tool bar, click the **Step-Execute** button to see a graphical display to identify the execution path of the script, and then click the **Execute** button to execute the session file.

6 On the **SQL** tool bar, click the **Query Plan** button to see a graphical display to identify the execution path of the script.

7 Click the **Execute** button.

DBArtisan executes the script.

For more information, see Session Recording.

Database Object Finder

The Database Object Finder is a search utility that lets you search through all your database objects. You can narrow your search by type of object and object owner, and you can set the search direction. The Database Object Finder also includes many options that let you further focus your search. The Database Object Finder displays your search results in the Database Object Found box at the bottom of the Database Object Finder dialog box. For successful searches, DBArtisan lets you open the appropriate object editor from the Database Object Finder.

>

Completing the Database Object Finder

1 Connect to a datasource.

2 On the **Datasource** menu, click **Find Object**.

The table below describes the options and functionality on the Database Object Finder dialog box.

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**, or **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.

NOTE: If the search is successful, DBArtisan displays the search string in the Database Object Found box. Click the Open button and DBArtisan opens the object editor for the target object.

Changing an ODBC Datasource to a Native Datasource

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

Microsoft SQL Server uses ODBC to connect to Microsoft SQL Servers. DBArtisan requires native connectivity. To connect through DBArtisan, 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.
Port number	In the box, type the port number of the target computer.

3 Click **OK**.

DBArtisan returns to the Client Network Utility dialog box.

>

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

DBArtisan adds the server.

5 Open **DBArtisan** and the Datasource Registration wizard. For more information, see Datasource Registration.

For more information, see Datasource Management.

Supported Objects

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

Objects	IBM DB2 for Windows/Unix	IBM DB2 for OS/390 and z/OS	Microsoft SQL Server	MySQL	Oracle	Sybase ASE
Aliases	X	X	X			X
Dump and Backup Devices			X			
Bufferpools	X	X				
Check Constraints	X	X	X		X	X
Clusters					X	
Control Files					X	
Database Devices			X			X
Database Links					X	X
Databases		X	X			X
Data Caches						X
DBRM		X				
Defaults			X			X
Directories					X	
Event Monitors	X					
Extended Procedures						X
Foreign Keys	X	X	X	X	X	X
Functions	X	X	X	X	X	X
Groups	X		X		X	X
Indexes	X	X	X	X	X	X
Java Classes					X	
Java Resources					X	
Java Sources					X	
Job Queues					X	
Libraries					X	
Linked Servers			X			
Logins			X			X
Materialized Query Tables	X					

Objects	IBM DB2 for Windows/Unix	IBM DB2 for OS/390 and z/OS	Microsoft SQL Server	MySQL	Oracle	Sybase ASE
Materialized Views					X	
Materialized View Logs					X	
Nodegroups	X					
Outlines					X	
Package Bodies					X	
Packages	X	X			X	
Plans		X				
Primary Keys	X	X	X	X	X	X
Procedures	X	X	X		X	X
Profiles					X	
Redo Log Groups					X	
Remote Servers			X			X
Roles			X		X	X
Rollback Segments					X	
Rules			X			X
Segments			X			X
Sequences	X	X			X	
Snapshots					X	
Snapshot Logs					X	
Stogroups		X				
Structured Types	X	X				
Summary Tables	X					
Synonyms		X			X	
Tables	X	X	X	X	X	X
Tablespaces	X	X			X	
Triggers	X	X	X		X	X
Type Bodies					X	
Types					X	
Unique Keys	X	X	X	X	X	X
User Datatypes	X	X	X			X
User Messages	X		X			X

Objects	IBM DB2 for Windows/Unix	IBM DB2 for OS/390 and z/OS	Microsoft SQL Server	MySQL	Oracle	Sybase ASE
Users	X	X	X	X	X	X
Views	X	X	X		X	X

Aliases

NOTE: This object is supported by IBM DB2 for Linux, Unix, and Windows, IBM DB2 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 user requires only temporary access to a database. You can also use an alias to mask a user's identity.

Datasource Explorer List

When you select the Aliases node of the Datasource Explorer, DBArtisan lists the aliases defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about aliases from the system catalog:

NOTE: Columns differ by platform.

Column Name	Description
Alias Owner	The identity of the alias' owner.
Alias Name	Name of the alias
Aliased Object Owner	User who owns the object that has been aliased
Aliased Object Name	Name of the object that has been aliased
Create Date	Date and time the alias was created

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Report	Transfer Ownership			

Bufferpools

NOTE: This object is supported by IBM DB2 for OS/390 and z/OS and IBM DB2 for Linux, Unix, and Windows only.

Bufferpools are areas of system memory to cache table and index data pages. DBArtisan retrieves table and index data pages from their physical storage location and places them in the bufferpool so that when you modify or access data, you can retrieve the data from memory rather than a physical device or file. Bufferpools improve database system performance because DBArtisan can retrieve the data quickly from memory.

Datasource Explorer List

When you select the Bufferpools node of the Datasource Explorer, DBArtisan lists the tablespaces defined in a database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about tablespaces from the system catalog:

IBM DB2 for Windows/Unix

Column Name	Description
Bufferpool Name	Name of the bufferpool
Node Group or Storage Group	Name of the nodegroups or storage groups that are assigned to the bufferpool
# Pages	The number of pages assigned to the bufferpool
Page Size	The size of the pages in the bufferpool
Extended Storage	Whether or not the bufferpool uses extended storage

IBM DB2 for OS/390 and z/OS

Column Name	Description
Buffer Pool Name	Buffer pool external name (BP0, BP1, ..., BP49, or BP32K, BP32K1, ...,BP32K9).
ID	Buffer pool internal identifier (0-49, 80-89, 100-109, 120-129)
# Use Count	Number of open table spaces or index spaces that reference this bufferpool. (Inactive pools have a zero use count.)
Size	User-specified virtual buffer pool size.
Type	Type of virtual buffer pool to be allocated. PRIMARY - Used to maintain the coherency of cached data. This structure is used for page registration and cross-invalidation. DATASPACE - A range of up to 2 gigabytes of contiguous virtual storage addresses that a program can directly manipulate. Unlike an address space, a data space can hold only data; it does not contain common areas, system data, or programs.
Hiper Pool Size	User-specified hiper pool size.
# VPSEQT	Sequential steal threshold for the virtual buffer pool.
# VPPSEQT	Parallel sequential threshold for the virtual buffer pool.
# VPXSEQT	Portion of the virtual buffer pool that is used for processing queries that originate on other members of the data sharing group.
# HPSEQT	Sequential steal threshold for the hiperpool.
# DWQT	Deferred write threshold for the virtual buffer pool. It is expressed as a percentage of the total virtual buffer pool size, and valid values range from 0 to 90. This threshold determines when deferred writes begin, based on the number of unavailable (non-stealable) buffers.
# VDWQT1	Vertical deferred write threshold for the virtual buffer pool expressed as a percentage of the total virtual buffer pool size, and valid values range from 0 to 90. This attribute determines when deferred writes begin, based on the number of updated pages for a given data set.
# VDWQT2	Vertical deferred write threshold for the virtual buffer pool expressed as an absolute number of buffers. VDWQT2 only applies when VDWQT1 is 0; VDWQT2 can range from0 to 9999.
CASTOUT	CASTOUT option for all hiperpools. The possible values are YES and NO.
PGSTEAL	Page stealing algorithm DB2 uses for the virtual buffer pool.LRU or FIFO are the possible options.

>

Available Functionality

DBArtisan offers the following functionality for this object:

NOTE: Functionalities differ by platform.

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Rename	Report			

Check Constraints

NOTE: This object is supported by all platforms.

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

Datasource Explorer List

When you select the Check Constraints node of the Datasource Explorer tree, DBArtisan lists the check constraints defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about check constraints from the system catalog:

Column Name	Description
Table Owner	Name of the table owner of the table with the check constraint
Table Name	Table with the check constraint
Constraint Name	Name of the check constraint

Available Functionality

DBArtisan offers the following functionality for this object:

NOTE: Functionalities differ by platform.

Drop	Extract	Migrate	Creating objects	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, 6 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.

Datasource Explorer List

When you select the Clusters node of the Datasource Explorer, DBArtisan lists the clusters defined in a database, and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about clusters from the system catalog:

Column Name	Description
Owner	Owner of the cluster.
Cluster Name	Name of the cluster
Tablespace	Tablespace on which the cluster is placed
Type	Whether the cluster is an index or a hash cluster
%Free	The percentage of space reserved on each data block for update operations
%Used	The minimum percentage of space that must be filled on each data block before another data block is allocated to the cluster
Key Size	Cluster key size in bytes

Available Functionality

DBArtisan offers the following functionality for this object:

Allocate Extent	Analyze	Deallocate Unused Space	Drop	Extract
Migrate	Creating objects	Modifying objects using editors	Report	Truncate

Control Files

NOTE: This object is supported by Oracle only.

Control files map to the structures that make up the database including the location of datafiles. Control files are read-only because they are only constructed during database creation.

Available Functionality

DBArtisan offers the following functionality for this object:

Backup Control Files	Open Control Files	Report		

Database Devices

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

Database Devices are the fundamental storage mechanism for Microsoft SQL Server. They bridge the physical world of your hard disks and the logical internal structure of Microsoft SQL Server. When you create a device, Microsoft SQL Server writes a file on your hard disk and pre-initializes the panels for quick allocation to databases. Microsoft SQL Server stores the resulting panel range in the system catalog and refers to it when determining how to allocate space on a device to a particular database.

Datasource Explorer List

When you select the Database Devices node located under the Storage node of the Datasource Explorer tree, DBArtisan lists the database devices and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about database devices from the system catalog:

Column Name	Description
Database Device Name	Name of the database device.
Size (MB)	Defined size of the database device.
Used (MB)	Space on the database device that has already been allocated to databases.
Free (MB)	Remaining space that has not been allocated on the database device.
Physical Name	Name and location of the *.dat file that stores the database device.

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Creating objects	Modifying objects using editors	Reinitialize
Disk Resize	Refit	Report		

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.

Datasource Explorer List

When you select the Database Links node of the Datasource Explorer, DBArtisan lists the database links defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about database links from the system catalog:

Column Name	Description
Owner	Owner of the database link

>

Column Name	Description
Database Link Name	Name of the database link
User	The user account for connecting to the remote database
Host	Instructions for connecting to the remote database
Created	The date and time that the database link was created

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Rename	Report			

Databases

DBArtisan supports the following DBMS databases:

- IBM DB2 for Linux, Unix, and Windows Instances
- Microsoft SQL Server Databases
- IBM DB2 for OS/390 and z/OS Instances
- 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.

For more information, see Available Functionality for Databases.

IBM DB2 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.

Datasource Explorer List

When you select the Instances node of the Datasource Explorer, DBArtisan lists the databases and displays pertinent information about each of them in the right pane. The tables below describes how this information is displayed:

IBM DB2 for Windows/Unix	
Column Name	**Description**
Name	Name of database
Creator	Name of creator

>

IBM DB2 for Windows/Unix	
Column Name	**Description**
Stogroup	Name of stogroup for the database, blank if it is a system database
Bufferpool	Name of default bufferpool for the database, blank if it is a system database
DB ID	Internal identification number for database
Timestamp	Indicates timestamp
Type	Indicates type of database; W indicates a work file
Group Member	The DB2 sharing member name
CreateEdts	Time when the create statement was created for the database
AlterEdts	Time when the most recent alter edit was applied
MIXED_CCSID	Indicates graphic characters
Index Bufferpool	Indicates name of the default bufferpool for indexes

For more information, see Available Functionality for Databases.

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.

NOTE: 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.

Datasource Explorer List

When you select the Databases node of the Datasource Explorer, DBArtisan lists the databases and displays pertinent information about each of them in the right pane. The tables below describes how this information is displayed:

Microsoft SQL Server	
Column Name	**Description**
Database Name	Name of the database
Database Status	Status of the database, such as online, offline, read-only or single-user
Create Date/Time	When the database was created

For more information, see Available Functionality for Databases.

IBM DB2 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.

Datasource Explorer List

When you select the Instances node of the Datasource Explorer, DBArtisan lists the databases and displays pertinent information about each of them in the right pane. The tables below describes how this information is displayed:

IBM DB2 for OS/390 and z/OS	
Column Name	**Description**
Name	Name of database
Creator	Name of creator
Stogroup	Name of stogroup for the database, blank if it is a system database
Bufferpool	Name of default bufferpool for the database, blank if it is a system database
DB ID	Internal identification number for database
Timestamp	Indicates timestamp
Type	Indicates type of database; W indicates a work file
Group Member	The DB2 sharing member name
CreateEdts	Time when the create statement was created for the database
AlterEdts	Time when the most recent alter edit was applied
MIXED_CCSID	Indicates graphic characters
Index Bufferpool	Indicates name of the default bufferpool for indexes

For more information, see Available Functionality for Databases.

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.

Datasource Explorer List

When you select the Databases node of the Datasource Explorer, DBArtisan lists the databases and displays pertinent information about each of them in the right pane. The tables below describes how this information is displayed:

>

Sybase ASE	
Column Name	**Description**
Database Name	Name of the database
Database Status	Status of the database, such as online or offline
Size	Size of the database in MB
Last Dump Tran	The last time that the transaction log for the database was dumped
Create Date/Time	When the database was created

For more information, see Available Functionality for Databases.

Available Functionality for Databases

DBArtisan offers the following functionality for databases:

NOTE: Functionalities differ by platform.

Database Backup Wizard	Checkpoint	Coalesce	DBCC	Detach/Attach
Drop	Extract	Migrate	Move Log	Creating objects
Modifying objects using editors	Quiesce Database	Rename	Report	Set Online/Offline
Shrink	Start Database	Stop Database	Unquiesce Database	Update Statistics

Data Caches

NOTE: This object is supported by Sybase only.

Data Caches are devices stored in memory. If configured properly, they can substantially boost performance by allowing users to read or write data in memory instead of to or from physical disk. With System 11, you can now “borrow” memory from the default data cache and allocate it to customized data caches.

Named data caches let you achieve substantially greater concurrency. Data caches also increase the likelihood that Sybase ASE reads data from memory and not disk. DBArtisan lets you bind particularly active databases (for example, tempdb), tables (for example, syslogs and sysindexes) or indexes to dedicated data caches. In addition, you can boost performance by creating memory pools, each configured for different I/O throughput (2 K to 16 K), within every named data cache.

In spite of their potential, named data caches are not a panacea for performance problems. You should still carefully optimize your queries and indexing strategy. Poor configuration of named data caches can lead to worse performance because precious memory resources are misallocated. To determine data cache performance, you should examine the cache hit ratio.

>

The Datasource Explorer List

When you select the Data Caches node located under Storage in the Datasource Explorer tree, DBArtisan lists the data caches and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about data caches from the system catalog:

Column Name	Description
Instance ID	Instance ID for Cluster Edition datasources only.
Data Cache Name	Name of the data cache
Type	Type of the data cache, whether it is for mixed use or log only
Status	Status of the data cache
Configured	Configured size of the data cache in bytes
Run	Run size of the data cache in bytes
Overhead	Overhead required to manage the data cache in bytes

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Creating objects	Modifying objects using editors	Report

DBRM

NOTE: This object is supported by IBM DB2 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.

The Datasource Explorer List

When you select the DBRM node of the Datasource Explorer, DBArtisan lists pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays DBRMs from the system catalog:

Column Name	Description
DBRM Name	Displays the name of the DBRM.
Version	Displays DBRM version.
Contoken	The contoken (consistency token) in the object compares to the contoken in the plan/package. If the contokens match, then the program will execute. If they do not match and autorebind = y, and the rebind is successful, then the program will execute. If they do not match and autorebind = n or fails, then an -818 is returned to the program - contoken mismatch, and the job ends. DB2 displays the mismatched contokens. The analysis associated with a -818 requires finding the DBRM or package contoken. This pane lets you easily expose the contokens and find the version mismatch.
Language	Displays the Host language.
Precompile Timestamp	Precompilation time of the DBRM.

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Creating objects	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.

Datasource Explorer List

When you select the Defaults node of the Datasource Explorer tree, DBArtisan lists the defaults defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about defaults from the system catalog:

Column Name	Description
Default Owner	User who owns the default
Default Name	Name of the default
Create/Date Time	When the default was created

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Rename	Report			

Directories

NOTE: This object is supported by Oracle only.

Directories create an alias to an external operating system directory to your database files, which can be used for storing large binary object files. When you create a directory, provide the full path name to the outside operating system where the BFILE's are stored. This object lets you store large files, such as video, outside of the database. The directory object lets you provide a simple alias to the full path name of an outside server's file system, which you can then use to point to the files when creating procedural logic objects. This saves the developer from having to type the full path name when coding.

To create a Directory object, you need CREATE ANY DIRECTORY system privileges. You also create or have in place a corresponding operating system directory to store the file. This directory must have the correct read permissions for Oracle processes.

>

Datasource Explorer List

When you select the Directories node of the Datasource Explorer, DBArtisan lists the directories defined in a database, and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about directories from the system catalog:

Column Name	Description
Name	Name of the directory
File Name	Physical path and name of the operating system file that corresponds to the directory

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Report				

Dump and Backup Devices

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

Dump devices store backups of databases and transaction logs. In Microsoft SQL Server, dump devices have been renamed to backup devices. Before you can dump a database or its transaction log, first create a dump device that stores the resulting backup file. Microsoft SQL Server and Sybase ASE lets you create tape and disk dump devices. Disk dump devices are a fast and convenient mechanism for backups. Tape dump devices facilitate off site storage for enhanced disaster recovery.

Datasource Explorer List

When you select the Dump Devices or Backup Devices node located under the Storage node of the Datasource Explorer tree, DBArtisan lists the devices and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about these devices from the system catalog:

Column Name	Description
Dump\Backup Device Name	Name of the dump or backup device.
Type	Type of the dump device, either disk or tape.
Physical Name	Name of the dump device recognized by the operating system.

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Creating objects	Modifying objects using editors	Report
View Contents				

Event Monitors

NOTE: This object is supported by IBM DB2 for Windows/Unix only.

Event Monitors lets you monitor specified events that occur on the database including deadlocks, connections, and statements. Because monitoring is configurable by type, you can record activity against all trackable objects or only those you consider most important.

NOTE: Event Monitors cannot be altered using the ALTER command. An Event Monitor has to be dropped and re-created.

Datasource Explorer List

When you select the Event Monitors node of the Datasource Explorer tree, DBArtisan lists the event monitors defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about event monitors from the system catalog:

Column Name	Description
Monitor Owner	User who owns the monitor
Monitor Name	Name of the monitor
Monitor Type	Type of the monitor
Autostarts	Indicates if event monitor autostarts
On Node	Indicates if event monitor is on node
Scope	Scope of the monitor

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Rename	Report	Transfer Ownership		

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.

Datasource Explorer List

When you select the Extended Stored Procedures node of the Datasource Explorer tree, DBArtisan lists the extended stored procedures defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about extended stored procedures from the system catalog:

Column Name	Description
Owner	User who owns the stored procedure
Procedure Name	Name of the Stored Procedure
DLL	The routine called by the procedure

>

Column Name	Description
Created	When the stored procedure was created

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Execute	Extract	Migrate	Creating objects
Modifying objects using editors	Rename	Report		

Foreign Keys

NOTE: This object is supported by all platforms.

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

DBArtisan offers the following functionality for this object:

NOTE: Functionality differs by platform.

Drop	Extract	Migrate	Creating objects	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 DBArtisan 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.

>

Function	Description
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 datasource function -Oracle, SQL Server, Sybase, etc. A function mapping needs to be created in conjunction with the template function.

Datasource Explorer List

When you select the Functions node of the Datasource Explorer, DBArtisan lists the functions defined in a database, and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about functions from the system catalog:

NOTE: Columns differ by platform.

Column Name	Description
Function Owner	Owner of the function
Function Name	Name of the function
Status	Whether the function is valid or invalid
Specific Name	Unique identity of the function
Origin	Indicates if the function is system-generated or user-defined
Language	The programming language which was used to create the function.
Implementation	The programming language which was used to create the function.
Created	The date and time that the function was created
Last Modified	Date that the function was last modified

Available Functionality

DBArtisan offers the following functionality for this object:

NOTE: Functionalities differ by platform.

Compile	Create Synonym	Drop	Extract	Creating objects
Modifying objects using editors	Report	Transfer Ownership		

Important Notes

With respect to functions and stored procedures, it can be necessary to have statements executed before and after creation of the procedure or function. This can be useful for example, if you need to create or drop temporary tables used by the function or procedure. DBArtisan supports the use of two tag pairs, ETStart and ETEnd, that let you embed statements in the first comment block of a stored procedure or function. The following shows the expected syntax:

```
create procedure dbo.procname(@a numeric) as
/*
<ETStart>SQL Statement</ETStart>
<ETEnd>SQL Statement</ETEnd>
*/
begin
...
```

>

Groups

NOTE: This object is supported by IBM DB2 for Windows/Unix, Oracle, and Sybase only.

Groups are a defined collection of database users. The primary use of groups is to consolidate the management of permissions. By batching together similar users into groups, you can greatly reduce the number of commands required to set permissions.

Every user automatically belongs to the public group. To assign a user to another group, add the user to that group. Then the user belongs to that group and public.

NOTE: A user can only belong to one group at a time other than public.

Datasource Explorer List

When you select the Groups node of the Datasource Explorer, DBArtisan lists the groups defined in a database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about the group:

Column Name	Description
Group Name	Name of the group

NOTE: In Oracle, there is only one group PUBLIC.

Available Functionality

DBArtisan offers the following functionality for this object:

NOTE: Functionalities differs by platform.

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Rename	Report			

Image Copies

NOTE: This object is supported by IBM DB2 for OS/390 and zOS.

Image Copies are a full image copy of all pages in a tablespace, partition, data set, or index space. Image Copies can also be an incremental image copy of only pages that have been modified since the last use of the Image Copy Wizard.

Available Functionality

DBArtisan offers the following functionality for this object:

- Recover
- Modify Recovery
- Copy to Copy
- MergeCopy

Indexes

NOTE: This object is supported by all platforms.

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 for Linux, Unix, and Windows Indexes

IBM DB2 for OS/390 and z/OS Indexes

Microsoft SQL Server Indexes

Oracle Indexes

Sybase ASE Indexes

Available Functionality for Indexes

IBM DB2 for Linux, Unix, and Windows Indexes

IBM DB2 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.

When you select the indexes node of the Datasource Explorer, DBArtisan lists the indexes defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about indexes from the system catalog:

Column Name	Description
Table Owner	User who owns the indexed table
Table Name	Name of the indexed table
Index Owner	Name of the index owner
Index Name	Name of the index
Type	Indicates if the index is primary, unique, or non-unique
#Levels	The number of levels of the index's B-Tree. If this value is greater than 3, then you should consider rebuilding the index.
#Columns	The number of table columns that are indexed.

For more information, see Available Functionality for Indexes.

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.

DBArtisan separates system indexes from user-defined indexes in the Datasource Explorer to ensure that system indexes are not accidentally altered or dropped.

Datasource Explorer List

When you select the indexes node of the Datasource Explorer, DBArtisan lists the indexes defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about indexes from the system catalog:

Microsoft SQL Server	
Column Name	**Description**
Table Owner	User who owns the indexed table
Table View/Name	Name of the indexed table or view
Index Name	Name of the index
On View	Indicates whether the indicates resides on a view.
Clustered	Whether the index is clustered or non-clustered
Unique	Whether or not the index is unique
File Group	Name of the file group on which the index is placed

For more information, see Available Functionality for Indexes.

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.

Datasource Explorer List

When you select the indexes node of the Datasource Explorer, DBArtisan lists the indexes defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about indexes from the system catalog:

Oracle	
Column Name	**Description**
Table Owner	User who owns the indexed table
Table Name	Name of the indexed table
Index Name	Name of the index
Status	Indicates if the index is valid or invalid
Unique	Indicates if the index is unique or non-unique
Levels	The number of levels of the index's B-Tree. If this value is greater than 3, then you should consider rebuilding the index.
Dist. Keys	The number of distinct values found in the table based on the index. A low value indicates that the index is not very specific and, therefore, could be unnecessary.

For more information, see Available Functionality for Indexes.

Index Partitions

Index partitions are similar to table partitions. There are three types of partitioned indexes that Oracle supports:

1 Local pefixed

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 for OS/390 and z/OS Indexes

IBM DB2 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.

Datasource Explorer List

When you select the indexes node of the Datasource Explorer, DBArtisan lists the indexes defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about indexes from the system catalog:

IBM DB2 for OS/390 and z/OS	
Column Name	**Description**
Index Name	Name of the index
Index Owner	Name of the index owner
Organization	Indicates the object reference
Table Name	Name of the indexed table
Table Owner	User who owns the indexed table
Database	Name of the database
Bufferpool	Name of the Bufferpool used for the index
Last RUNSTATS	Indicates the time of the last RUNSTAT
Levels	The number of levels of the index's B-Tree. If this value is greater than 3, then you should consider rebuilding the index.

For more information, see Available Functionality for Indexes.

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.

DBArtisan separates system indexes from user-defined indexes in the Datasource Explorer to ensure that system indexes are not accidentally altered or dropped.

Datasource Explorer List

When you select the indexes node of the Datasource Explorer tree, DBArtisan lists the indexes defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about indexes from the system catalog:

Column Name	Description
Table Owner	User who owns the indexed table
Table Name	Name of the indexed table
Index Name	Name of the index
Clustered	Whether the index is clustered or non-clustered
Unique	Whether or not the index is unique
Segment	The segment on which the index is placed

For more information, see Available Functionality for Indexes.

Available Functionality for Indexes

DBArtisan offers the following functionality for indexes:

NOTE: Functionality differs by platform.

Analyze	Allocate Extent	Check Index	Convert to Partitioned	Copy to Copy
Deallocate Unused Space	Disable Index	Drop	DBCC	Estimate Size
Extract	Copying Images	Migrate	Modify Recovery	Modify Statistics
Creating objects	Modifying objects using editors	Place	Rebuild Index	Rename
Reorganize	Report	Runstats Index	Transfer Ownership	Update Statistics

Instance

NOTE: This object is support for IBM DB2 and Oracle.

DBArtisan places Instance as the first level of information under the Datasource node in the Datasource Explorer. Instance includes:

- DB Manager Configuration
- datasources

>

Available Functionality

DBArtisan offers the following functionality for this object:

Instance Node	Database Wizard for IBM DB2 (IBM DB2 LUW)	Quiesce Database	Unquiesce Database	

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

DBArtisan offers the following functionality for this object:

Run	Creating objects	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.

Datasource Explorer List

When you select the libraries node of the Datasource Explorer, DBArtisan lists the libraries defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about indexes from the system catalog:

Column Name	Description
Owner	User who owns the library
Library Name	Name of the library
File Name	Name and location of the shared operating system library file
Dynamic	Whether the referenced operating system file supports dynamic linking
Status	Whether or not the library is valid. Libraries are not validated upon creation, but rather when they are called in a procedure or function.

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Rename	Report			

Linked Servers

NOTE: This object is supported by Microsoft SQL Server or later.

Linked servers let you create links to OLE DB datasources. With a linked server you can execute commands against OLE DB datasources on different servers. Using linked servers you can:

- Access remote servers
- Issue enterprise-wide distributed queries, updates, commands, and transactions.

Distributed Queries

Before you run a distributed query, write the following as the first statement of your SQL script:

```
SET ANSI_NULLS ON

GO

SET ANSI_WARNINGS ON

GO
```

Datasource Explorer List

When you select the Linked Servers node located under the Storage node of the Datasource Explorer, DBArtisan lists the linked servers and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about linked servers from the system catalog:

Column Name	Description
Linked Server Name	Name of the server.
Net Name	Reserved (name of the server).
Product Name	Product name for the linked server.
Ole DB Provider Name	Name of the OLE DB provider accessing the server.

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
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.

For more information, see the following topics:

- Microsoft SQL Server Logins
- Sybase ASE Logins

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.

Datasource Explorer List

When you select the Logins node located under the Storage node of the Datasource Explorer tree, DBArtisan lists the logins and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about logins from the system catalog:

Microsoft SQL Server	
Column Name	**Description**
Login Name	Name of the login.
Type	Type of login Standard, NT Group, or NT User.
Access	Indicates if the login has been granted access to the server or not.
Default Database	Database that is assigned as the default for the login.
Server Role	Shows whether or not a server role is assigned to the login.
Date/Time Created	When the login was created.

For more information, see Available Functionality for Logins.

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.

The Datasource Explorer List

When you select the Logins node located under the Storage node of the Datasource Explorer tree, DBArtisan lists the logins and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about logins from the system catalog:

Column Name	**Description**
Login Name	Name of the login
Default Database	Database that is assigned as the default for the login
CPU	Cumulative CPU usage by the login since statistics were last flushed
I/O	Cumulative I/O usage by the login. Since statistics were last flushed
Create Date/Time	When the login was created

Available Functionality for Logins

DBArtisan offers the following functionality for this object:

Change Password	Drop	Extract	Migrate	Creating objects

Modifying objects using editors	Report			

Materialized Query Tables

NOTE: This object is supported by IBM DB2 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.

Datasource Explorer List

When you select the Materialized Query Tables node of the Datasource Explorer, DBArtisan lists the materialized query tables defined in a database and displays pertinent information about each of them in the right pane.

The table below describes the information DBArtisan derives and displays about materialized query tables from the system catalog:

Column Name	Description
Table Owner	User who owns the table
Table Name	Name of the table
# Rows	Number of rows of data the table includes
% Free	The amount of remaining space allocated to the table
Tablespace	The name of the tablespace on which the table resides
Index TS	The name of the tablespace on which the table's indexes reside
Long TS	The name of the tablespace on which the long objects (LOB,CLOB) stored on the table reside
Refresh Mode	Either Deferred or Immediate
Refresh Time	Last time table was refreshed
Create Time	When the table was created

Available Functionality

DBArtisan offers the following functionality for this object:

Create Insert Statements	Create Synonym	Drop	Edit Data	Extract
Migrate	Creating objects	Modifying objects using editors	Refresh 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.

Datasource Explorer List

When you select the Materialized Views node of the Datasource Explorer, DBArtisan lists the materialized views defined in a database and displays pertinent information about each of them in the right pane.

The table below describes the information DBArtisan derives and displays about materialized views from the system catalog:

Column Name	Description
Materialized View Owner	User who owns the materialized view
Materialized View Name	Name of the materialized view
Last Refresh	Date the materialized view was last refreshed.
Refresh Type	The method by which the materialized view refreshes itself (fast, force or complete).

Available Functionality

DBArtisan offers the following functionality for this object:

Compile	Create Synonym	Drop	Extract	Migrate
Creating objects	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.

Datasource Explorer List

When you select the Materialized View Logs node of the Datasource Explorer, DBArtisan lists the materialized views logs defined in a database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about materialized view logs from the system catalog:

Column Name	Description
Materialized View Log Owner	User who owns the materialized view log

>

Column Name	Description
Materialized View Log Name	Name of the materialized view log
Log Table	A system-generated table that tracks changes to data in the master table
Log Trigger	A system-generated trigger that maintains data in the Log Table

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Report				

Nodegroups

NOTE: This object is supported by IBM DB2 for Windows/Unix only.

Nodegroups are a named subset of one or more partitions in a database. DBArtisan places Nodegroup under the Storage node in the Datasource Explorer.

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Rebind Packages	Redistribute	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.

Datasource Explorer List

When you select the Oracle Job Queues node of the Datasource Explorer, DBArtisan lists the job queues defined in the database and displays pertinent information about each of them in the right pane. DBArtisan displays the following columns:

- Logon User
- Job ID
- Job
- Running
- Start Date/Time
- Next Date/Time

>

- Failures
- Broken

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	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.

Datasource Explorer List

When you select the Outlines node of the Datasource Explorer, DBArtisan lists the outlines defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about stored outlines from the system catalog:

Column Name	Description
Owner	User who owns the stored outline.
Name	Name of the stored outline.
Category	Indicates the category of the stored outline.
Used	Indicates if the outline is used or unused.
Timestamp	Date and time the outline was created.
Version	Indicates the version of Oracle for the stored outline.

Available Functionality

DBArtisan offers the following functionality for this object:

Change Category	Drop	Extract	Migrate	Creating objects
Modifying objects using editors	Reassign by Category	Rebuild Outlines	Rename	Report

Packages

NOTE: This object is supported by IBM DB2 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.

NOTE: DBArtisan splits out package headers and package bodies in the Datasource Explorer; however, you create both a package header and body from the packages node.

Datasource Explorer List

When you select the Packages node of the Datasource Explorer, DBArtisan lists the packages defined in a database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about packages from the system catalog:

NOTE: Columns differ by platform.

Column Name	Description
Owner	User who owns the package
Package Name	Name of the package
Collection	Server generated identification
Collection ID	Server generated identification
Contoken	The consistency token in the object is compared to the contoken in the package. This pane lets you easily expose the contokens and find the version mismatch.
Creator	The user who generated the package
Qualifier	Name of package creator
Status	Whether the package is valid or invalid
Explain Snapshot	Whether or not the snapshot information should be gathered for each static SQL statement in the calling application
Binder	The user who generated the package

>

Column Name	Description
Sections	The number of sections that comprise the package
Last Bind Date	The date that the package was last bound to an external program
Valid	Whether the package is valid or invalid
Operative	Whether the package is operative or not
Created	Date and time the package was created
Last Modified	Date and time the package was last modified

Available Functionality

DBArtisan offers the following functionality for this object:

NOTE: Functionalities differ by platform.

Bind Package	Compile	Create Synonym	Drop	Extract
Flush Cache	Migrate	Creating objects	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 more information, see Package Bodies Editor (Oracle).

Available Functionality

DBArtisan offers the following functionality for this object:

Compile	Create Synonym	Drop	Extract	Migrate
Creating objects	Modifying objects using editors	Report		

Plans

NOTE: This object is supported by IBM DB2 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.

>

The Datasource Explorer List

You can select the Plan node in the Datasource Explorer and view the corresponding statistics. The table below describes the information DBArtisan derives and displays about plans from the system catalog:

Column Name	Description
Plan Name	Name of plan
Owner/Creator	User name of plan owner or creator
Qualifier	Name of qualifier
Valid	Indicates if plan is valid or invalid
Operative	Indicates if plan is operative
Isolation	Level of isolation
Reoptvar	
Packlists	Number of packlists
Environments	Number of environments
Bind Date/Timestamp	Time stamp of bind

Available Functionality

DBArtisan offers the following functionality for this object:

Bind Plan	Copy Object Names	Drop	Free Plans	Modifying objects using editors
Rebind Plans	Report			

Primary Keys

NOTE: This object is supported by all platforms.

Primary Keys are a set of table columns that can uniquely identify every row of a table.

Available Functionality

DBArtisan offers the following functionality for this object:

NOTE: Functionalities differ by platform.

Disable Index	Drop	Extract	Migrate	Creating objects
Modifying objects using editors	Rebuild Index (SQL Server)	Rename	Reorganize	Report
Status	Transfer Ownership			

Procedures

NOTE: This object is supported by all platforms.

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 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.

NOTE: DBArtisan also supports extended procedures. For details, see Extended Procedures.

Datasource Explorer List

When you select the Procedures node of the Datasource Explorer, DBArtisan lists procedures defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about procedures from the system catalog:

NOTE: Columns differ by platform.

Column Name	Description
Owner	User who owns the procedure
Procedure Name	Name of the procedure
Procedure Owner	User who owns the procedure
Schema	Indicates the schema of the procedure
Owner	User who owns the procedure
External Name	The load module name for the procedure
Language	The programming language in which the procedure body was written
WLM Environment	Indicates the WLM application environment
Create Time	Indicates the time the procedure was created
Specific Name	The name that uniquely identifies the procedure
Tran Mode	Whether the procedure is chained or unchained
Origin	Indicates if the procedure is was created externally or internally
Language	The programming language in which the procedure body was written
Implementation	How the procedure is implemented or called
Status	Indicates if the procedure is valid or invalid
Created	Date and time the procedure was created
Last Modified	Date and time the procedure was last modified

>

Available Functionality

DBArtisan offers the following functionality for this object:

NOTE: Functionalities differ by platform.

Compile	Create Synonym	Drop	Extract	Creating objects
Modifying objects using editors	Rename	Report	Transfer Ownership	

Important Notes

With respect to functions and stored procedures, it can be necessary to have statements executed before and after creation of the procedure or function. This can be useful for example, if you need to create or drop temporary tables used by the function or procedure. DBArtisan supports the use of two tag pairs, ETStart and ETEnd, that let you embed statements in the first comment block of a stored procedure or function. The following shows the expected syntax:

```
create procedure dbo.procname(@a numeric) as
/*
<ETStart>SQL Statement</ETStart>
<ETEnd>SQL Statement</ETEnd>
*/
begin
...
```

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.

Datasource Explorer List

When you select the Profiles node of the Datasource Explorer, DBArtisan lists the profiles defined in a database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about profiles from the system catalog:

Column Name	Description
Profile Name	Name of the profile

>

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Report				

Redo Log Groups

NOTE: This object is supported by Oracle only.

Redo Logs save information on the state of the database after an update. If the database fails before the update is written to the datafiles, the update can be re-done, bringing the changes back after a crash, so it appears the update did happen.

TIP: Have sufficient redo logs available so that other Oracle threads do not have to wait until one is free before performing an update.

Redo logs are arranged in groups, and when a particular group is being written to, each member receives the same data.

TIP: To avoid crashes, have two logfiles in a group, each one being on a different (physical) disk.

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Report				

Remote Servers

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

Remote Servers are servers on a network that you can access through your local server. Instead of logging directly into a remote server, you can log into your local server and execute a remote procedure call to it.

The local server routes all commands to the remote server. To communicate with a remote server, you log into a local server and submit your remote procedure calls for remote servers. The remote server receives remote procedure calls from the local server and processes them accordingly.

Datasource Explorer List

When you select the Remote Servers node located under the Storage node of the Datasource Explorer tree, DBArtisan lists the remote servers and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about remote servers from the system catalog:

Microsoft SQL Server	
Column Name	**Description**
Remote Server Name	Name of the server

>

Microsoft SQL Server	
Column Name	**Description**
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.

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Report				

Roles

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. For more information, see the following topics:

- Microsoft SQL Server Roles
- Oracle Roles
- Sybase ASE Roles
- Available Functionality for Roles.

Microsoft SQL Server Roles

Roles let you manage permissions for users to perform a set of activities based on job functions. Roles are a mechanism for streamlining the assignment of object and system privileges to multiple users. Instead of granting sets of privileges to individual users, you can create a role, grant system and object privileges to that role, and then grant that role to all the users who should share the same privileges. You can grant one or more roles to the same user. Therefore, if you need to change the permissions for a certain role you can do so, and not have to grant or revoke the permissions for each individual user.

>

Datasource Explorer List

When you select the Roles node of the Datasource Explorer tree, DBArtisan lists the roles defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about roles from the system catalog:

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

For more information, see Available Functionality for Roles.

Oracle Roles

Roles are a mechanism for streamlining the assignment of object and system privileges to multiple users. Instead of granting sets of privileges to individual users, you can create a role, grant system and object privileges to that role, then simply grant that role to all the users who should share the same privileges. You can grant one or more roles to the same user.

The Datasource Explorer List

When you select the Roles node of the Datasource Explorer, DBArtisan lists the roles defined in the database, displaying pertinent information about each of them in the right pane. DBArtisan derives and displays the following information about roles from the system catalog:

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).

For more information, see Available Functionality for Roles.

Sybase ASE Roles

Roles are a mechanism for streamlining the assignment of object and system privileges to multiple users. Instead of granting sets of privileges to individual users, you can create a role, grant system and object privileges to that role, then simply grant that role to all the users who should share the same privileges. You can grant one or more roles to the same user.

The Datasource Explorer List

When you select the Roles node of the Datasource Explorer, DBArtisan lists the roles defined in the database, displaying pertinent information about each of them in the right pane. DBArtisan derives and displays the following information about roles from the system catalog:

Column Name	Description
Role Name	Name of the role

>

For more information, see Available Functionality for Roles.

Available Functionality for Roles

DBArtisan offers the following functionality for roles:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Report				

Rollback Segments

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.

Datasource Explorer List

When you select the Rollback Segments node of the Datasource Explorer, DBArtisan lists the rollback segments defined in a database, displaying pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about rollback segments from the system catalog:

Column Name	Description
Rollback Segment Name Owner	Name of the Rollback Segment
Tablespaces	Name of the tablespace where the rollback segment is placed

Column Name	Description
Status	Indicates if the rollback segment is on-line or off-line
Gets	Number of rollback segment header requests. *It is an indicator of activity levels on the rollback segment.
Waits	Number of rollback segment header requests that resulted in waits. It is an indicator of concurrency. *If you see a significant number of waits, then you should consider adding rollback segments to spread concurrency.
Extents	A logical unit of database storage space allocation that comprises a number of contiguous data blocks. Oracle allocates a new extent for the segment when the existing is completely used.
Size	The space in bytes allocated for the rollback segment
Bytes Written	Current datafile size in bytes
Current Xacts	
Shrinks	The cumulative number of times Oracle has truncated extents from the rollback segment
Optimal Size	The OPTIMAL size in bytes of the rollback segment

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	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.

Datasource Explorer List

When you select the Rules node of the Datasource Explorer tree, DBArtisan lists the rules defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about rules from the system catalog.

Column Name	Description
Owner	User who owns the rule.
Rule Name	Name of the rule.
Created	When the rule was created

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Rename	Report			

Schema

A Schema is a container of objects that can be owned by any user.

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

Datasource Explorer List

When you select the Schemas node of the Datasource Explorer tree, DBArtisan lists the schemas defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about schemas.

Column Name	Description
Name	Name of the schema.
Owner	User who owns the schema.
Created	When the schema was created

Available Functionality

DBArtisan offers the following functionality for this object:

Copy Object Names	Drop	Extract	Migrate	Creating objects
Modifying objects using editors	Report	Transfer Ownership		

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

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	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.

Datasource Explorer List

When you select the Sequences node of the Datasource Explorer, DBArtisan lists the sequences defined in a database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about sequences from the system catalog:

Column Name	Description
Sequence Owner	User who owns the sequence
Sequence Name	Name of the Sequence
Last Number	Last value generated by the sequence
Max Value	Maximum allowable value of the sequence. Unless the sequence is allowed to cycle, it stops generating values once it reaches the maximum value.
Created	Date and time the sequence was created

Available Functionality

DBArtisan offers the following functionality for this object:

Create Alias	Create Synonym	Drop	Extract	Migrate
Creating objects	Modifying objects using editors	Rename	Report	Restart
Transfer Ownership				

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

DBArtisan offers the following functionality for this object:

Create Synonym	Drop	Extract	Migrate	Creating objects

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

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Report				

Stogroups

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

Stogroups are storage groups, named sets of volumes on which DB2 for OS/390 data is stored.

When you select the Stogroups node of the Datasource Explorer, DBArtisan lists pertinent information about each of them in the right pane.

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Creating objects	Modifying objects using editors	Report

Structured Types

NOTE: This object is supported by IBM DB2 for Linux, Unix, and Windows and IBM DB2 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.

>

Datasource Explorer List

When you select the Types node of the Datasource Explorer, DBArtisan lists the types defined in the database, displaying pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about types from the system catalog.

Column Name	Description
Type Owner	User who owns the structured type
Name	Name of the structured type
Status	Whether the structured type is valid or invalid
Created	Date and time the structured type was created
Last Modified	Date and time the structured type was last modified

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Creating objects	Report	Transfer Ownership

Summary Tables

NOTE: This object is supported by IBM DB2 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.

Datasource Explorer List

When you select the SummaryTables node of the Datasource Explorer, DBArtisan lists the summary tables defined in a database and displays pertinent information about each of them in the right pane.

The table below describes the information DBArtisan derives and displays about summary tables from the system catalog:

Column Name	Description
Table Owner	User who owns the table
Table Name	Name of the table
# Rows	Number of rows of data the table includes
% Free	The amount of remaining space allocated to the table
Tablespace	The name of the tablespace on which the table resides
Index TS	The name of the tablespace on which the table's indexes reside
Long TS	The name of the tablespace on which the long objects (LOB,CLOB) stored on the table reside
Refresh Mode	Either Deferred or Immediate
Refresh Time	Last time table was refreshed
Create Time	When the table was created

>

Available Functionality

DBArtisan offers the following functionality for this object:

Compile	Create Synonym	Drop	Extract	Migrate
Creating objects	Modifying objects using editors	Rename	Report	

Synonyms

NOTE: This object is supported by IBM DB2 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.

Datasource Explorer List

When you select the Synonyms node of the Datasource Explorer, DBArtisan lists the synonyms defined in a database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about synonyms from the system catalog:

Column Name	Description
Synonym Owner	User who owns the synonym
Synonym Name	Name of the synonym
Base Object	Base object referenced by the synonym
Created	Date and time the synonym was created
Creator	Name of the owner of the synonym
Created By	Name of the user who created the synonym
Object Name	Indicates the name of the object
Object Creator	Name of the owner of the object

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Rename	Report			

Tables

NOTE: This object is supported by all platforms.

Tables are a 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: DBArtisan separates system tables from user-defined tables in the Datasource Explorer to ensure that system tables are not accidentally altered or dropped.

Datasource Explorer List

When you select the Tables node of the Datasource Explorer, DBArtisan lists the tables defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about tables from the system catalog:

IBM DB2 for Windows/Unix	
Column Name	**Description**
Table Owner	User who owns the table
Table Name	Name of the table
# Rows	Number of rows of data the table includes
% Free	The amount of remaining space allocated to the table
Tablespace	The name of the tablespace on which the table resides
Index TS	The name of the tablespace on which the table's indexes reside
Long TS	The name of the tablespace on which the long objects (LOB,CLOB) stored on the table reside
Type	Object type
Created	When the table was created
Last Runstats	When the script was last run

Microsoft SQL Server	
Column Name	**Description**
Table Owner	User who owns the table
Table Name	Name of the table
# Rows	Number of rows of data the table includes
File Group	The file group on which the table is placed
Created	When the table was created

Oracle	
Column Name	**Description**
Table Owner	User who owns the table
Table Name	Name of the table
# Rows	Number of rows of data the table includes
% Free	The percentage of space reserved on each data block for update operations
% Used	The minimum percentage of space that must be filled on each data block before another data block is allocated to the table
Tablespace	Name of the tablespace on which the table resides
Created	When the table was created
Last Modified	**Oracle 7 or later** When the table was last changed
Last Analyzed	**Oracle 8 or later** When the table was last analyzed

IBM DB2 for OS/390 and z/OS	
Column Name	**Description**
Table Name	Name of the table
Schema	Schema of the table
# Rows	Number of rows of data the table includes
Database Name	Indicates the name of the database on which the table resides
Tablespace	The name of the tablespace on which the table resides
Parents	Number of parents
Children	Number of children
Type	The type of the object
Created	When the table was created
Last Modified	When the table was last changed
Last Runstats	Indicates the time of the last runstat

Sybase ASE	
Column Name	**Description**
Table Owner	User who owns the table
Table Name	Name of the Table
# Rows	Number of rows of data the table includes
Table Segment	The segment on which the table is placed
Created	When the table was created

Available Functionality

DBArtisan 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
DBCC	Deallocate Unused Space	Drop	Drop Clone
Edit Data	Estimate Size	Exchange Data With Clone	Extract
Indexes	Lock	Migrate	Creating objects
Modifying objects using editors	Place	Quiesce Tablespaces	Recompile
Rename	Reorganize	Report	Schema
Select * From	Set Integrity	Transfer Ownership	Triggers
Truncate	Update Statistics		

Tablespaces

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

Datasource Explorer List

When you select the Tablespaces node on the Datasource Explorer, DBArtisan lists the tablespaces defined in the database, displaying pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about triggers from the system catalog:

NOTE: Columns differ by platform.

Column Name	Description
Tablespace Name	Name of the tablespace
Status	Status of the tablespace
Extent Management	Indicates LOCAL or DICTIONARY extent management type.
Segment Space Management	Indicates AUTO or MANUAL segment space management type.
Size	Total size (MB) and data blocks allocated to a tablespace.
Used	Total space used and used data blocks in a tablespace
Free	Total unused space and unused data blocks in a tablespace
Percent Free	Minimum percentage of free space in a block
Data	The type of data stored on the table space
Type	Whether the table space is system manage or database managed
Extent	The number of extents assigned to the tablespace. Extents are issued in 4K pages.
Prefetch	The number of pages to be fetched prior to being referenced
Overhead	The estimated time of the waiting period to initialize a new i/o operation
Transfer	The estimated time required to read one extent
Database	Name of the database selected for the table space
Creator	Indicates who created the tablespace
# Partitions	Indicates the number of partitioned table spaces
Layout	Type of table space layout

>

Column Name	Description
Type	Type of table space
Last STOSPACE	Indicates the last STOSPACE utility ran
Space(KB)	Indicates the space allocated to the table space
# Tables	Number of tables defined in the table space
Bufferpool	Name of the bufferpool assigned to the tablespace

Available Functionality

DBArtisan offers the following functionality for this object:

NOTE: Functionalities differ by platform.

Change Status	Check Data	Check Index	Coalesce	Copy to Copy
Drop	Extract	Copying Images	Migrate	Modify Statistics
Creating objects	Modifying objects using editors	Quiesce Tablespaces	Rebuild Index (DB2 z/OS)	Rename
Reorganize	Report	Runstats Index	Runstats Tablespace	Set Default
Start Database	Stop Database	Switch Online	Transfer Ownership	

Triggers

NOTE: This object is supported by all platforms.

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.

Datasource Explorer List

When you select the Triggers node on the Datasource Explorer, DBArtisan lists the triggers defined in the database, displaying pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about triggers from the system catalog:

NOTE: Columns differ by platform.

Column Name	Description
Table Owner	User who owns the target table

Column Name	Description
Trigger Owner	User who owns trigger
Trigger Name	Name of the trigger
Event	Data modification event(s) that causes the trigger to fire
Created	When the trigger was created.
Type	Whether the trigger fires once for an entire operation (STATEMENT) or for each row of data (ROW)
Enabled/Disabled	Whether the Trigger is enabled or disabled
Valid/Invalid	Whether the Trigger is valid or invalid
Status	Indicates if the trigger is valid or invalid
Time	Indicates if the trigger fires before or after a data modification
Schema	Name of the schema
Created By	Indicates the user who created the table
Table Name	Name of the table to which the trigger belongs
Granularity	Indicates whether the granularity is for each row or for each statement
Text	Indicates the trigger text
Create Time	Indicates the trigger time, either after or before
Remarks	Comments pertaining to the trigger

Available Functionality

DBArtisan offers the following functionality for this object:

NOTE: Functionalities differ by platform.

Compile	Drop	Extract	Migrate	Creating objects
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.

>

Datasource Explorer List

When you select the Types node of the Datasource Explorer, DBArtisan lists the types defined in the database, displaying pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about types from the system catalog.

Column Name	Description
Type Owner	User who owns the object type
Name	Name of the object type
Status	Whether the object type is valid or invalid
Created	Date and time the object type was created
Last Modified	Date and time the object type was last modified

Available Functionality

DBArtisan offers the following functionality for this object:

Compile	Drop	Extract	Creating objects	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.

Available Functionality

DBArtisan offers the following functionality for this object:

Compile	Drop	Extract	Migrate	Creating objects
Modifying objects using editors	Report			

Unique Keys

NOTE: This object is supported by all platforms.

Unique keys can enforce logical keys that are not chosen as the primary key. They enforce uniqueness for specified columns in a table.

Datasource Explorer List

When you select the Unique Keys node of the Datasource Explorer, DBArtisan lists the unique keys defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about unique keys from the system catalog:

Column Name	Description
Unique Index Name	Name of the unique index
Index Owner	Name of the index owner
Table Name	Name of the indexed table
Table Owner	User who owns the indexed table
Database	Name of the database
Bufferpool	Name of the Bufferpool used for the index
Last RUNSTATS	Indicates the time of the last RUNSTAT
Levels	The number of levels of the index's B-Tree. If this value is greater than 3, then you should consider rebuilding the index.

Available Functionality

DBArtisan offers the following functionality for this object:

NOTE: Functionalities differ by platform.

Disable Index	Drop	Extract	Migrate	Creating objects
Modifying objects using editors	Rebuild Index	Rename	Reorganize	Report
Status	Transfer Ownership			

User Datatypes

NOTE: This object is supported by IBM DB2 for Linux, Unix, and Windows, IBM DB2 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.

Datasource Explorer List

When you select the User Datatypes node of the Datasource Explorer, DBArtisan lists the user datatypes defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about user datatypes from the system catalog:

NOTE: Columns differ by platform.

>

Column Name	Description
Datatype Owner	User who owns the datatype
Datatype Name	Name of the user datatype
Base Type	Base datatype parameters for the user datatype
Bound Rule	The rule, if any, bound to the user datatype
Bound Default	The default, if any, bound to the user datatype
Length	The length of the datatype
Remark	Any comments that pertain to the datatype
Schema	Schema of the datatype
Created By	Who created the user datatype
Source Schema	Schema of the source data type
Source Type	Source datatype parameters for the user datatype
Datatype Owner	User who owns the datatype
Length	The length of the datatype
Remarks	Any comments that pertain to the datatype

Available Functionality

DBArtisan offers the following functionality for this object:

NOTE: Functionalities differ by platform.

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Rename	Report	Transfer Ownership		

User Messages

NOTE: This object is supported by IBM DB2 for Windows/Unix, Microsoft SQL Server, and Sybase only.

User Messages lets you catalog error messages that your database applications can re-use. Microsoft SQL Server stores your error messages in a system table, sysmessages. To return error messages from stored procedures and triggers, you need to call a system stored procedure and pass an error number as a parameter.

NOTE: User Messages can only be accessed from the master database for a specific server.

Datasource Explorer List

When you select the User Messages node located under the master database of the Datasource Explorer tree, DBArtisan lists the user messages defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about user messages from the system catalog:

Column Name	Description
Msg#	Number assigned to the message

Column Name	Description
Language	Language in which the user message is written
Severity	Level of severity assigned to the user message
Description	Error text description associated with the user message

Available Functionality

DBArtisan offers the following functionality for this object:

Drop	Extract	Migrate	Creating objects	Modifying objects using editors
Report				

Users

NOTE: This object is supported by all platforms.

A user is an individual with access to the DBMS.

For more information, see:

- IBM DB2 for Linux, Unix, and Windows Users
- IBM DB2 for OS/390 and z/OS Users
- Microsoft SQL Server Users
- Users/Privileges
- Oracle Users
- Sybase ASE Users
- Available Functionality for Users

IBM DB2 for Linux, Unix, and Windows Users

IBM DB2 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.

NOTE: DBArtisan lets you set a default SQLID user for a DB2 datasource. This lets the user set the implicit schema for the unqualified object references to a schema different from the user's login id.

Datasource Explorer List

When you select the Users node of the Datasource Explorer, DBArtisan lists the users defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about users from the system catalog:

Column Name	Description
User Name	Name of the user
Connect	Indicates if the user has connect privileges
Create Table	Indicates if the user has create table privileges

Column Name	Description
Create package	Indicates if the user has create package privileges
Register Function	Indicates if the user can register functions
Admin	Indicates if the user has admin privileges
Create Schemas Implicitly	Indicates if the user can create objects on which schemas are created implicitly

For more information, see Available Functionality for Users.

IBM DB2 for OS/390 and z/OS Users

IBM DB2 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.

NOTE: DBArtisan lets you set a default SCHEMA user for a DB2 datasource. This lets the user set the implicit schema for the unqualified object references to a schema different from the user's login id.

Datasource Explorer List

When you select the Users node of the Datasource Explorer, DBArtisan lists the users defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about users from the system catalog:

IBM DB2 for OS/390 and z/OS	
Column Name	**Description**
User Name	Name of the user
System	Aggregate number of system privileges
Database	Indicates the aggregate number of privileges for the database
Plan	Indicates the aggregate number of privileges for the plan
Packages	Indicates the aggregate number of privileges for the packages
Table	Indicates the aggregate number of table privileges
Column	Indicates the aggregate number or column privileges
Schema	Indicates the aggregate number of privileges for the schema

For more information, see Available Functionality for Users.

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.

>

Datasource Explorer List

When you select the Users node of the Datasource Explorer tree, DBArtisan lists the users defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about users from the system catalog:

Microsoft SQL Server	
Column Name	**Description**
User Name	Name of the User
Login Name	Name of the login that maps to the user
Created	When the user was created
Updated	When the User was last updated

For more information, see Available Functionality for Users.

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

Datasource Explorer List

When you select the Users node of the Datasource Explorer, DBArtisan lists the users defined in the database, displaying pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about users from the system catalog.

Column Name	**Description**
User Name	Name of the User
Default Tablespace	Name of the default tablespace assigned to the user
Temp Tablespace	Name of the temporary tablespace the user is placed on
Profile	Profile assigned to the user to govern resource utilization
Create Date	Date and time the user was created

For more information, see Available Functionality for Users.

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.

The Datasource Explorer List

When you select the Users node of the Datasource Explorer tree, DBArtisan lists the users defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about users from the system catalog.

Column Name	Description
User Name	Name of the user
Login Name	Name of the login that maps to the user
Group Name	Name of the group to which the user belongs

For more information, see Available Functionality for Users.

Available Functionality for Users

DBArtisan offers the following functionality for users:

NOTE: Functionalities differ by platform.

Change Password	Drop	Extract	Migrate	Creating objects
Modifying objects using editors	Report			

Views

NOTE: This object is supported by all platforms.

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.

Datasource Explorer List

When you select the Views node of the Datasource Explorer, DBArtisan lists the views defined in the database and displays pertinent information about each of them in the right pane. The table below describes the information DBArtisan derives and displays about views from the system catalog:

NOTE: Columns differ by platform.

Column Name	Description
View Owner	User who owns the view
View Name	Name of the view
Created	Date and time the view was created

>

Column Name	Description
Owner	User who owns the view
Created By	Who created the view
Status	Whether the view is valid or invalid
Created	Date and time the view was created
Last Modified	Date and time the view was last modified

Available Functionality

DBArtisan offers the following functionality for this object:

NOTE: Functionalities differ by platform.

Build Query	Compile	Copy Object Names	Create Alias	Create Synonym
Drop	Extract	Creating objects	Modifying objects using editors	Rename
Report	Select * From	Transfer Ownership	Update Statistics	

>

Creating objects

DBArtisan offers easy-to-use wizards and dialog boxes for creating objects. 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

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:

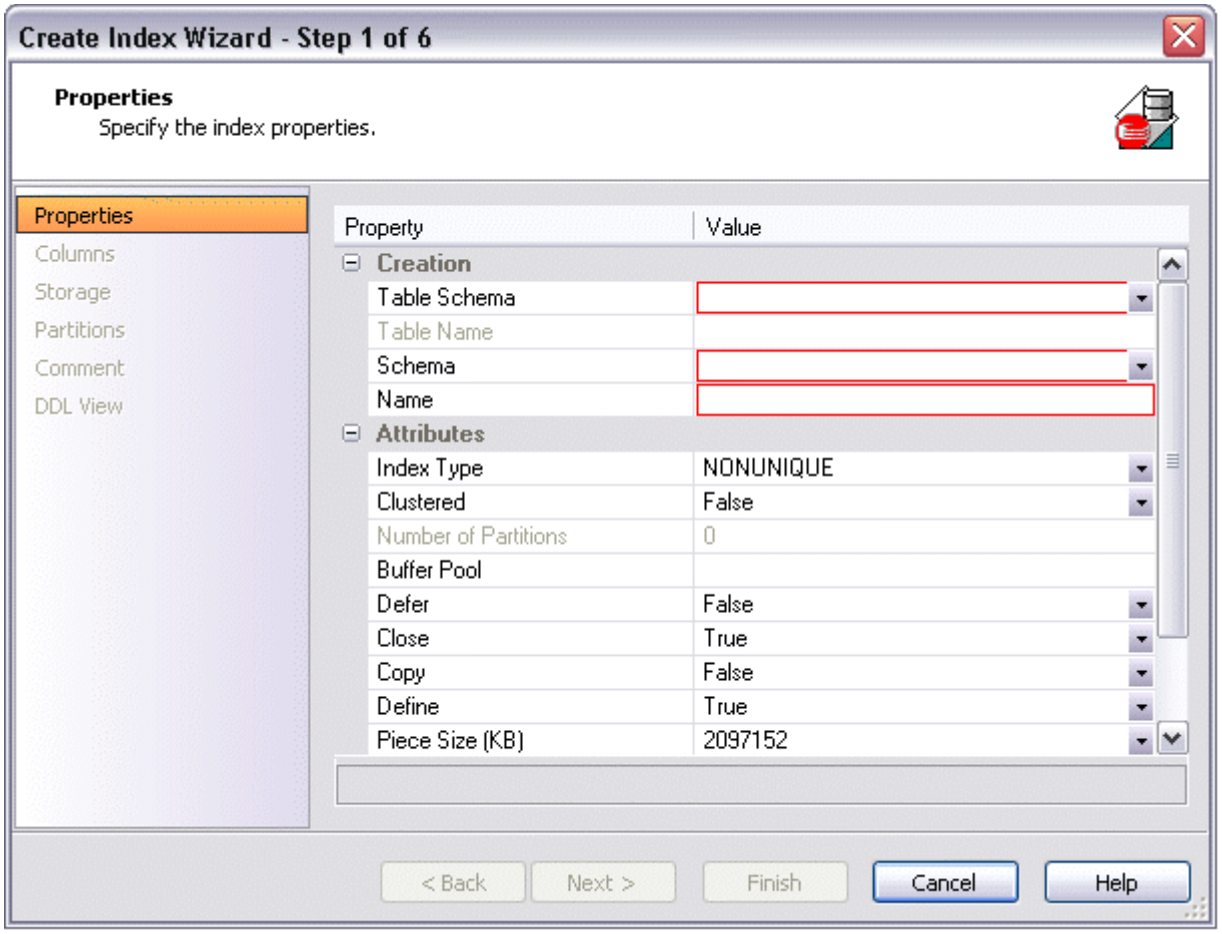

In order to work with object wizards, you must know how to:

- 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 DBArtisan 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.

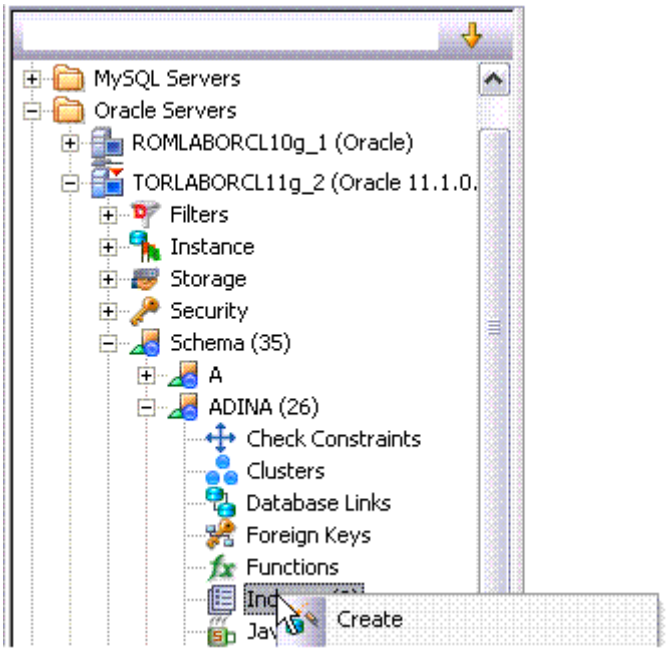

DBArtisan opens the object wizard for that object type.

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.

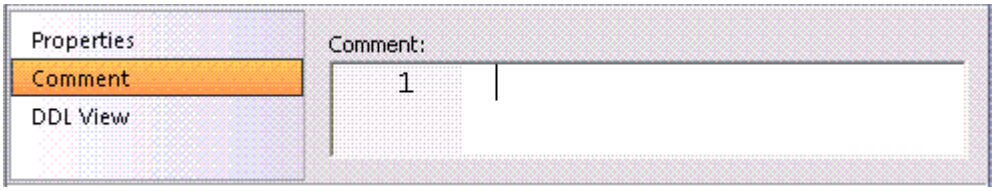

3. In the **Comment** area, type an explanatory note of up to 254 characters long.

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

DBArtisan lets you create DB2 LUW objects using the following wizards:

- Alias Wizard for IBM DB2 (IBM DB2 LUW)
- Bufferpool Wizard for IBM DB2 (IBM DB2 LUW)
- Database Wizard for IBM DB2 (IBM DB2 LUW)
- Event Monitor Wizard for IBM DB2 for Linux, Unix, and Windows
- Foreign Key Wizard for IBM DB2 (IBM DB2 LUW)
- Function Wizard for IBM DB2 (IBM DB2 LUW)
- Group Wizard for IBM DB2 for Windows/Unix
- Index Wizard for IBM DB2 (IBM DB2 LUW)
- Materialized Query Table Wizard for IBM DB2 (IBM DB2 LUW)
- Nodegroup Wizard for IBM DB2 (IBM DB2 LUW)
- Primary Key Wizard for IBM DB2 (IBM DB2 LUW)
- Procedure Wizard for IBM DB2 (IBM DB2 LUW)
- Schema Wizard for IBM DB2 (IBM DB2 LUW)
- Sequence Wizard for IBM DB2 (IBM DB2 LUW)
- Structured Type Wizard for IBM DB2 (IBM DB2 LUW)
- Table Wizard for IBM DB2 (IBM DB2 LUW)
- Tablespace Wizard for IBM DB2 (IBM DB2 LUW)
- Trigger Wizard for IBM DB2 (IBM DB2 LUW)
- Unique Key Wizard for IBM DB2 for Windows, Unix, and Linux
- User Datatype Wizard for IBM DB2 LUW
- User Wizard for IBM DB2 LUW
- View Wizard for IBM DB2 LUW

In addition, see Create Synonym..

Alias Wizard for IBM DB2 (IBM 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, DBArtisan 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 - 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 - Properties

When creating or editing an alias, this tab/panel lets you work with the following settings:

Setting	Description
Schema	Select the schema that is to own the alias.
Name	Provide a name for the alias
Target Owner	Select the owner of the object to which you are creating an alias.
Target Type	Select the type of object to which you want to create an alias.
Target Name	Select the specific object to which you are creating an alias.

Bufferpool Wizard for IBM DB2 (IBM DB2 LUW)

The Bufferpool Wizard lets you create a bufferpool (an area of memory where data pages are read, altered, and held during an operation) without knowing the underlying commands. As you complete the Bufferpool Wizard, DBArtisan constructs the necessary CREATE BUFFERPOOL statement based on the information that you supply. Note that the bufferpool you're creating to store the data from the tables in selected database won't be active until you restart the database.

The Bufferpool Wizard lets you:

- Name the tablespace and specify the size of the bufferpool.
- Set storage options for UNIX systems.
- Set Bufferpool disbursement parameters.
- Override the default bufferpool size for partitions (or nodes).

To Open the Bufferpool Wizard

1 On the **Explorer**, find the database where you want to add the new bufferpool.

2 On the **Buffer Pools** branch of the **Performance** node, right-click and select **New**.

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 Buffer Pools - Properties.
- **Nodegroups** panel - for details, see Buffer Pools - Nodegroups.
- **Partitions** panel - for details, see Buffer Pools - Partitions.
- **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.

Buffer Pools - Properties

When creating or editing a buffer pool, this tab/panel lets you work with the following settings:

Setting	Description
Create deferred	Lets you defer changes until the next database activiation.
Pagesize	The page size, in KB, is available in the increments 4, 8, 16, or 32.
Automatic Size	Enables automatic size adjustment to account for workload.
Number of Pages	Indicate the number of pages you want to include in the new bufferpool. If the database is partitioned database, the default size is of all the partitions (or nodes) where the bufferpool exists. A value of -1, the default, means the size will be taken from the database configuration parameter (BUFFPAGE).
No. of Block Pages and **Block Size**	Lets you specify NUMBLOCKPAGES and BLOCKSIZE values used when creating or modifying buffer pools.

Buffer Pools - Nodegroups

Lets you select the Nodegroups on which to create the buffer pool. For more information, see Nodegroup Wizard for IBM DB2 (IBM DB2 LUW).

Buffer Pools - Partitions

Lets you modify the default size of assocaited partitions.

Database Wizard for IBM DB2 (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 DBArtisan 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 **Preview SQL** dialog box. For more information, see Preview.

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?	**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. **Use Database Managed Space:** If you select this option, the database manager controls the storage space for catalog tables. Add or Edit - For details, see Add/Edit Container for Tablespace.

>

Required Information	Description
What optional default storage parameters should be used.	Optionally identify values for extent size, pre-fetch size, transfer rate, and overhead.
Finish	Opens the **Preview SQ**L dialog box. For more information, see Preview.

Add/Edit Container for Tablespace

The table below describes the options and functionality on the Add and Edit Container for Tablespace dialog box:

Event Monitor Wizard for IBM DB2 for Linux, Unix, and Windows

An event monitor is used to collect information about a particular database and connected appliances when one or more events you specify take(s) place. Where a snapshot monitor captures information for a given moment in time, the event monitor lets you trace the activity you specify.

The Event Monitor Wizard lets you provide basic properties, select event types, and specify conditions where appropriate, for an event monitor.

To create a new event monitor using a wizard:

1. Open a creation wizard for an event monitor. 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 Event Monitors - Properties.
   - **Definition** panel - for details, see Event Monitors - Definition.
   - **Conditions** panel - for details, see Event Monitors - Conditions.
   - **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.

Event Monitors - Properties

When creating or editing an event monitor, this tab/panel lets you work with the following settings:

Setting	Description
Name	The name used to refernce the object.
Auto Start	Indicates whether the event monitor is activated each time the database starts.
Target Type	Indicates where the output is written: FILE, PIPE, or TABLE.
Buffer Size	For a **Target Type** of TABLE or FILE, indicates the buffer size
Blocked	For a **Target Type** of TABLE or FILE, indicate whether the event monitor is to be blocked.
Target	For a **Target Type** of PIPE, a named pipe identifier.
Write Mode	For a **Target Type** of FILE, the value can be APPEND or REPLACE.
Max Files	For a **Target Type** of FILE, indocate the maximum number of files.
Max FIle Size	For a **Target Type** of FILE, indicates the maximum file size.

>

Event Monitors - Definition

This tab/panel lets you specify the **Event Types** used to construct the CREATE EVENT MONITOR statement FOR clause.

The base set of event types is BUFFERPOOLS, CONNECTIONS, DATABASE, DEADLOCKS, DEADLOCK WITH DETAILS, STATEMENTS, TABLES, TABLESPACES, and TRANSACTIONS.

For IBM DB2 for Linux, Unix, and Windows version 9.0, DEADLOCK WITH DETAILS HISTORY and DEADLOCKS WITH DETAILS HISTORY VALUES event types are available.

For IBM DB2 for Linux, Unix, and Windows version 9.5, ACTIVITIES, THRESHOLD VIOLATIONS, and STATISTICS event types are available.

NOTE: If your choices are all from the DEADLOCKS... types, you can also specify **Deadlock Monitoring** of **Global** or **Local**. If you select Local, specify the **Database Partitions** the event monitor is to run against.

Event Monitors - Conditions

Available only for CONNECTION, STATEMENT, and TRANSACTION **Event Type**s, this tab lets you identify a condition used to determine which CONNECTION, STATEMENT, or TRANSACTION the event monitor will collect data for.

Foreign Key Wizard for IBM DB2 (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 DBArtisan'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 - Properties
   - **Column Mapping** panel - for details, see Foreign keys - 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 - Properties

When creating or editing a foreign key, this tab/panel lets you work with the following settings:

Setting	Description
Table Schema	The schema owning the table where the foreign key is being created.
Table Name	This is the table where the foreign key link originates--the child table.
Name	Lets you select a constraint name. System Generated Name - DB2 automatically generates a name. User Specified Constraint Name - You type the name.

>

Setting	Description
Enabled	TRUE enables the Foreign Key while FALSE disables the Foreigh Key.
Delete Rule	Select an action: **NO ACTION** - ensures that referenced values cannot be updated or deleted if to do so would violate referential integrity. **CASCADE** 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. **RESTRICT** prevents a command from executing if changes to referential data prompts an error.
Update Rule	Select an action. **NO ACTION** - ensures that referenced values cannot be updated or deleted if to do so would violate referential integrity. **RESTRICT** - prevents a command from executing if changes to referential data prompts an error

Foreign keys - 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. Then, under **Referenced Table**, select the corresponding column check boxes.

Function Wizard for IBM DB2 (IBM DB2 LUW)

To create a relationship between one set of values and another, DBArtisan 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, DBArtisan 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 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 .Functions - Properties
- **Advanced** panel - for details, see .Functions - Advanced
- **Source** panel (only available for a **Function Type** of SOURCED) - for details, see Functions - Source.
- **Parameters** panel (only available for a **Function Type** of SQL, SOURCED, or TEMPLATE) - for details, see Functions - Parameters.
- **Return Scalar** panel - for details, see Functions - Return Scalar.
- **Return Columns** panel (only available when you choose a **Function Type** of EXTERNAL TABLE) - for details, see Functions - Return Columns.
- **Body** panel - for details, see Functions - 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 - Properties

When creating or editing a foreign key, this tab/panel lets you work with the following settings:

Setting	Description
Schema, **Name**, and **Specific Name**	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.
Function Type	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 About function types.
Language	If you chose a **Function Type** of EXTERNAL SCALAR or EXTERNAL TABLE, specify a language of C, JAVA, or OLE. For more information, see About function types.
Return Type	For a **Function Type** of SQL, select ROW, TABLE, or SCALAR. For other **Function Type** choices, this indicates the default return type for that choice.
External Name	Provide the External Name of the function.
SQL Access Level	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 - 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 - 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 - 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 - 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 - 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 - 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.

Group Wizard for IBM DB2 for Windows/Unix

The Group Wizard lets you define new groups and to assign the appropriate object and system level privileges. Any user who is a part of the group you create will inherit the authorities and privileges assigned to the group.

To Open the Group Wizard

1. On the **Explorer**, find the datasource where you want to add the new group.
2. Expand the **Security** branch, right-click **Groups**, and select **New**.

The table that follows describes the fields you may encounter as you complete the Group Wizard:

Required Information	Description
What is the name of the group?	Lets you enter the name of the group.
What database authority(ies) do you want to give this group?	Select all or some of the authorities: BINDADD: User can create new database packages CONNECT: Gives users access to the database CREATETAB: A user can create new tables in the database CREATE_NOT_FENCED: Allows a user to create user-defined functions IMPLICIT_SCHEMA: Enables any user to implicitly create a schema by creating an object with a CREATE statement using a schema name that doesn't already exist. DBADM: Gives users administrative privileges. LOAD: Users can load data into a table. **NOTE**: The authorities listed here are the system privileges. You can specify object privileges after the new group has been created using the Groups Editor (IBM DB2 LUW).

Index Wizard for IBM DB2 (IBM 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 - Properties.
- **Columns** and **Include Columns** panels - for details, see Indexes - 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 - Properties

When creating or editing an index, this tab/panel lets you work with the following settings:

Setting	Description
Table Schema and **Table Name**	Choose the owner and name of the table in which the index is being created.
Schema and **Name**	Choose the owner and name of the index being created.
Index Type	Index enforces uniqueness on the values of the table's index key.
Clustered	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.
Percent Free	Lets you type or select the percentage of each index page to leave as free space when building the index, from 0 to 99.

Indexes - 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 for IBM DB2 (IBM 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
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 **Add**, **Insert** or **Edit** to add or modify table columns. Click **Drop** to deletea 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, DBArtisan 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, DBArtisan opens the Material Query Tables Editor.

Required Information	Description
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.
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 for IBM DB2 (IBM 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 for IBM DB2 (IBM 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 - Properties.
- **Columns** panel - for details, see Primary Keys - 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 - Properties

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
Table Schema and **Table Name**	Choose the owner and name of the table in which the primary key is being created.
Name	Choose the owner and name of the primary key being created.

Primary Keys - 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 for IBM DB2 (IBM 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 - Properties.
- **Advanced** panel - for details, see Procedures - Advanced.
- **Parameters** panel - for details, see Procedures - 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 - Properties

When creating or editing a procedure, this tab/panel lets you work with the following settings:

Setting	Description
Schema	Select the owner for the procedure.
Name	Provide the name of the function.
Specific Name	Optionally, provide the unique name of the procedure.
Language	Select among C, JAVA, COBOL, OLE, or SQL. The database manager will call the procedure accordingly.
External Name	Provide the external name of the procedure.
SQL Access Level	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 - Advanced

When creating or editing a procedure, this tab/panel lets you work with the following settings:

Setting	Description
Results Sets	Indicate the estimated upper bound of returned result sets. 0 is the default.
External Action	Select the External Action option.
New Save Point	Lets you specify a NEW SAVEPOINT LEVEL clause for the procedure.
Threadsafe	Specify whether the procedure is safe to run within the same process as other routines.
Fenced	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. 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.
Parameter Style	Lets you select an option: DB2DARI, DB2GENERAL, DB2SQL, GENERAL, GENERAL WITH NULLS, JAVA, and SQL. DB2GENERAL is for Java Language only. DB2SQL is for C, COBOL, or OLE Language only. GENERAL is for C Language only. GENERAL WITH NULLS is for C or COBOL Language only. JAVA is for Java Language only. SQL is for C, COBOL, or OLE Language only.
Program Type	**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. **SUBROUTINE**: the procedure expects the parameters to be passed as separate arguments.

Setting	Description
Inherit Special Registers	Lets you specify this optional clause dictating that the procedure will inherit initial values from the environment of the invoking statement.
DBINFO	Specific information contains such information as the database name, application ID, database code page, and so on.
Deterministic	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 - 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 for IBM DB2 (IBM 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, DBArtisan constructs the necessary CREATE SEQUENCE statement from the information that you have supplied. When finished, you can instruct DBArtisan to compile the sequence on the target DBArtisan 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 DBArtisan 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.

Required Information	Description
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 for IBM DB2 (IBM 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 for IBM DB2 (IBM 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 - Properties.
   - **Attributes** panel - for details, see Structured Types - 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 - Properties

When creating or editing a structured type, this tab/panel lets you work with the following settings:

Setting	Description
Attributes	**Instantiable** - if disabled, no constructor function can be generated, nor can a non-instantiable type be used as the type of a table or view. **Final Type** - enabled, indicates that the structured type cannot be used as a supertype. **With Function Access** - 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. **Without Comparisons** - If enabled, indicates that there are no comparison functions supported for instances of the structured type. **Inline Length** - 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).
Supertype	**Supertype Schema** and **Supertype Name** - let you designate this structured type as a subtype by providing details of the owning supertype.
Reference	**Cast (Source as Ref) With** and **Cast (Ref as Source) With** - 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. **Reference Using**, **Size**, **Precision**, and **Scale** - Define the built-in data type used as the underlying data type for the structured type you are creating and all its subtypes.

Structured Types - 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 for IBM DB2 (IBM 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 - Properties.
- **Columns** panel - for details, see Tables - Columns.
- **Partition** panel - for details, see Tables - Partition.
- **Tablespaces** panel - for details, see Tables - Tablespaces.
- **Dimension** panel - for details, see Tables - Dimensions.
- **Distribution Key Columns** panel - for details, see Tables - Distribution Key Columns.
- **Indexes** panel - for details, see Tables - Indexes.
- **Constraints** panel - for details, see Tables - 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 - Properties

When creating or editing a table, this tab/panel lets you work with the following settings:

Setting	Description
Schema	Select the schema that is to own the table.
Name	Provide a name for the table.
Percent Free	
Lock Size	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.
Append Data	Enable or Disable appending data to the end of the table.
Volatile	Enable this fature if a table contents may fluctuate from empty to very large. The access plan will not depend on the existing statistics for that table.
Compress	Enable or disable value compression
Row Compression	Enable or disable row compression.
Security Policy	Lets you add a security policy to a table.
RestrictDrop	Corresponds to the DB2 Restrict on Drop attribute.
Log Index Build	Enables this level of logging wne creating, recreating, or reorganizing an index.
CCSID	Specify ASCII or UNICODE or leave unspecified. If specified, this is the encoding scheme for string data. If unspecified, CCSID default encoding is used.
Tablespace, Index Tablespace, and Long Data Tablespace	Select a tablespace, an index tablespace, and a tablespace for Long or LOB table columns.

Setting	Description
Do not initially log	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.
Data Capture	Specify additional information logged by selecting DATACAPTURE NONE, DATA CAPTURE CHANGES, or DATA CAPTURE CHANGES INCLUDE LONGVAR.

Tables - 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 - 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 - 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 - 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 - 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 - Indexes

Click **Add** to open the Index wizard. For more information, see Index Wizard for IBM DB2 (IBM DB2 LUW).

Tables - Constraints

Selecting a constraint type and clicking **Add** opens the object wizard for that object type. For details see:

- Primary Key Wizard for IBM DB2 (IBM DB2 LUW)
- Unique Key Wizard (IBM DB2 Z/OS)
- Foreign Key Wizard (IBM DB2 Z/OS)

>

Tablespace Wizard for IBM DB2 (IBM 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 - Properties.
   - **Container** panel - for details, see Tablespaces - 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 - Properties

When creating or editing a tablespace, this tab/panel lets you work with the following settings:

Setting	Description
Tablespace properties	**Type** - Select REGULAR, LARGE, TEMPORARY, or USER TEMPORARY. **Use Automatic Storage** and **Managed By** let you specify whether storage is managed automatically, by the database, or by the system. **Database Partition Group** - lets you select a database partition group. **Buffer Pool** - lets you select a buffer pool. **Drop Recovery** - For REGULAR type tablesapces, lets you enable/disable drop recovery.
Performance properties	This group lets you specify or select the **Page SIze**, **Extent Size**, **Prefetch Automatic**, **Prefetch Size**,**Overhead**, **Transfer Rate**, and **File System Caching** properties.
Automatic Storage properties	This group lets you specify or select the **AutoResize**, **Initial Size**, **Increase Size**, **Max Size Unlimited**, and **Max Size** attributes.

Tablespaces - 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 for IBM DB2 (IBM 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 dandy 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 - Properties.
- **Column Selection** panel - for details, see Triggers - Column Selection.
- **Definition** panel - for details, see Triggers - Definition.

3 Finally, use the **Execute** button to create the object. For more information, see Create the new object.

Triggers - Properties

When creating or editing a trigger, this tab/panel lets you work with the following settings:

Setting	Description
Table Schema and **Table Name**	Choose the owner and name of the table for which the trigger is being created.
Schema and **Name**	Choose the owner and name of the trigger being created.
Trigger Timing	**BEFORE**: 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. **AFTER:** 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.
Trigger Events	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. An **UPDATE** trigger can be associated with specific columns of the base table and will only be activated if those columns are updated.
Trigger Type	**STATEMENT**: (only fires once). **ROW** (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.
Old Table Alias	Type the name of a temporary table of rows as they exist before they're updated or deleted.
New Table Alias	Type a name for a temporary table of rows as they exist after they're inserted or updated.
Old Row Alias	Type a name for the rows as they are before they've been deleted or updated.
New Row Alias	Type a name for the rows as they are after they've been inserted or updated.

Triggers - 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 - 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 for IBM DB2 for Windows, Unix, and Linux

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 - Properties.
   - **Columns** panel - for details, see Unique Keys - 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 - Properties

When creating or editing a unique key, this tab/panel lets you work with the following settings:

Setting	Description
Table Schema and **Table Name**	Choose the owner and name of the table in which the unique key is being created.
Name	Provide a name for the unique key being created.

Unique Keys - 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 Wizard for IBM 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 - 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 - Properties

When creating or editing a user datatype, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the owner of the user datatype.
Datatype	Provide a name for the datatype.
Type	Select the base datatype.
Size	Provide the size of the datatype.
Allow Bit Data	The option is only available for certain datatypes. A check means you want to store the data in a bit format.

User Wizard for IBM 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 - 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 - Properties

When creating or editing a user datatype, this tab/panel lets you provide the user **Name**.

View Wizard for IBM 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 - Properties.
- **Definition** panel - for details, see Views - Definition.

3 Finally, use the **Execute** button to create the object. For more information, see Create the new object.

Views - Properties

When creating or editing a view, this tab/panel lets you work with the following settings:

Setting	Description
Schema	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.
Name	Provide a name for the view.
Check Type	**CHECK_NONE** - No search conditions must be satisfied for insert or update operations. **CHECK_LOCAL** - 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. **CHECK_CASCADED** - 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.

Views - Definition

Complete the CREATE VIEW statement by typing or pasting in AS, SELECT, FROM, and WHERE statements or clauses. Finally, use the **Execute button** to Create the new object.

IBM DB2 for z/OS Object Wizards

DBArtisan lets you create DB2 OS390 objects using the following wizards:

- Alias Wizard (IBM DB2 Z/OS)
- Database Wizard (IBM DB2 Z/OS)
- Foreign Key Wizard (IBM DB2 Z/OS)
- Funcion Wizard (IBM DB2 Z/OS)
- Index Wizard (IBM DB2 Z/OS)
- Plan Wizard (IBM DB2 Z/OS)
- Primary Key Wizard (IBM DB2 Z/OS)
- Procedure Wizard (IBM DB2 Z/OS)
- Stogroup Wizard (IBM DB2 Z/OS)
- Synonym Wizard (IBM DB2 Z/OS)
- Table Wizard (IBM DB2 Z/OS)
- Tablespace Wizard (IBM DB2 Z/OS)
- Trigger Wizard (IBM DB2 Z/OS)
- Unique Key Wizard (IBM DB2 Z/OS)
- User Datatype Wizard (IBM DB2 Z/OS)

>

- User Wizard (IBM DB2 Z/OS)
- View Wizard (IBM DB2 Z/OS)

In addition, see Create Synonym.

Alias Wizard (IBM 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, DBArtisan 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 - 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 - Properties

When creating or editing a view, this tab/panel lets you work with the following settings:

Setting	Description
Schema	Select the schema that is to own the alias.
Name	Provide a name for the alias.
Target Owner	Select the owner of the object to which you are creating an alias.
Target Type	Select the type of object (TABLE, VIEW) to which you are creating an alias.
Target Name	Select the specific object to which you are creating an alias.

Database Wizard (IBM 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 - 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 - Properties

When creating or editing a database, this tab/panel lets you work with the following settings:

Setting	Description
Name	Provide a unique name for the database.
Type	**Workfile**: This option in 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.
Group Member	Specifies the member for which this database is being created. Use this only in a shared environment.
Tablespace Buffer Pool	Select the default buffer pool to be used for tablespaces created within the database.
Index Buffer Pool	Select the default buffer pool name to be used for indexes created within the database.
Storage Group	Select the default storage group to support the DASD space requirements for tablespaces and indexes within the database.
Encoding Scheme	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 (IBM 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 DBArtisan'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 - Properties.
- **Column Mapping** panel - for details, see Foreign Keys - 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 - Properties

When creating or editing a foreign key, this tab/panel lets you work with the following settings:

Setting	Description
Table Schema	Select the owner of the referencing, or child, table.
Table Name	Select the name of the referencing, or child, table.
Name	If you do not want to use the system-generated name, provide a new one.
Delete Rule	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 - 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.

Funcion Wizard (IBM DB2 Z/OS)

To create a relationship between one set of values and another, DBArtisan 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, DBArtisan 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, seeFunctions - Properties.
   - **Source** panel - (Only available for a Function Type of SOURCED.) for details, see Functions - Source.
   - **Parameters** panel - for details, see Functions - Parameters.
   - **Return Scalar** panel - for details, see Functions - Return Scalar.
   - **Return Columns** panel - (Only available when you choose a Function Type of EXTERNAL TABLE.) for details, see Functions - Return Columns.
   - **Body** panel - for details, see Functions - 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 - Properties

When creating or editing a function, this tab/panel lets you work with the following settings:

Setting	Description
Schema, **Name**, and **Specific Name**	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.
Function Type	**External Scalar**: This allows you to extend the function by adding your own or another party's definition for the function. **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. **Sourced**: Here you are creating a function that is based on an existing scalar or table function with an application server. **SQL**: This type of function returns a single value when the function is invoked if the SQL statement that defines it is valid.
Language	If you chose a Function Type of EXTERNAL SCALAR or EXTERNAL TABLE, specify a language of ASSEMBLE, C, COBOL, or PLI.
Return Type	Identifies the return type of the function.
External Name	Provide the External Name of the function.
SQL Access Level	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.
WLM Environment	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').
WLM For Nested	Self-explanatory

Functions - Source

NOTE: Only available for a **Function Type** of SOURCED.

Select the **Schema**, **Name**, and **Specific Name** of the source function.

Functions - 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 - 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 - 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 - Body

Enter the return statement for the function.

Index Wizard (IBM 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 - Properties
   - **Columns** panel - for details, see Indexes - Columns
   - **Indexes** panel - for details, see Indexes - Storage
   - **Partitions** panel - for details, see Indexes - 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 - Properties

When creating or editing an index, this tab/panel lets you work with the following settings:

Setting	Description
Table Schema and **Table Name**	Choose the owner and name of the table in which the index is being created.
Schema and **Name**	Choose the owner and name of the index being created.
Index Type	**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. **Non-Unique** (default)
Clustered	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.
Number of Partitions	If you enabled Clustering, specify the number of partitions.

Setting	Description
Buffer Pool	Provide the buffer pool in which this index should reside
Defer, **Close**, **Copy**, and **Define**	Enable or disable these DB2 options as required.
Piece Size	The maximum addressability of the dataset for a non-partitioned index.
Compress	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 **Buffer Pool** is 8k, 16k, or 32k.
Index on expression	Setting the check box enables the Index on Expression feature. The key expression is then provided on the **Columns** tab/panel. For details, see Indexes - Columns.
Padded	Setting this check box specifieat 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 **Columns** tab/panel.

Indexes - 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 - 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 - 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:

Free Page	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)
Percent Free	The percentage of free space you want to leave in every page when entries are added to an existing index. The default is 10%.
GBP Cache	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 - 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 (IBM 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 **Member Name**, **PDS Name** (partitioned data set) and click **Add** to enter each member and PDS name.	
3	Lets you select the **Location** to connect to, the **Collection** (location of the DBMS where the plan binds and where the description of the plan resides.) and a **Package**.	

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 **Commit** point or **Deallocate** 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.
6	**Enable** or **Disable** system connection types that can use the plan or package, select a **System** or **Cname** option.	

>

Primary Key Wizard (IBM 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 - Properties.
   - **Columns** panel - for details, see Primary Keys -Columns.
   - **Storage** panel - for details, see Primary Keys - Storage - Edit button (manage datasets) and Primary Keys - Storage - Attributes.
   - **Partitions** panel - for details, see Primary Keys - 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 - Properties

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
Table Schema and **Table Name**	Choose the owner and name of the table in which the index is being created.
Schema and **Name**	Choose the owner and name of the index being created.
Clustered	Enable or disable clustering.
Number of Partitions	If you enabled Clustering, specify the number of partitions.
Buffer Pool	Provide the buffer pool in which this index should reside
Defer, **Close**, **Copy**, and **Define**	Enable or disable these DB2 options as required.
Piece Size	The maximum addressability of the dataset for a non-partitioned index.

Primary Keys -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 - Storage - Attributes

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
Storage Group, **Primary Space Allocation**, **Secondary Space Allocation**, **Erase**, and **VCAT catalog**	The ability to set these options depends on the dataset management options you chose.
Free Page	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)
Percent Free	The percentage of free space you want to leave in every page when entries are added to an existing index. The default is 10%.
GBP Cache	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 - 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 (IBM 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 - Properties.
- **Parameters** panel - for details, see Procedures - 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 - Properties

When creating or editing a procedure, this tab/panel lets you work with the following settings:

Setting	Description
Schema, **Name**, and **Specific Name**	Select the schema that is to own the procedure, provide a name for the procedure, and provide the Specific Name for the procedure.
Language	The database manager will call the procedure accordingly assuming the program is designed to run in the server's environment. **Assemble:** a stored procedure written in Assembler **C:** a stored procedure written in C or C++ **COBOL:** a stored procedure written in COBOL **CompJAVA**: CompJAVA is no longer supported. Stored procedures should alternatively be written in JAVA. **JAVA** **PLI:** A stored procedure written in PL/I. **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. **SQL**
SQL Access Level	**MODIFIES SQL DATA (Default):** 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 **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.
WLM Environment	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').
WLM For Nested	Self-explanatory

>

Procedures - 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 (IBM 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 Add Volume Dialog Box.
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 Volumes Info Dialog Box.
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 (IBM 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 - 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 - Properties

When creating or editing a synonym, this tab/panel lets you work with the following settings:

Setting	Description
Schema	Select the schema that is to own the synonym.
Name	Provide a name for the synonym.
Referenced Object Owner	Select the owner of the object to which you are creating a synonym.
Referenced Object Type	Select the type of object (TABLE, VIEW) to which you are creating a synonym.
Referenced Object Name	Select the specific object to which you are creating a synonym.

Table Wizard (IBM 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 - Properties.
   - **Columns** panel - for details, see Tables - Columns.
   - **Indexes** panel - for details, see Tables - Indexes.
   - **Constraints** panel - for details, see Tables - 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 - Properties

Select the **Schema**, provide a **Name**, and provide or select other table properties.

Tables - 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 - Indexes

Click Add to open the Index Wizard (IBM DB2 Z/OS).

Tables - Constraints

Selecting a constraint type and clicking **Add** opens the object wizard for that object type. For details see:

- Foreign Key Wizard (IBM DB2 Z/OS)
- Primary Key Wizard (IBM DB2 Z/OS)
- Unique Key Wizard (IBM DB2 Z/OS)
- Create Synonym

>

Tablespace Wizard (IBM 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 - 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.

Tablespaces - Properties

When creating or editing a tablespace, this tab/panel lets you work with the following settings:

Group	Settings	Description
Creation	**Name** and **Database**	Provide a name and select an assocaited 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 http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp. Search on a term such as **universal tablespaces**.
Tablespace management	**Management type**	Lets you select how you want the tablespace managed: DEFAULT STORAGE GROUP, STORAGE GROUP, or VCAT CATALOG-NAME.
	Storage group, **Minimum primary space allocation**, **Minimum secondary space allocation**, and **Erase rule**	If you selected a **Management type** 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 **Management type** of VCAT catalog-name, select a catalog name.
Bufferpool	**Buffer Pool**	Select the buffer pool where the tablespace is to reside.

Group	Settings	Description
Partitions and size	**Number of partitions**	If you selected a **Creation Type** of PARTITIONED or RANGE-PARTITIONED UNIVERSAL, specify the number of partitions.
	Partition size (DSSIZE)	If you did not select a **Creation Type** of SEGMENTED, select a partition size.
	Segment size	If you selected a **Creation Type** 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.
Space Management	**Free space portion of each page (%)**	If you did not select a **Creation Type** 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 **Creation Type** of LOB, specify how often to leave a free page when a tablespace is reloaded or reorganized.
Other parameters	**GBPCACHE**	Select a group buffer pool cache scheme of CHANGED, ALL, SYSTEM, or NONE.
	Compress	Enable this setting to allow compression.
	Track modified pages	Enable this seting to track modified pages in the space map pages.
	Encoding scheme	Select an encoding scheme of EBCDIC, UNICODE, ASCII, or NONE.
	Log	Enabling this check box specifies that changes to a LOB column are to be written to the log.
	Define	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.
	Member Cluster	Enable this seting to manage space for inserts on a member-by-member basis.
	Close rule	Enable this seting to specify CLOSE YES, dictating the data set close priority.
	Lock Size	Specify a lock sizeof ANY, TABLESPACE, PAGE, or ROW.
	Maximum locks	Specify the number of locaks that are allowed before escalating.

Trigger Wizard (IBM 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 - Properties.
- **Column Selection** panel - for details, see Triggers - Column Selection.
- **Definition** panel - for details, see Triggers - Definition.

3 Finally, use the **Execute** button to create the object. For more information, see Create the new object.

>

Triggers - Properties

When creating or editing a trigger, this tab/panel lets you work with the following settings:

Setting	Description
Table Schema and **Table Name**	Choose the owner and name of the table for which the trigger is being created.
Schema and **Name**	Choose the owner and name of the trigger being created.
Trigger Timing	**BEFORE**: 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. **AFTER:** 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.
Trigger Events	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. An **UPDATE** trigger can be associated with specific columns of the base table and will only be activated if those columns are updated.
Trigger Type	**STATEMENT**: (only fires once). **ROW** (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.
Old Table Alias	Type the name of a temporary table of rows as they exist before they're updated or deleted.
New Table Alias	Type a name for a temporary table of rows as they exist after they're inserted or updated.
Old Row Alias	Type a name for the rows as they are before they've been deleted or updated.
New Row Alias	Type a name for the rows as they are after they've been inserted or updated.

Triggers - 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 - 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 (IBM 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 - Properties.
- **Columns** panel - for details, see Unique Keys -Columns.
- **Storage** panel - for details, see Unique Keys - Storage - Edit button (manage datasets) and Unique Keys - Storage - Attributes.
- **Partitions** panel - for details, see Unique Keys - 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 - Properties

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
Table Schema and **Table Name**	Choose the owner and name of the table in which the unique key is being created.
Schema and **Name**	Choose the owner and name of the index being created.
Clustered	Enable or disable clustering.
Number of Partitions	If you enabled Clustering, specify the number of partitions.
Buffer Pool	Provide the buffer pool in which this unique key should reside
Defer, **Close**, **Copy**, and **Define**	Enable or disable these DB2 options as required.
Piece Size	The maximum addressability of the dataset for a non-partitioned index.

Unique Keys -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 - 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 - 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
Storage Group, **Primary Space Allocation**, **Secondary Space Allocation**, **Erase**, and **VCAT catalog**	The ability to set these options depends on the dataset management options you chose.
Free Page	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)
Percent Free	The percentage of free space you want to leave in every page when entries are added to an existing index. The default is 10%.
GBP Cache	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 - 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 (IBM 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:

>

3

- **Properties** panel - for details, see User Datatypes - Properties.
- **Comment** panel - for details, see Add a comment to an object.
- **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.

User Datatypes - Properties

When creating or editing a user datatype, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the owner of the user datatype.
Datatype	Provide a name for the datatype.
Type	Select the base datatype.
Size	Provide the size of the datatype.
For Data	Select the MIXED, SBCS, or BIT option for the datatype.
CCSID	Select the NONE, ASCII, EBCDIC, or UNICODE option for the datatype.

User Wizard (IBM 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 - 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 - Properties

When creating or editing a user datatype, this tab/panel lets you provide the user **Name**.

>

View Wizard (IBM 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 - Properties.
- **Definition** panel - for details, see Views - Definition.

3 Finally, use the **Execute** button to create the object. For more information, see Create the new object.

Views - Properties

When creating or editing a view, this tab/panel lets you work with the following settings:

Setting	Description
Owner	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.
Name	Provide a name for the view.
Check Type	**CHECK_NONE** - No search conditions must be satisfied for insert or update operations. **CHECK_LOCAL** - 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. **CHECK CASCADED** - 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.

Views - Definition

Complete the CREATE VIEW statement by typing or pasting in AS, SELECT, FROM, and WHERE statements or clauses. Finally, use the **Schedule** or **Execute** buttons to Create the new object.

Microsoft SQL Server Object Wizards

DBArtisan lets you create SQL Server objects using the following wizards:

Backup Device Wizard (SQL Server)

The Backup Device Wizard lets you add a disk or tape backup device to a SQL Server instance. Backup devices store backups of databases and transaction logs.

Before you can dump a database or its transaction log, first link to a backup device that stores the resulting backup file.

To create a new backup device using a wizard:

1. Open a creation wizard for a backup device. 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 Backup Devices - Properties.
   - **DDL View** panel - for details, see Create the new object.
3. Finally, use the **Execute** button to create the object.

Backup Devices - Properties

When creating or editing a backup device, this tab/panel lets you work with the following settings:

Setting	Description
Name	Lets you provide a logical name for the backup device.
Device Type	For Microsoft SQL Server 2005, DISK and TAPE are supported. For Microsoft SQL Server 2000, DISK, TAPE, and PIPE are supported.
Skip ANSI Labels	Available only with a **Device Type** of TAPE, ANSI tape labels are ignored if this check box is enabled.
Physical Name	The physical name of the device. This value can be any simple file name (sqltable.dat for example) or full path (a:\sqltable.dat for example) that is recognized by the operating system.

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 - Options.
   - **Placement** panel - for details, see Databases - Placement.
   - **Transaction Log** panel - for details, see Databases - 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 - Options

When creating or editing a database, this tab/panel lets you work with the following settings:

Setting	Description
Name	Provide a name for the database.
Attach existing OS files	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.
Compatible Level	Select a version compatibility level.

Setting	Description
Properties group	Select the following settings: **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**.

Databases - 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 - 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 - Properties.
   - **Dependencies** panel - Defaults - 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 - Properties

When creating or editing a default, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the schema that is to own the default.
Name	Provide a name for the default.
Value	Provide the value of the default.

Defaults - 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 - 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 - Properties

When creating or editing an extended procedure, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the owner of the extended procedure.
Name	Provide a name for the extended procedure.
Library	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 DBArtisan'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 - Properties.
   - **Column Mapping** panel - Foreign Keys - 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 - Properties

When creating or editing a foreign key, this tab/panel lets you work with the following settings:

Setting	Description
Table Owner and **Table Name**	Select the owner and name of the table for which the foreign key is being created.
Name	Provide a name for the foreign key.
Enabled	Enables or disables the foreign key.
Not For Replication	Replication copies and distributes data and database objects from one database to another and then synchronizes information between databases for consistency.
Delete Rule	If you choose the CASCADE option, all rows containing data involved with the foreign key will be deleted after a delete operation.
Update Rule	If you choose the CASCADE option, all rows containing data involved with the foreign key will be deleted after an update operation.

Foreign Keys - 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 - Properties.
   - **Definition** panel - for details, see Functions - Definition.
3. Finally, use the **Execute** button to create the object. For more information, see Create the new object.

Functions - Properties

When creating or editing a function, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the owner of the function.
Name	Provide a name for the function.
Schema Binding	Choose whether the function is bound to database objects that it references.
Encryption	Choose whether SQL Server encrypts table columns that contain the text of the CREATE FUNCTION statement.

Functions - 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 - Properties.
   - **Columns** panel - Indexes - 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 - Properties

When creating or editing an index, this tab/panel lets you work with the following settings:

Setting	Description
Parent Type	Select TABLE or VIEW.

>

Setting	Description
Parent Owner	Select the owner of the table or view.
Parent Name	Select the specific table or view containing the columns you want to index.
Name	Provide a name for the index.
Build Online	Enabling this feature specifies that the ONLINE=ON clause is used when creating this object and can susequently 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.
Index Type	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.
Clustered	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. In a nonclustered index, the physical order of rows is independent of the indexed order of rows. For an indexed view, you can only create a nonclustered index if there is already a clustered index extant. With this check box enabled, subsequent rebuild and drop operations offer an online option.
Ignore Duplicate Key	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.
Statistics Recompute	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.
File Group	
Fill Factor	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.
Pad Index	Enable or disable padding of index pages.
Sort in TempDB	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.
Allow Row Locks and **Allow Page Locks**	Lets you enable locking granularity at the page and row level (Default=TRUE) **NOTE:** You cannot reorganize an index (primary key, or unique key) that has an **Allow Page Locks** property set to FALSE. For information on reorganizing indexes, see Reorganizing SQL Server indexes, primary keys, and unique keys.

Indexes - 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.

Linked Server Wizard (SQL Server)

A linked server object represents a database system other than Microsoft SQL Server. The links are created through OLE DB datasources. OLE DB is a COM-based application programming interface (API) for accessing data. OLE DB supports accessing data stored in any format (databases, spreadsheets, text files, and so on) for which an OLE DB provider is available. With a linked server you can execute commands against OLE DB datasources on different servers. Linked servers let you issue distributed queries, commands, and so on.

The wizard prompts you to provide much of the same information you give when you initially register a server.

To create a new linked server using a wizard:

1. Open a creation wizard for a linked server. 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 Linked Servers- Properties.
   - **DDL View** panel - for details, see Create the new object.
3. Finally, use the **Execute** button to create the object.

Linked Servers- Properties

When creating or editing a linked servere, this tab/panel lets you work with the following settings:

Group	Settings and descriptions
General	If defining a linked server defined through Microsoft SQL Server, set the **Is SQL Server** check box. The name must be the server's network name. Any tables retrieved from the server are from the default database defined for the login on the linked server. If defining a linked server not defined through Microsoft SQL Server, clear the **Is SQL Server** check box and use the following options to provide OLE DB details: **Product Name** - The product name of the OLE DB datasource to which you are linking. **Datasource** - The name of the OLE DB provider managing access to the linked server. **OLE DB Provider Name** - The name of the OLE DB provider managing access to the specified linked server. **Provider Location** - The OLE DB location property corresponding to the linked server. **Provider Connection String** - The OLD DB provider-specific connection string that identifies a unique datasource. **Catalog** - The name of the catalog to be used when making a connection to the OLE DB provider.

Group	Settings and descriptions
Server Options	**RPC Out Server** - Lets you select the server with RPC encryption. **RPC** - Enables RPC encryption from the server. **Collation Compatible Server** - Indicates Change DIstributed Query execution against linked servers. If you select this option, Microsoft SQL Server assumes that all characters in the linked server are compatible with the local server, with regard to character set and collation sequence (or sort order). This option is for Microsoft SQL Server 8.0 or later. **Lazy Schema Validation** - This option lets you optimize performance by ensuring the query processor does not request metadata for any of the linked tables until data is actually needed from the remote member table. This option is for Microsoft SQL Server 8.0 or later. **Use Remote Collation** and **Collation Name** - Let you indicate that the linked server should use remote collation and select the name of the collation. This option is for Microsoft SQL Server 8.0 or later. **Query Timeout** - Type the query timeout value. This option is for Microsoft SQL Server 8.0 or later. **Connection Timeout** - Type the connection timeout value. This option is for Microsoft SQL Server 8.0 or later.

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 - Properties.
   - **Columns** panel - Primary Keys - 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 - Properties

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
Table Owner and **Table Name**	Choose the owner and name of the table in which the primary key is being created.
Name	Provide a name of the primary key being created.
Build Online	Enabling this feature specifies that the ONLINE=ON clause is used when creating this object and can susequently 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.
Clustered	Enable or disable clustering. With this check box enabled, subsequent rebuild and drop operations offer an online option.
File Group	If you do not specify a filegroup, Microsoft SQL Server creates the index in the default filegroup.
Fill Factor	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 - 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 - Properties.
- **Definition** panel - for details, see Procedures - Definition.

3 Finally, use the **Execute** button to create the object. For more information, see Create the new object.

Procedures - Properties

When creating or editing a procedures, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the owner of the procedure.
Name	Provide a name for the procedure
Procedure Number	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.
Replication	This option creates a procedure that is used as a stored procedure filter and is executed only during replication.
Recompile	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.
Encryption	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 - 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.

Remote Server Wizard (SQL Server)

Remote Servers are servers on a network that you can access through your local server. Instead of logging directly into a remote server, you can log into your local server and execute a remote procedure call to it.

NOTE: Only stored procedures are allowed against remote servers. You can use a linked server to execute distributed queries and stored procedures.

>

Before beginning

Before you can set up a remote server, you first configure both the local and remote servers to accept remote access. To configure the Microsoft SQL Servers for remote access, you need to log into each server. The table below describes the parameters you need to configure on each server:

Parameter	Description
Remote Access	1(enabled)
Remote Connections	Number of remote connections required
Remote Logins	Number of remote logins required
Remote Sites	Number of remote sites required

After setting these server configuration parameters, shutdown and restart the server so the new values can take effect.

To create a new remote server using a wizard:

1. Open a creation wizard for a remote server. 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 Remote Servers - Properties.
   - **DDL View** panel - for details, see Create the new object.
3. Finally, use the **Execute** button to create the object.

Remote Servers - Properties

When creating or editing a remote server, this tab/panel lets you work with the following settings:

Setting	Description
Name	The name for the remote server.
Server Type	Lets you specify a server type of LOCAL or REMOTE. The local server routes all commands to the remote server. To communicate with a remote server, you log into a local server and submit your remote procedure calls for remote servers. The remote server receives remote procedure calls from the local server and processes them accordingly.
Publication Server	Select if the server is to publish data for replication.
Distribution Server	Select If the server is to manage the distribution database.
Subscription Server	This option applies if you want to deploy a subscription management application from the remote server.
Publisher/Subscriber Server	Select if the server is both a subscriber and publisher of replicated data.
Collation Compatible Server	Select if the server is to be collation compatible.
RPC Out Server	Select if the server is to serve as a RPC Out server.

Role Wizard (SQL Server)

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. Roles can help you comply with Sarbanes Oxley regulations by limiting which users can have access to what privileges, for example a Human Resources Role versus an Accounting Role.

To create a new role using a wizard:

1. Open a creation wizard for a role. 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 Roles - Properties.
   - **DDL View** panel - for details, see Create the new object.
3. Finally, use the **Execute** button to create the object.

After you click **Finish** and the role has been created, the Roles Editor open. You can then assign object and system privileges to the role and determine which users can take part in the role. For more information, see Roles Editor (SQL Server).

Roles - Properties

When creating or editing a role, this tab/panel lets you work with the following settings:

Setting	Description
Name	The name used to refernce the role.
RoleType	The value can be NONE, APPLICATION, or STANDARD.
Authorization Owner	Only required with a **Role Type** of STANDARD.
Password	Only required with a **Role Type** of APPLICATION.

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 - Properties.
   - **Dependencies** panel - Rules - 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 - Properties

When creating or editing a rule, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the Owner of the rule.
Name	Provide a name for the rule.
Restriction	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 - 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)

DBArtisan 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 - Properties.
   - **Permissions** panel - Schema - 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 - Properties

When creating or editing a schema, this tab/panel lets you provide a **Name** and select the **Owner** of the schema.

Schema - 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 - Properties.
- **Columns** panel - Tables - Columns.
- **Indexes** panel - Tables - Indexes.
- **Constraints** panel - Tables - 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 - Properties

When creating or editing a table, this tab/panel lets you work with the following settings:

Setting	Description
Schema	Select the owner of the table.
Name	Provide a name for the table
ANSI_NULLS option	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.
Filegroup	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.
Text Image Filegroup	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 - 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** lsit. 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 - Indexes

Click Add to open the Index Wizard (IBM DB2 Z/OS).

Tables - 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 - Properties.
   - **Definition** panel - for details, see Triggers - 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 - Properties

When creating or editing a trigger, this tab/panel lets you work with the following settings:

Setting	Description
Parent Type	Select the TABLE or VIEW on which the trigger is to be created.
Parent Schema	Select the owner of the table or view on which the trigger is to be created.
Parent Name	Select the specific table or view in which the trigger is to be created.
Name	Provide a name for the trigger.
Trigger Timing	**INSTEAD OF:** 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. **AFTER**: 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.
Fire On Insert	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.
Fire On Update	An **UPDATE** trigger can be associated with specific columns of the base table and will only be activated if those columns are updated.
Fire On Delete	A **DELETE** trigger is associated with a DELETE operation.
Encrypted	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 - 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 - Properties.
- **Columns** panel - Unique Keys - 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 - Properties

When creating or editing a unique key, this tab/panel lets you work with the following settings:

Setting	Description
Table Owner and **Table Name**	Choose the owner and name of the table in which the unique key is being created.
Name	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 susequently 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.
Clustered	Enable or disable clustering. With this check box enabled, subsequent rebuild and drop operations offer an online option.
File Group	If you do not specify a filegroup, Microsoft SQL Server creates the unique key in the default filegroup.
Fill Factor	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 - 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 - Properties.
   - **Information** panel - for details, see User Messages - Information.
   - **DDL View** panel - for details, see Create the new object.
3. Finally, use the **Execute** button to create the object.

User Messages - Properties

When creating or editing a user message, this tab/panel lets you work with the following settings:

Setting	Description
Message Number	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.
Severity	Lets you select a predefined severity level between 001 and 025. The levels and their meanings are indicated in the User Message editor.
Write to NT Event Log	Lets you specify that the message is always written to the Windows NT Event Log.

User Messages - 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 - 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 - Properties

When creating or editing a user, this tab/panel lets you work with the following settings:

Setting	Description
Login Name	Select the server login associated with this user.
Name	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 - 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 - Properties

When creating or editing a user datatype, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the owner of the user datatype.
Datatype	Provide a name for the datatype.
Type	Select the base datatype.
Size	Provide the size of the datatype.
Allow Nulls	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.
Default Binding	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.
Rule Binding	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 - Properties.
- **Definition** panel - for details, see Triggers - Definition.

3 Finally, use the **Execute** button to create the object. For more information, see Create the new object.

Views - Properties

When creating or editing a view, this tab/panel lets you work with the following settings:

Setting	Description
Owner	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.
Name	Provide a name for the view.
Encryption	If you select this option, the view will not be published with SQL Server replication.
Schema Binding	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
Check Condition	When a row is modified through a view, this option makes sure the data remains visible through the view after the modification is committed.
Owner	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.

Oracle Object Wizards

DBArtisan 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)
- Job Queue 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)
- Redo Log Group 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

Required Information	Description
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”	**Add** or **Edit** Button - For more information, see Add or Modify Cluster Column.. **Drop** 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.
What is the cluster type?	**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. **Hash**: Rows with the same hash key value are stored together. This is helpful if the tables are static.
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?	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. **Default** is the Oracle internal hash function, otherwise specify the hash expression you want to use.
How many transaction entries are allowed for each data block in the cluster?	Each transaction that updates a data block requires a transaction entry. **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. **Maximum (1-255)**: The maximum parameter limits concurrency on a data block.
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?	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 (1-99)
How large are the cluster’s extents?	The unit of space allocated to an object whenever the object needs more space. Initial Extent - The initial space extent (in bytes) allocated to the object. Next Extent - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required.
Specify the number of free lists	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. **Free Lists:** The default and minimum value is 1; this option should be set higher if multiple processes access the same data block.
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.

Required Information	Description
Define a default buffer pool for this cluster	Default - Select to retain the default. Keep - Select to retain the object in memory to avoid I/O conflicts. **ORACLE 8i ONLY:** Recycle - Select to rid data blocks from memory as soon as they are no longer in use.
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**.

   DBArtisan 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 - Properties.
   - **Column Mapping** panel - Foreign Keys - 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 - Properties

When creating or editing a foreign key, this tab/panel lets you work with the following settings:

Setting	Description
Table Owner	The owner of the table where the foreign key is being created.
Table Name	This is the table where the foreign key link originates--the child table.
Name	Lets you select a constraint name. System Generated Name - DB2 automatically generates a name. User Specified Constraint Name - You type the name.
Enabled	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.
Delete Rule	Select an action: **NO ACTION** - ensures that referenced values cannot be updated or deleted if to do so would violate referential integrity. **CASCADE** 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 - 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 - 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 - Properties

When creating or editing a function, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the owner of the function.
Name	Provide a name for the function.

Functions - 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. DBArtisan 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 - Properties.
- **Columns** panel - Indexes - Columns.
- **Storage** panel - Indexes - Storage.
- **Partition** panel - Indexes - 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 - Properties

When creating or editing an index, this tab/panel lets you work with the following settings:

Setting	Description
Table Owner and **Table Name**	Choose the owner and name of the table in which the index is being created.
Owner and **Name**	Choose the owner and provide the name of the index being created.
Index Type	NONUNIQUE - In a non-unique index, the ROWID is treated as part of the key. Oracle treats a constraint as deferrable. UNIQUE - Select if the index is a unique constraint.The values in the indexed columns must be distinct. 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.
No Sort	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.
Logging	Enabling logs this operation to the redo file.
Reverse	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.
Function-Based	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.
No Parallel Execution	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. 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.
Parallel Degree	The value you select indicates the number of query server processes that should be used in the operation.
Parallel Instances	The value you select indicates how you want the parallel query partitioned between the Parallel Servers.

>

Indexes - 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 - Storage

When creating or editing an index, this tab/panel lets you work with the following settings:

Setting	Description
Data Block Storage group	Select the DEFAULT **Tablespace** 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.) 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. **Percent Free** identifies how much space you want to allocate for new rows or updates to existing rows. **Initial Transactions** ensures that a minimum number of concurrent transactions can update an index block, avoiding the overhead of allocating a transaction entry dynamically. **Maximum Transactions** limits concurrency on an index block.
Extents group	An extent is the unit of space allocated to an object whenever the object needs more space. **Initial Extent** - The initial space extent (in bytes) allocated to the object. **Next Extent** - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required. **Percentage Increase** - Lets you type the percentage. **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. **Minimum Extents** - 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. **Maximum Extents** - 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.
Freelists group	**Free lists** 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. The default and minimum value is 1. You should increase this number if multiple processes access the same data block. **Free List Groups** is the number of groups of free lists. **NOTE:** This option is only applicable for the parallel server option.
Buffer Pool	**DEFAULT** - Choose this if you want to use the default bufferpool. **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. **RECYCLE** - Select this option to save cache space by ridding data blocks from memory as soon as they are no longer in use.

Indexes - Partition

Clicking Create Partition opens the Add, Insert, or Modify Partition wizard.

Job Queue Wizard (Oracle)

Job Queues are built-in mechanisms that let you schedule a variety of SQL-based or command-line driven tasks.

The Job Queue Wizard lets you:

- Specify the PL/SQL code that will run in the job.
- Specify when the job will run, if it will run again, and if it should be enabled to run.

To Open the Job Queue Wizard

1. On the **Explorer**, find the datasource where you want to create an Oracle job and expand the **Instance** node.
2. Right-click the **Oracle Job Queue** node, and select **New**.

The table that follows describes the fields you may encounter using this wizard:

Required Information	Description
Enter the PL/SQL code you would like submitted	Type PL/SQL code or retrieve a previously saved PL/SQL script.
No Parse/Parse	Parsing can be expensive resource-wise, so think carefully before you choose the parse option. For a parse call, Oracle checks the syntactic and semantic validity of the SQL statement, whether the process has the correct permissions, and allocates a private SQL area for the statement. Oracle will also check to see whether the statement exists in a shared library cache.
When would you like for the job to begin execution?	Self-explanatory.
Would you like to have the job run on an ongoing basis?	Self-explanatory.
Would you like to have the job submitted as disabled?	Self-explanatory.

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.
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?	**Fast** - Using the information logged on the materialized view logs or a partition maintenance operation, the refresh applies incremental changes. See Fast Refresh Requirements for more information. **Complete** - This refresh recalculates the materialized view's defining query. **Force** - Applies fast refresh when feasible, otherwise uses a complete refresh. **Never** - Materialized view will not be refreshed with the optional refresh mechanisms.
Choose a refresh mechanism	**On Demand** - This option requires that all refreshes be manually executed. **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. **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.
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	**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. **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.
How many transaction entries are allowed for each datablock in the materialized view?	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. 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. Maximum (1-255) - Limits concurrency on a data block.
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?	**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 views extents?	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. **Initial Extent** (KB) - The initial space extent (in bytes) allocated to the object. **Next Extent** - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required.
How many extents should be allocated to the materialized view?	**Minimum Extents** - The appropriate minimum extents value for the object. **Maximum Extents** - The appropriate maximum extents value for the object.
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**

Required Information	Description
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
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 **Add Partition** dialog box may open. For details, see Add, Insert, or Modify Partition.

Required Information	Description
Specify number of subpartions	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
STDDEV with COUNT and SUM		X	X
WHERE clause includes join predicates which can be ANDed bit 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

	When the Materialized View has:		
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?	**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. **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.
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?	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. 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. Maximum (1-255) - Limits concurrency on a data block.
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?	**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.

Required Information	Description
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 - Properties.
   - **Definition** panel - Outlines - 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 - Properties

When creating or editing an outline, this tab/panel lets you work with the following settings:

Setting	Description
Stored Outline Name	Provide a unique name for the outline.
Category	The category is a name you want to use to group stored outlines. (You can type over the word Default that appears automatically.)

Outlines - 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 - Properties.
   - **Columns** panel - Primary Keys - Columns.
   - **Storage** panel - Primary Keys - Storage.
   - **Partition** panel - Primary Keys - 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 - Properties

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
Table Owner and **Table Name**	Choose the owner and name of the table in which the primary key is being created.
Name	Provide a name for the primary key being created.
No Sort	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.
Logging	Enabling logs this operation to the redo file.
Reverse	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.
Validate	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 cheeked against the constraint.
Deferrable	Dictates whether constraint checking can be deferred until the end of a transaction.
Deferred	This option is enabled only if you enabled the **Deferrable** 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.
Enabled	Enables or disables the primary key.

>

Primary Keys - 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 - Storage

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
Data Block Storage group	Select the DEFAULT **Tablespace** 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.) 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. **Percent Free** identifies how much space you want to allocate for new rows or updates to existing rows. **Initial Transactions** ensures that a minimum number of concurrent transactions can update an index block, avoiding the overhead of allocating a transaction entry dynamically. **Maximum Transactions** limits concurrency on an index block.
Extents group	An extent is the unit of space allocated to an object whenever the object needs more space. **Initial Extent** - The initial space extent (in bytes) allocated to the object. **Next Extent** - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required. **Percentage Increase** - Lets you type the percentage. **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. **Minimum Extents** - 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. **Maximum Extents** - 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.
Freelists group	**Free lists** 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. The default and minimum value is 1. You should increase this number if multiple processes access the same data block. **Free List Groups** is the number of groups of free lists. **NOTE:** This option is only applicable for the parallel server option.
Buffer Pool	**DEFAULT** - Choose this if you want to use the default bufferpool. **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. **RECYCLE** - Select this option to save cache space by ridding data blocks from memory as soon as they are no longer in use.

Primary Keys - Partition

Clicking **Create Partitio**n 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 - 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 - Properties

When creating or editing a procedure, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the owner of the procedure.
Name	Provide a name for the procedure

Procedures - 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)
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)

Required Information	Description
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

Redo Log Group Wizard (Oracle)

A redo log makes it possible to replay SQL statements because a redo log stores all changes made to a database as they occur. If something unanticipated happens to a datafile, you can restore a backed-up datafile and use the redo log to bring the datafile back to the state it was in before the unanticipated activity. The redo log can consist of two parts: the online redo log and the archived redo log. In archive log mode, redo log files will not be overwritten unless they have been archived somewhere. Each log is a member of the group.

NOTE: Redo should not be confused with Undo, which is a transaction based process.

If you find that you have problems with redo log group availability and Oracle consistently has to wait for redo log groups to become free, you can create additional redo log groups. DBArtisan lets you create additional redo log groups to ensure that there is always a recycled redo log group available.

The Redo Log Group Wizard lets you:

- Assign a number to the new redo log group and determine the file size for the group members.
- Choose a thread for the redo log group for Oracle Parallel Server.
- Add redo log members with the Add Redo Log Member dialog box.

To Open the Redo Log Wizard

1. On the **Explorer**, find the datasource where you want to create a redo log and expand the **Storage** node.
2. Right-click the **Redo Log Groups** node, and select **New**.

The table that follows describes the fields you may encounter as you complete the wizard.

Required Information	Description
Do you want to use Oracle Managed Files (OMF) for the redo log group? You must be set up to use OMF to use this option	Lets you specify to use Oracle Managed Files to automatically create the redo log group. Oracle automatically creates and removes operating system files as you create and drop storage objects. **NOTE:** If you select OMF, you do not have to enter a redo log group number or file size.
What is the group number for the redo log group?	When numbering your redo log groups, it is better not to skip numbers (e.g., 5, 10, 15). Skipping numbers causes Oracle to consume extra space in the control files of the database.

Required Information	Description
What file size should the group's redo log members be? Note: All members of the group must be the same. With OMF, you can let Oracle assign a size or enter it manually.	size in KB or MB
To which thread should the redo log group be added? For this option to be available, you need to be using Oracle with the parallel server option in parallel mode.	**OPTIONAL:** If you are using Oracle Parallel Server, type the thread. A thread is an operating system component that allows the logic of multiuser applications to be performed as several separate, asynchronous execution paths. Threads help ensure that work is being performed for some user connections even when other connections are blocked (for example, when waiting for a disk read or write operation to complete).
Add redo log group members	The name of each group member should be a file name. For example, D:\ORACLE\ORADATA\FINANCEDB\REDO02.LOG is a valid group member name. If you do not specify a full path, DBArtisan creates files in either the default or current directory of the database server.

Role Wizard (Oracle)

As you complete the Role Wizard, DBArtisan 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?	Not identified - Selecting this means the role you are creating will be enabled immediately. Identified: If this is your choice, you're indicating one of the following authorization methods will be followed: Globally - Select to indicate that Oracle permits access to the user by obtaining user name and password information from the security domain central authority. Externally - Select to indicate that Oracle should verify the database user name against an existing operating system user name. 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.

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 a 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.

Required Information	Description
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.

>

Required Information	Description
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
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 - 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 - Properties

When creating or editing a synonym, this tab/panel lets you work with the following settings:

Setting	Description
Owner and **Name**	Choose the owner and provide a name for the synonym being created.
Referenced Object Owner	Select the owner of the object to which the synonym is to refer.
Referenced Object Type	Select the type of the object to which the synonym is to refer.
Referenced Object Name	Select the specific object to which the synonym is to refer.
Database Link	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, DBArtisan 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 Cluster Wizard (Oracle).

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 - Properties.
- **Columns** panel - for details, see Tables - Columns.
- **Indexes** panel (not available with a **Row Organization** of **EXTERNAL**) - for details, see Tables - Columns.
- **Constraints** panel (not available with a **Row Organization** of **EXTERNAL**) - for details, see Tables - Constraints.
- **Storage** panel (not available with a **Row Organization** of **EXTERNAL**) - for details, see Tables - Storage.
- **IOT Properties** panel (not available with a **Row Organization** of **EXTERNAL**) - for details, see Tables - IOT Properties.
- **Partition** (not available with a **Row Organization** of **EXTERNAL**) - for details, see Tables - 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 - Properties

When creating or editing a table, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the owner of the table.
Name	Provide a name for the table
Cache	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.
Row Movement	Enabling this option permits the migration of a row to a new partition if its key is updated.
Parallel Degree	A value indicating the number of query server processes that should be used in the operation.
Parallel Instances	A value indicating how you want the parallel query partitioned between the Parallel Servers.
Physical group	Choose a Row Organization of **INDEX**, **HEAP**, or **EXTERNAL**. If you chose **INDEX** or **HEAP**, enable or disable **Logging**. If you chose **EXTERNAL**, provide an **External Type**, **Default Directory**, **Access Parameters**, **Reject Limit**, and **Location**..
Logging	Redo logs minimize the loss of data in the event that an uncontrolled shutdown happens.

Tables - 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 - Indexes

NOTE: This tab/panel is not available with a **Row Organization** of **EXTERNAL**.

Click **Add** to open the Index Wizard (Oracle).

Tables - 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 for IBM DB2 (IBM DB2 LUW)
- Unique Key Wizard for IBM DB2 for Windows, Unix, and Linux
- Foreign Key Wizard for IBM DB2 (IBM DB2 LUW)

Tables - 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
Data Block Storage group	Select the DEFAULT **Tablespace** only if you want the partitions in the same tablespace as the partitions in the underlying table. **Percent Free** identifies how much space you want to allocate for new rows or updates to existing rows. **Initial Transactions** ensures that a minimum number of concurrent transactions can update a primary key block, avoiding the overhead of allocating a transaction entry dynamically. **Maximum Transactions** limits concurrency on a primary key block.
Extents group	An extent is the unit of space allocated to an object whenever the object needs more space. **Initial Extent** - The initial space extent (in bytes) allocated to the object. **Next Extent** - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required. **Percentage Increase** - Lets you type the percentage. **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. **Minimum Extents** - 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. **Maximum Extents** - 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.

>

Settings	Description
Freelists group	**Free lists** 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. The default and minimum value is 1. You should increase this number if multiple processes access the same data block. **Free List Groups** is the number of groups of free lists. **NOTE:** This option is only applicable for the parallel server option.
Buffer Pool	**DEFAULT** - Choose this if you want to use the default bufferpool. **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. **RECYCLE** - Select this option to save cache space by ridding data blocks from memory as soon as they are no longer in use.

Tables - 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 - Partition

NOTE: This tab/panel is not available with a **Row Organization** of **EXTERNAL**.

Prior to working with partitions, you should be familiarwith the material in DBArtisan and Oracle Partitioning.

Click **Create Partition** to Partition a table.

DBArtisan 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).

>

Partition Type	Description
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
Properties	Select a **Partition Type** 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.
Subpartition Columns (only available with a **Partition Type** of **RANGE** and **Subpartition Type** of **HASH** or **LIST**)	For each column, click the New button and select a name from the Column dropdown. Use the Delete button to drop a selected column.
Subpartitions (only available with a **Subpartition Type** of **HASH**)	Specify a **Default number of partitions**. For each partition, click the New button and then select a tablespace from the dropdown.
Range Definitions (only available with a **Partition Type** of **RANGE**)	Click the New button to Add a partition definition.
Partition Definition (only available with a **Partition Type** of **HASH**)	To specify a partition method other than **None**, take one of the following actions: (1) Select the **Number Of Partitions** radio box, specify the **Number Of Partitions**, and for each partition, click the New button and choose a tablespace from the dropdown, or (2) select the **By Partition Name** 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 **Partition Type** of **List**)	Click New to Add a partition definition.

Add a partition definition

Use the following topics as a guide in completing the settings in this wizard:

Step	Settings and tasks	
Partition Definition	**Name**	Provide a name.
	Tablespace	Select a tablespace from the drodown
	Logging	Enable or disable logging.
Subpartitions	To specify a partition method other than **None**, select the **By Subartition Name** radio box and for each partition, click the New button to open a dialog that lets you provide subpartition values. **NOTE:** When you split a range-list partition, you cannot specify the new partitions' subpartition information.	
Storage	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 to set 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?	**Permanent**: Contains schema objects you want to keep on an ongoing basis. The objects are stored in datafiles. **Temporary**: Schema objects will last only as long as your session continues. **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.
Do you want space management to be performed through the data dictionary or managed locally?	**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. **Locally Managed:** Each tablespace manages its own free and used space so locally managed tablespaces are generally much more efficient.
Will the extents be allocated by the system or do you want to choose a uniform extent size?	**Autoallocate extents:** Users cannot specify an extent size because the tablespace is system managed. This option is not available for temporary tablespaces. **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.
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?	**NOTE:** This option is not available for temporary tablespaces. **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. **Offline**: This is a tablespace that's unavailable to users even when the database is open. The datafiles in the tablespace are offline too.
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 Adding or modifying a tablespace datafile

Required Information	Description
What type of segment management should be used?	This is an option available only to permanent, locally-managed tablespaces. **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). **Manual (Oracle free lists)**: The database will manage the segment free space using free lists.
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-*n*K_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.

   DBArtisan 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**.

   DBArtisan adds the datafile to the tablespace definition.
6. Click **Close**.

   DBArtisan 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 - Properties.
- **Column Selection** panel - for details, see Triggers - Column Selection.
- **Action** panel - for details, see Triggers - 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 - Properties

When creating or editing a trigger, this tab/panel lets you work with the following settings:

Setting	Description
Parent Type, **Parent Owner**, and **Parent Name**	Select the type (TABLE or VIEW), owner and specific object for which the trigger is being created.
Name	Provide a name for the trigger being created.
Enabled	Enable or disable the trigger.
Trigger Timing	**BEFORE**: 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. **AFTER:** 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.
Fire On Insert, **Fire On Update**, and **Fire On Delete**	Enable the events that fire the trigger. 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. An **UPDATE** trigger can be associated with specific columns of the base table and will only be activated if those columns are updated. A **DELETE** trigger fires automatically after items in the table are deleted.
Trigger Type	**STATEMENT**: (only fires once). **ROW** (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.
Old Table Alias	Type the name of a temporary table of rows as they exist before they're updated or deleted.
New Table Alias	Type a name for a temporary table of rows as they exist after they're inserted or updated.
When Clause	Type a WHEN clause or open a WHEN CLAUSE wizard to further qualify the trigger behavior.

Triggers - 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 - 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 - Properties.
- **Columns** panel - for details, see Unique Keys - Columns.
- **Storage** panel - for details, see Unique Keys - Storage.
- **Partition** panel - for details, see Unique Keys - 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 - Properties

When creating or editing a unique key, this tab/panel lets you work with the following settings:

Setting	Description
Table Owner and **Table Name**	Choose the owner and name of the table in which the unique key is being created.
Name	Provide a name for the unique key being created.
No Sort	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.
Logging	Enabling logs this operation to the redo file.
Reverse	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.
Validate	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.
Deferrable	Dictates whether constraint checking can be deferred until the end of a transaction.
Deferred	This option is enabled only if you enabled the **Deferrable** 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.
Enabled	Enables or disables the primary key.

Unique Keys - 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 - Storage

When creating or editing a primary key, this tab/panel lets you work with the following settings:

Setting	Description
Data Block Storage group	Select the DEFAULT **Tablespace** 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.) 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. **Percent Free** identifies how much space you want to allocate for new rows or updates to existing rows. **Initial Transactions** ensures that a minimum number of concurrent transactions can update an index block, avoiding the overhead of allocating a transaction entry dynamically. **Maximum Transactions** limits concurrency on an index block.
Extents group	An extent is the unit of space allocated to an object whenever the object needs more space. **Initial Extent** - The initial space extent (in bytes) allocated to the object. **Next Extent** - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required. **Percentage Increase** - Lets you type the percentage. **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. **Minimum Extents** - 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. **Maximum Extents** - 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.
Freelists group	**Free lists** 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. The default and minimum value is 1. You should increase this number if multiple processes access the same data block. **Free List Groups** is the number of groups of free lists. **NOTE:** This option is only applicable for the parallel server option.
Buffer Pool	**DEFAULT** - Choose this if you want to use the default bufferpool. **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. **RECYCLE** - Select this option to save cache space by ridding data blocks from memory as soon as they are no longer in use.

Unique Keys - 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 - Properties.
   - **Roles** panel - for details, see Users - Roles.
   - **Quotas** panel - for details, see Users - 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 - Properties

When creating or editing a user, this tab/panel lets you work with the following settings:

Setting	Description
Name	Provide a name for the user.
Default Tablespace and **Temporary Tablespace**	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 Set Default.
Profile	The profile you choose affects the amount of database resources available to the user.
Identified By	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.
Password	If you specified an **Identified By** value of REQUIRED_YES, provide a password for the user.
External Name	If you specified an **Identified By** value of REQUIRED_EXTERNAL, provide an external name for the user.
Account Locked	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.
Password Expired	Marks the password as expired, forcing the user to change the password before being allowed to connect to the database.

Users - Roles

For each role to be assigned to the user, select the check box beside that role.

Users - 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 - 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 - Properties

When creating or editing a view, this tab/panel lets you work with the following settings:

Setting	Description
Owner	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.
Name	Provide a name for the view.

Views - 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

- DBArtisan lets you create Sybase objects using the following wizards:
- Alias Wizard (Sybase)
- Database Device Wizard (Sybase)
- Database Wizard (Sybase)
- Data Cache Wizard (Sybase)

- Default Wizard (Sybase)
- Dump Device 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)
- Remote Server Wizard (Sybase)
- Role Wizard (Sybase)
- Rule Wizard (Sybase)
- Segment Wizard (Sybase)
- Table Wizard (Sybase)
- Trigger Wizard (Sybase)
- Unique Key Wizard (Sybase)
- User Datatype Wizard (Sybase)
- User Message 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 - 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 - Properties

When creating or editing an alias, this tab/panel lets you work with the following settings:

Setting	Description
Name	Provide a name for the alias being created.
User	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 Device Wizard (Sybase)

A database device is part or all of a disk device used for storing a database and its transaction log.

The Database Device Wizard lets you specify basic creation attributes and override many obscure, rarely-used parameters of the DISK INIT statement.

> **NOTE:** Make sure that your Sybase ASE is configured for enough devices that it can accommodate the creation of another device. If not, increase that parameter from the **Server Configuration** dialog box. For details, see Configuring the Server.

To Open the Database Device Wizard

1. On the **Explorer**, find the datasource where you want to add a database device and expand the **Storage** node.
2. Right-click the **Database Devices** node, and select **New**.
3. Use the following topics as a guide to understanding and setting options on the panels of this wizard:
   - **Properties** panel - for details, see Database Devices - Properties.
   - **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.

Database Devices - Properties

When creating or editing a database device, this tab/panel lets you work with the following settings:

Setting	Description
Creation	**Name** - the abstract name you want to assign to the device. If for any reason the physical device is down, you can redefine the logical name to a new device.
Attributes	**Physical Name** - specify the full path for the file, including the hard drive letter or machine name, directory and full file name with the *.DAT extension (for example, D:\SYB\DATA\TEST.DAT) **Device Number** - DBArtisan automatically calculates the next open device number and defaults its value in the box. Numbering is sequential. **Controller number** - Normally, Sybase ASE uses 0 **Size** - in kilobytes, megabyes, gigabytes, or tera bytes. **Default Device** - The default device is where all new DBArtisan databases are placed unless the user specifies an exact device in the CREATE DATABASE statement. **Virtual Address**, **Dsync**, and **DirectIO** - Let you specify DISK INIT and DISK REINIT parameters.
Autogrowth	Lets you enable/disable **Auto Growth** and if enabled, to specify **Growth rate** and **Maximum Size** values.

>

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 - Properties.
   - **Placement** panel - for details, see Databases - 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 - Properties

When creating or editing a database, this tab/panel lets you work with the following settings:

Setting	Description
Database Name	Provide a name for the database.
For Load	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.
With Override	Enable this option to specify whether logs and data are to be kept on the same logical device
Type	Optionally, use the TEMP option to designate this as a temporary database.
Options group	Enable or disable the following Sybase database options: **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**, and **unique auto_identity index**.

Databases - 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.

Data Cache Wizard (Sybase)

The data cache holds the log, data, and index pages that are currently being used or were recently accessed. There's a default cache loaded by Sybase, but you can add additional caches without reducing the size of the default. You can bind a table, index, image, or text to the cache you are creating.

>

To create a new data cache using a wizard:

1 Open an object wizard for a data cache.

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 Data Caches - 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.

Data Caches - Properties

When creating or editing a data cache, this tab/panel lets you work with the following settings:

Setting	Description
Name	Lets you provide a name for the cache.
General properties	**Type** - lets you specify a cache status of DEFAULT, MIXED, or LOGONLY. **Policy** - lets you specify a replacement policy of STRICT or RELAXED. **Status** - displays the current status of the cache: ACTIVE, PENDING ACTIVE, or PENDING DELETE. You can only change a mixed data cache if it does not have any non-log objects bound to it. **Partition** - displays the local cache partition number.
Size properties	**Configured Size** - Lets you specify the size of the cache after the server is restarted. **Run Size** - displays the size that the cache is currently using. For the default cache, this is the amount of all cache space not explicitly configured for another cache

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 - Properties.
- **Dependencies** panel - for details, see Defaults - 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 - Properties

When creating or editing a default, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the schema that is to own the default.

Setting	Description
Name	Provide a name for the default.
Value	Provide the value of the default.

Defaults - 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**.

Dump Device Wizard (Sybase)

Your system writes its memory image to a dump device so you can retrieve data if your system crashes. All load and dump operations require that you name a device before the operation will complete. When you create a new dump device, ASE adds a row to the sysdevices table in the master database. Sybase recommends that you back up the master database before you create a new device and again after you have created the device. Be sure you have enough disk space on the physical drive to perform these operations.

To Open the Dump Device Wizard

1. On the **Explorer**, find the datasource where you want to add a dump device and expand the **Security** node.
2. Right-click the **Dump Devices** 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 device type?	Tape or disk.
What is the dump device Name?	This is the device's logical name.
What is the physical name?	The device or file name that is recognized by the operating system. (For example, a:sqltable.dat).

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 - 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 - Properties

When creating or editing an extended procedure, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the owner of the extended procedure.
Name	Provide the name of the procedure.
Library Name	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 DBArtisan'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 - Properties.
   - **Column Mapping** panel - for details, see Foreign Keys - 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 - Properties

When creating or editing a foreign key, this tab/panel lets you work with the following settings:

Setting	Description
Table Owner	Select the owner of the referring, or child, table
Table Name	Select the name of the referring, or child, table.
Match Full	Specify a referential integrity option.

Foreign Keys - 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 - Properties.
   - **Columns** panel - for details, see Indexes - Columns.
   - **Partition** panel - for details, see Indexes - 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 - Properties

When creating or editing an index, this tab/panel lets you work with the following settings:

Setting	Description
Table Owner, **Table Name**, and **Name**	Select or provide identifying information.
Attributes	Provide or select values for the following Sybase options: **Index Type**, **Clustered**, **Ignore Duplicate Key**, **Ignore Duplicate Rows**, and **Maximum Rows Per Page** properties.
Storage	Provide or select values for the following Sybase options: **Reserve Page Gap**, **Segment Name**, **Fill Factor**, **Prefetch Strategy**, and **MRU Replacement Strategy**.

Indexes - 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 - 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	
Properties	Select a **Locality**.	
Partition Definitions	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 **Add** 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, DBArtisan 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.

Required Information	Description
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 - Properties.
   - **Columns** panel - for details, see Primary Keys - 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 - 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 - 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 - Properties.
   - **Definition** panel - for details, see Procedures - Definition.
3. Finally, use the **Execute** button to create the object. For more information, see Create the new object.

Procedures - Properties

When creating or editing a procedure, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the owner of the procedure.
Name	Provide the name of the procedure.
Procedure Number	Provide a procedure number if you want to group identically-named procedures.
Recompile	Optionally, enable this feature to prevent Sybase from saving a plan for the procedure.

Procedures - Definition

Complete the provided CREATE PROCEDURE statement by typing or pasting the body of the procedure.

Remote Server Wizard (Sybase)

Remote Servers are servers on a network that you can access through your local server. Instead of logging directly into a remote server, you can log into your local server and execute a remote procedure call to it.

NOTE: Only stored procedures are allowed against remote servers. You can use a linked server to execute distributed queries and stored procedures.

Before beginning

- Before you can set up a remote server, you first configure both the local and remote servers to accept remote access. To configure the Microsoft SQL Servers for remote access, you need to log into each server. The table below describes the parameters you need to configure on each server:

Parameter	Description
Remote Access	1(enabled)
Remote Connections	Number of remote connections required
Remote Logins	Number of remote logins required
Remote Sites	Number of remote sites required

- After setting these server configuration parameters, shut down and restart the server so the new values can take effect. When you have configured both servers, proceed to the steps below.

To Open the Remote Server Wizard

1 On the **Explorer**, expand the server where you want to add a remote server.

2 Expand the **Security** node

3 Right-click the **Remote Servers** node and 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 server?	Self-explanatory
What is the server's physical/network name?	Optional: Self-evident.
Do you want to add a remote server or identify the local server?	Self-evident.
Indicate whether SQL server should time out inactive physical connections with this remote server.	Yes, time out No, run until shut down
Indicate whether encryption should be used to protect passwords sent to this remote server	Yes, encrypt No, use plain text
Which security model would you like to use for remote procedure calls	**Security Model A**: This is equivalent to the level of security in release 11.0.x and prior releases. **Security Model B:** Implements security using DCE, Kerberos, or NT LAN Manager.
Specify a security mechanism	Optional: Self-evident.
If you selected security Model B on the previous page you may enable the following option	Mutual authentication: The local server authenticates the remote server by retrieving the remote server credential Message confidentiality: Encrypts messages to and from the remote server Message integrity: Ensures messages between servers will be checked for tampering.

Role Wizard (Sybase)

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. Roles can help you comply with Sarbanes Oxley regulations by limiting which users can have access to what privileges, for example a Human Resources Role versus an Accounting Role.

To Open the Role Wizard

1. On the **Explorer**, expand the server where you want to add a remote server.
2. Expand the **Security** node.
3. Right-click the **Roles** node and 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 role?	Self-explanatory
Enter the password, if any, that must be entered to activate the role	Seven characters is the default.
Enter the password expiration if any	OPTIONAL Self-explanatory
Enter the minimum password length, if any	OPTIONAL Self-explanatory
Enter the minimum number of login attempts, if any	OPTIONAL Self-explanatory

After you click **EXECUTE** and the Role is successfully created, the Role Editor opens where you can go about the serious business of assigning object and system privileges.

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 - Properties.
   - **Dependencies** panel - for details, see Rules - 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 - Properties

When creating or editing a rule, this tab/panel lets you work with the following settings:

Setting	Description
Owner and **Name**	Select an **Owner** and provide a **Name** for the rule.
Restriction	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).
Type	Choose STANDARD_RULE, AND_ACCESS_RULE, or OR_ACCESS_RULE.

Rules - 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 devise 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**, DBArtisan 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 - Properties.
- **Columns** panel - for details, see Tables - Columns.
- **Indexes** panel - for details, see Tables - Indexes.
- **Constraints** panel - for details, see Tables - Constraints.
- **Partitions** panel - for details, see Tables - Partition.
- **Permissions** panel - for details, see Tables - 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 - Properties

When creating or editing a table, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the owner of the table.
Name	Provide a name for the table
Segment Name	Specify the segment on which you want to place the table.
Maximum Rows Per Page	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.
Reserve Page Gap	
Identity Gap	Specify the size of the identity gap for the table.
MRU Replacement Strategy	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.
Prefetch Strategy	Enabling this feature allows you to fetch as many as eight 2K data pages simultaneously instead of one at a time (the default).
Lock Scheme	Select a locking scheme of ALLPAGES, DATAPAGES, or DATAROWS.
Expected Row Size	If you specified a **Lock Scheme** of DATAROWS or DATAPAGES, provide an expected row size

Tables - 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 - Indexes

Click **Add** to open the Index Wizard (Sybase).

Tables - 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 - Partition

Click **Add** to open the Table Partition wizard for Sybase ASE.

Tables - 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
Properties	Select a **Partition Type** of ROUNDROBIN, RANGE, HASH, or LIST
Columns (not available with a **Partition Type** of ROUNDROBIN)	For each column, click the Insert a new 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.
Partition Definition (only available with a **Partition Type** of HASH or ROUNDROBIN)	Take one of the following actions: (1) Select the **Number Of Partitions** radio box, specify the **Number Of Partitions**, and for each partition, click the New button and choose a tablespace from the dropdown, or (2) select the **By Partition Name** radio box and for each partition, click the New button provide a name and then choose a tablespace from the dropdown.
Range Definitions (only available with a **Partition Type** of RANGE)	Click the Define a new partition button to open a dialog that lets you provide a name, select a segment and **Add** the partition definition.
List Definitions (only available with a **Partition Type** 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 **Add** 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 - Properties.
   - **Definition** panel - for details, see Triggers - Definition.
3. Finally, use the **Execute** button to create the object. For more information, see Create the new object.

Triggers - Properties

When creating or editing a trigger, this tab/panel lets you work with the following settings:

Setting	Description
Table Owner, and **Table Name**	Select the owner and specific table for which the trigger is being created.
Name	Provide a name for the trigger being created.
Fire On Insert, **Fire On Update**, and **Fire On Delete**	Enable the events that fire the trigger. 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. An **UPDATE** trigger can be associated with specific columns of the base table and will only be activated if those columns are updated. A **DELETE** trigger fires automatically after items in the table are deleted.

Triggers - 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 - Properties.
   - **Columns** panel - for details, see Unique Keys - 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 - 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 - 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 - 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 - Properties

When creating or editing a user datatype, this tab/panel lets you work with the following settings:

Setting	Description
Datatype	Provide a name for the datatype.
Type	Select the base datatype.
Size	Provide the size of the datatype.
Allow Nulls	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.
Identity	This allows you to specify whether or not you want to limit access to the new datatype based on the user's privileges.
Default Binding	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.
Rule Binding	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 Message Wizard (Sybase)

A user message let's you write the error message users will see when a user-defined event transpires. You have the option of saving the message to the event log. You can also create messages in multiple languages, but you have to create the message in US English before you can write it in other languages.

To Open the User Message Wizard

1 On the **Explorer**, expand the datasource where you want to write a user message.

2 Open the database, right-click the **User Messages** Node, and then click **New**.

The table that follows describes the fields you will encounter as you complete the wizard.

U

Required Information	Description
What is the message number?	A message number must be greater than 2,000. The wizard automatically loads with the next unclaimed message number, but you can use another number if you prefer.
Add the text for this message	Click **Add** and write the message you want to attach to the message. Otherwise, what's the point of using the wizard?

When you have successfully created the new message, the User Messages Editor opens where you can bind the message to foreign key or check constraints as well as edit the message text.

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 - 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 - Properties

When creating or editing a user, this tab/panel lets you work with the following settings:

Setting	Description
Login Name	Choose a login ID from the drop-down list.
Name	Provide the user name.
Group Name	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 - Properties.
   - **Definition** panel - for details, see Views - 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 - Properties

When creating or editing a view, this tab/panel lets you work with the following settings:

Setting	Description
Owner	Select the owner of the view.
Name	Provide a name for the view.
Check Type	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 - Definition

Complete the CREATE VIEW statement by typing or pasting in the SELECT, FROM, and WHERE statements. Finally, use the **Execute button** to Create the new object.

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
- 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:

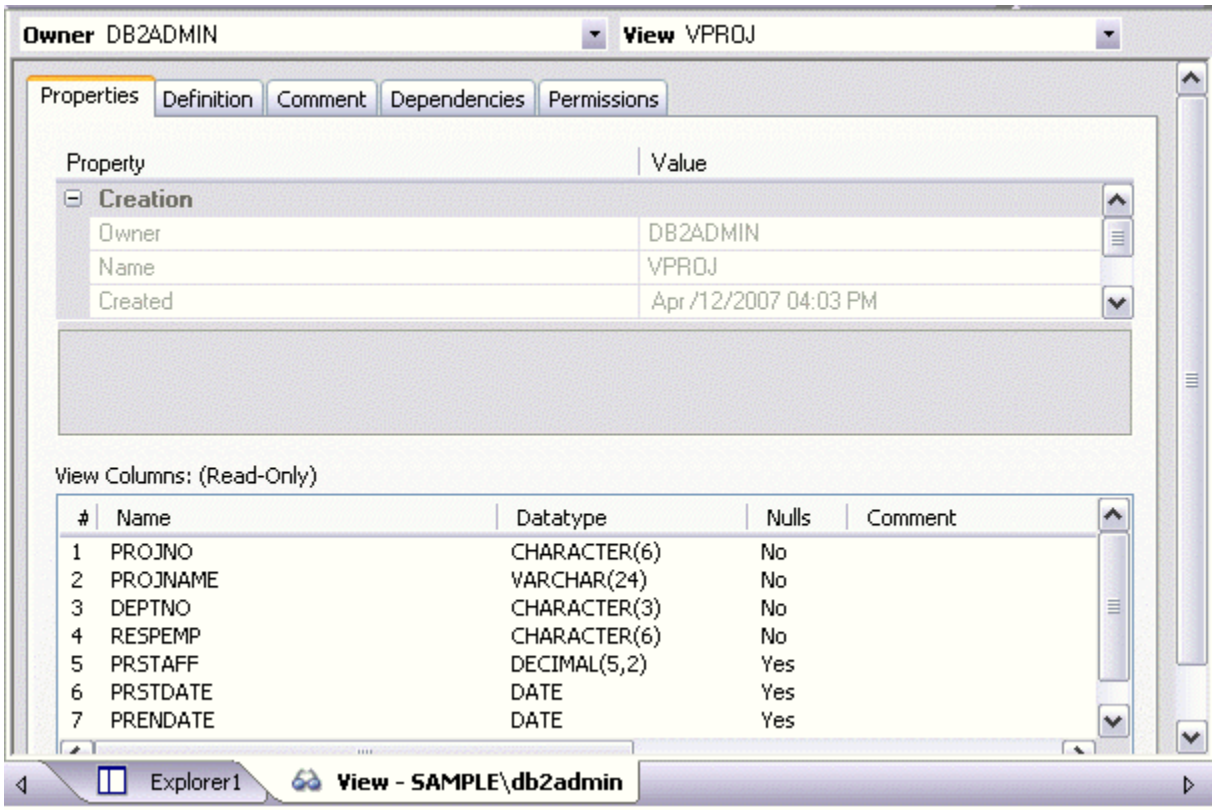

In order to work with object editors, you must be familiar with the following tasks:

- Opening an object editor on a server or database object.
- Viewing and modifying database 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 DBArtisan 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 under the datasource until the type of object you want to view or modify is visible.
4. On the **Datasource Explorer**, click on the specific type of object that you want to view or modify.

   All objects of that type on the datasource are displayed in the right pane of the explorer window.

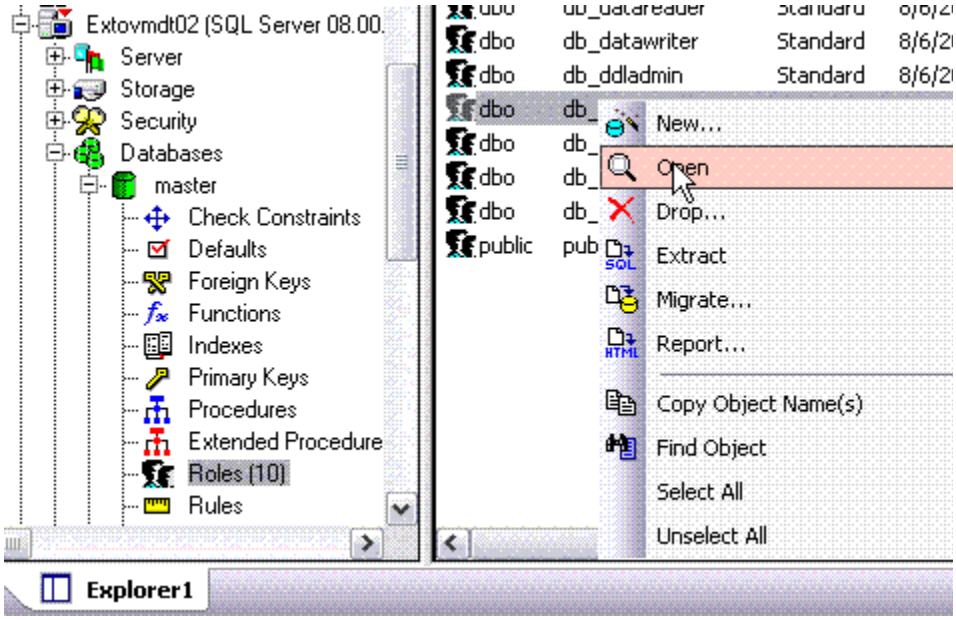

5 In the right pane of the Explorer window, right-click the specific object you want to view or modify, and then select **Open**.

DBArtisan opens an editor for that object type on the object you selected.

For information on the editing objects, see Viewing and modifying database object properties.

Viewing and modifying database 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/DDL 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:

Create	Launches a DBArtisan object creation wizard. For more information, see Creating objects.
Alter	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 Preview.
Drop	Lets you drop one or more database objects and remove their definition from the system catalog. For more information, see Drop.
Extract	DBArtisan 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 Extract.
Migrate	Lets you move object schema between DBMS or servers by copying the DDL between the source and target databases. For details, see Migrate.
Report	Lets you generate detailed or summary reports on database objects. For details, see Report.
Command Menu	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.
Refresh	Refreshes or clears Object editor contents.
Close	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:

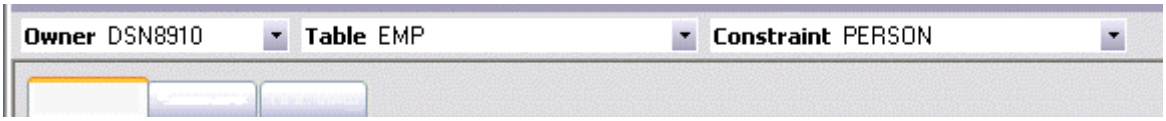

These lists allow you to qualify any object of the same type as the object currently displayed in the Object editor, and display information for that object in the current Object editor. The number of, and specific dropdown lists, differ according to the type of Object editor but there are always sufficient 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/DDL 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.

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.

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.

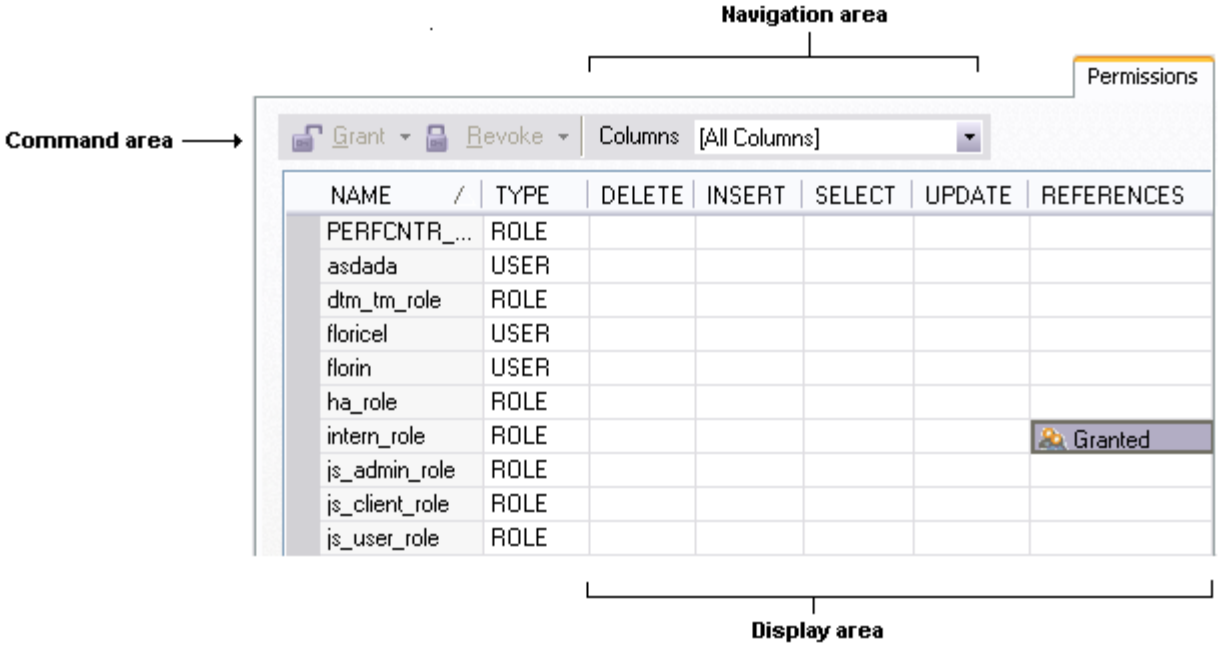

The Navigation area lets you change the content of the Display area to view more specific privilege details. For example, when viewing privileges for a database object such as a table, the dropdown lists in the navigation area, if present, let you drill down to populate the display area with privileges 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.

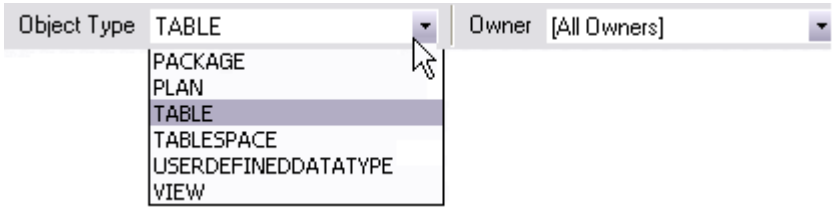

The Display area shows privilege details for an object or grantee. When viewing privileges for a grantee such as a user or role, the display area shows privileges for that recipient on objects currently qualified by the current selection in the Navigation area. When viewing privileges for an object such 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 Privilege(s) To or Revoke Privilege(s) From, database objects/grantees.
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, DBArtisan lets you view all dependencies. The Object editors for referencing or referenced objects have a **Dependencies** tab that shows all dependencies for an object.

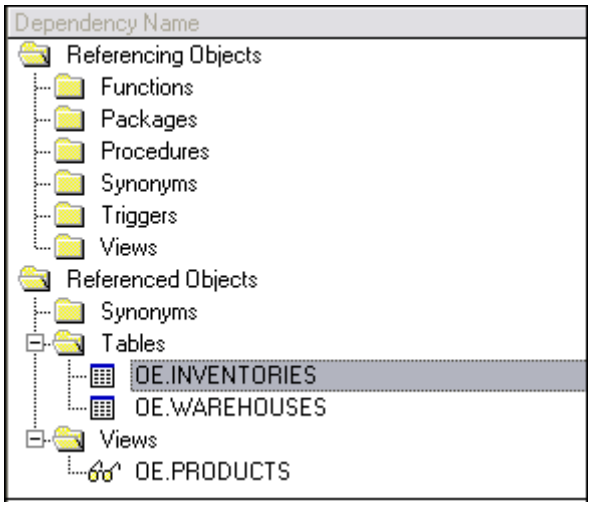

In addition to letting you view dependencies, Object editors let you drop a referencing or referenced object, or open an Object editor on that object.

To manage dependences 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.

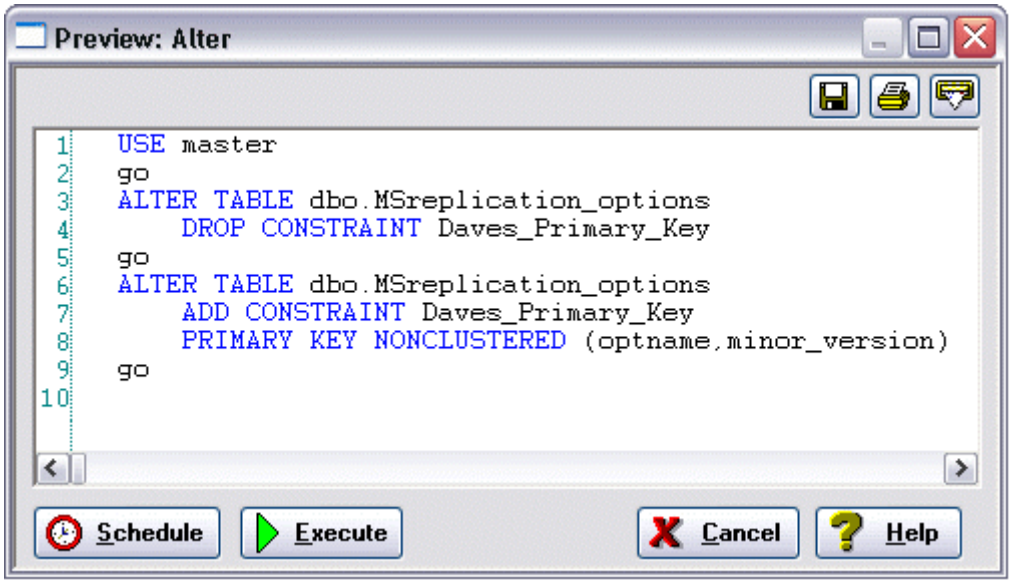

2 Use one of the following options to submit and execute your changes:

- **Execute** button - executes the changes immediately.
- **Schedule** button - opens a dialog that lets you schedule the change. For more information, see Scheduling.
- **Batch** button - (IBM DB2 for z/OS only) lets you specify the job cards used to execute the changes as a batch job. For information on how to check status, see Database Monitor.

IBM DB2 for Linux, Unix, and Windows Object Editors

DBArtisan 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)**
- **Bufferpool Editor (IBM DB2 LUW)**
- Check Constraints Editor (IBM DB2 LUW)
- Databases Editor (IBM DB2 LUW)
- Event Monitor Editor (IBM DB2 LUW)
- Foreign Keys Editor (IBM DB2 LUW)
- Functions Editor (IBM DB2 LUW)
- Groups Editor (IBM DB2 LUW)
- Indexes Editor (IBM DB2 LUW)
- Materialized Query Tables Editor (IBM DB2 LUW)

- Packages Editor (IBM DB2 LUW)
- Nodegroups 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
Properties	In addition to the owning **Schema**, this tab lets you view the **Target Owner**, **Target Type**, and **Target Name**, of the referenced object.
Comment	For details on using this tab, see Adding a comment to an object.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Bufferpool Editor (IBM DB2 LUW)

The Bufferpool editor lets you manage the basic properties, nodegroups, and partitions size of a bufferpool.

>

To edit a bufferpool:

1 Open an editor on the bufferpool. 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	**Deferred** - controls IMMEDIATE or DEFERRED update of the bufferpool. **Page size** - displays the current page size. **Automatic Size** and **Number of Pages** - Lets you specify the size of the bufferpool. You can use the **Automatic Size** option to enable resizing based on self-tuning or you can specify a fixed **Number of Pages**. **No. of Block Pages** and **Block Size** - Let you specify NUMBLOCKPAGES and BLOCKSIZE values.
Nodegroups	Displays all defined nodegroups and lets you associate nodegroups with this bufferpool by selecting the relevant check boxes.
Partitions	Displays all partitions for the bufferpool. Optionally you can click in the **Size** field of an entry to activate editing and provide a new size.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Definition	Use the **Enabled** control to enable/disable the check constraint. You can also edit the condition in the **Check Condition** box. The **Table Column** button opens a dialog that lets you select and paste column names into the check condition expression.
Comment	For details on using this tab, see Adding a comment to an object.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Placement	Select a tablespace and click the **Edit** button to open an object editor on that tablespace. For more information, see Tablespaces Editor (IBM DB2 LUW).
Options	Lets you modify the DB2 Database Manager configuration file parameters for the database. Select a parameter and click the **Edit** button to open a dialog that lets you change the current value. **NOTE:** To set database options for all future databases, set the database options on the model database.
DDL	For details on using this tab, see Viewing the SQL/DDL for an object.

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Event Monitor Editor (IBM DB2 LUW)

The event monitor editor lets you manage the event types associated with an event monitor, modify the condition for each event type, and specify where and how records are written.

To edit an event monitor:

1 Open an editor on the event monitor. 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	Lets you work with the following settings: **Auto Start** - lets you specify whether the event monitor is activated each time the database starts. **Target Type** - lets you specify where the output is written: FILE, PIPE, or TABLE. If you specify a **Target Type** of TABLE, you also have options to provide a **Buffer SIze** and specify whether the event monitor is to be **Blocked**. If you specify a **Target Type** of PIPE, you must also specify the **Target** named pipe identifier. If you specify a **Target Type** of FILE, you also have options to provide a **Buffer SIze** and whether the event monitor is to be **Blocked**. You also specify a **Write Mode** of APPEND or REPLACE, provide a **Max Files** value and a **Max File Size** value.
Definition	Lets you modify the **Event Types** used to construct the CREATE EVENT MONITOR statement FOR clause. The base set of event types is BUFFERPOOLS, CONNECTIONS, DATABASE, DEADLOCKS, DEADLOCK WITH DETAILS, STATEMENTS, TABLES, TABLESPACES, and TRANSACTIONS. For IBM DB2 for Linux, Unix, and Windows version 9.0, DEADLOCK WITH DETAILS HISTORY and DEADLOCKS WITH DETAILS HISTORY VALUES event types are available. For IBM DB2 for Linux, Unix, and Windows version 9.5, ACTIVITIES, THRESHOLD VIOLATIONS, and STATISTICS event types are available. NOTE: If your choices are all from the DEADLOCKS... types, you can also specify **Deadlock Monitoring** of **Global** or **Local**. If you select Local, specify the **Database Partitions** the event monitor is to run against.

Tab	Settings and tasks
Conditions	Available only for CONNECTION, STATEMENT, and TRANSACTION **Event Type**s, this tab lets you modify a condition used to determine which CONNECTION, STATEMENT, or TRANSACTION the event monitor will collect data for.
DDL	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Column Mapping	The existing column mapping for the foreign key is represented by selected columns in the **Main Table** and **Referenced Table** lists. Additional candidates in the **Referenced Table** list are indicated by enabled column check boxes. If necessary, use the **Table** dropdown in the **Referenced Table** group to choose a new table for this foreign key. Select or deselect columns in the **Main Table** list and **Referenced Table** list to form the referential constraint between the two tables.
Properties	Lets you modify the **Delete Rule** (NO ACTION, RESTRICT, CASCADE, SET NULL) and **Update Rule** (NO ACTION, RESTRICT) actions. This tab also lets you deselect the **Enforced** check box to specify a NOT ENFORCED option.
Comment	For details on using this tab, see Adding a comment to an object.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Properties	Displays the **Language**, **Return Type**, **External Name**, **SQL Access Level**, **Origin**, **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.

Tab	Settings and tasks
Parameters	Lets you manage function parameters. On opening, this tab shows the existing parameters. Optionally you can:
	Select a parameter from the list, modify the **Type** and if appropriate, the **Precision**, **Size**, and **Scale** 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.
Return Scalar (only available for functions created with a **Function Type** of EXTERNAL SCALAR, SOURCED or TEMPLATE and for SQL **Function Type** functions with an explicit **Return Type** of SCALAR)	Lets you change the **Type** of a return scalar and depending on the type specified, modify the **Precision**, **Scale**, and **Size** of the return scalar.
Return Columns (only available for functions created with a **Function Type** of OLEDB and EXTERNAL TABLE and for SQL **Function Type** functions with an explicit **Return Type** 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 **Type** and if appropriate, the **Precision**, **Size**, and **Scale** of the parameter. You can also specify that the return column is to be treated **As Locator**.
	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.
Body	Lets you edit the body of the function.
Comment	For details on using this tab, see Adding a comment to an object.
Dependencies	For details on using this tab, see Working with object dependencies.
Permissions	For details on using this tab, see Working with privileges and permissions.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Groups Editor (IBM DB2 LUW)

The group editor lets you manage privileges for a group.

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
Object Permissions and **System Permissions**	For details on using this tab, see Working with privileges and permissions.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 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	
Columns	Lets you manage columns that make up the index. On opening, this tab shows the existing columns. Optionally you can:	
		Change the **Sort** 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.
Properties	Lets you work with properties in the following categories:	
	Attributes	Lets you view the **Defined By** property. Lets you set the **Index Type** (UNIQUE or NONUNIQUE), specify the index as **Clustered**, and **Allow Reverse Scans**.
	Storage	Lets you set **Percent Free** and **Minimum Percent Used** settings.
Include Columns	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.
Space	Lets you view the following property groups:	
	Attributes	**Make Unique**, **System Required**, **Total Keys**, **Page Fetch Pairs**, and **Distinct Keys**
	Statistics	**Index Level**, **Cluster Ratio**, **Cluster Factor**, **Leaf Pages**, **Sequential Pages**, and **Density**
	Cards	**First Key, First 2 Keys, First3 Keys, First 4 Keys,** and **Full Keys**
Comment	For details on using this tab, see Adding a comment to an object.	
Permissions	For details on using this tab, see Working with privileges and permissions.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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	
Columns	Displays the currently defined columns for the materialized Query Table. Optionally, you can **Add**, **Insert**, or **Edit** columns for the materialized query table.	
Definition	**Base Query**	Lets you view the query that this materialized query table is based on.
	Materialized Query Table Options	Lets you specify an immediate or deferred **Refresh**, and explicitly set a Materialized Query Table as **Enabled** or **Disabled** or go with the **Default**. This tab also displays whether the Materialized Query Table is maintained by the system or by a user.
Performance	**Tablespace Placement**	Displays **Table Tablespace**, **Index Tablespace**, and **Long Tablespace** settings for the materialized query table.
	Log Options	Lets you select logging options: **Logged initially** **Data Capture** - Lets you select none or change to **Include Longvar Columns**. **NOTE:** If you select the **Data Capture** option, the table name / column name cannot be longer than 18 bytes.
	Options	Let you set the **Percent Free**, **Locksize**, **Append**, and **Volatile** settings for the materialized query table.
	Partition Columns	Use the **Add** and **Delete** buttons to manage the partition columns for the materialized query table.
Space	**Page Information**	Lets you view **Row Count** and **Num. of Overflow Rows** values.
	Row Information	Lets you view **Num. of Pages with Rows**, **Num. of Pages**, **Percent Free** values.
Comment	For details on using this tab, see Adding a comment to an object.	
Permissions	For details on using this tab, see Working with privileges and permissions.	
Dependencies	For details on using this tab, see Working with object dependencies.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Nodegroups Editor (IBM DB2 LUW)

The nodegroup editor lets you add partitions to, and delete partitions from the nodegroup.

>

To edit a nodegroup:

1 Open an editor on the nodegroup. 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
Partitions	Use the **Join Nodegroup** and **Leave Nodegroups** buttons to move partitions between the **Partitions Not in Nodegroup** and **Partitions in Nodegroup** lists.
Comment	For details on using this tab, see Adding a comment to an object.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Definition	Lets you view the following settings: **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**.
Statements	Shows the contents of a package statement. Optionally, you can select a statement and click **Explain** to copy the statement to an ISQL Editor window. For more information, see ISQL Editor.
Comment	For details on using this tab, see Adding a comment to an object.
Dependencies	For details on using this tab, see Working with object dependencies.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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	
Columns	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.
Properties	Displays the owning **Table Schema** and **Table Name**, the **Name** of the primary key, a **Defined By** property, and **Percent Free** and **Minimum Percent Used** properties.	
Space	Lets you view the following property groups:	
	Attributes	**Make Unique**, **System Required**, **Total Keys**, **Page Fetch Pairs**, and **Distinct Keys**
	Statistics	**Index Level**, **Cluster Ratio**, **Cluster Factor**, **Leaf Pages**, **Sequential Pages**, and **Density**
	Cards	**First Key, First 2 Keys, First3 Keys, First 4 Keys,** and **Full Keys**
Comment	For details on using this tab, see Adding a comment to an object.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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:

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
Properties	Lets you modify the following settings: **SQL Access Level** (MODIFIES SQL DATA, CONTAINS SQL, READS SQL DATA, NO SQL), number of **Result Sets**, **External Action**, **New Save Point**, **Inherit Special Registers**, **Deterministic**, and **Parameter CCSID**. Lets you view the following settings: **Language** (C, COBOL, JAVA, OLE, SQL CLR), **External Name**, **Threadsafe**, **Fenced**, **Parameter Type** (GENERAL, DB2, GENERAL WITH NULLS, SQL), **Program Type** (MAIN, SUB and not enabled for a **Parameter Style** of SQL), and **DBInfo**.

Tab	Settings and tasks
Parameters	Lets you manage function parameters. On opening, this tab shows the existing parameters. Optionally you can:
	Select a parameter from the list, modify the **Type**, and if appropriate, the **Precision**, **Size**, and **Scale** of the parameter. You can also set the **Parameter Mode** 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.
	NOTE: You cannot use host variables in the CALL statement for the name of the procedure.
Comment	For details on using this tab, see Adding a comment to an object.
Body	Lets you modify the body of the function.
Dependencies	For details on using this tab, see Working with object dependencies.
Permissions	For details on using this tab, see Working with privileges and permissions.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 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:

1 Open an editor on the sequence. 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	
Definition	Lets you work with settings in the following categories:	
	Parameters	**Increment By** - 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. **Minimum Value** - Lets you specify the minimum value of the sequence. This integer value can have 28 or fewer digits. **Maximum Value** - Lets you specify the maximum value the sequence can generate. This integer value can have 28 or fewer digits.
	Next Sequence Numbers	Lets you work with sequence cycle numbers.
	Sequence Datatype	Lets you specify the datatype and width for the sequence.
	Options	Lets you specify **Cache Size**, **Cycle When Reach Max/Min**, and **Generate Numbers in Order** (useful when you are using the sequence number as a timestamp) values.
Dependencies	For details on using this tab, see Working with object dependencies.	

Tab	Settings and tasks
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 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	
Properties	Displays the initialy-defined values for the following properties: **Instantiable**, **Final Type**, **With Function Access**, **Without Comparisons**, **Inline Length**, **Supertype Schema**, **Supertype Name**, **Cast (Source as Ref) With**, **Cast (Ref as Source) With**, **Reference Using**, **Size**, **Precision**, and **Scale**. For details on these properties, see Structured Type Wizard for IBM DB2 (IBM DB2 LUW).	
Attributes	Displays the currently defined attributes for the structured type. Optionally, take one of the following actions:	
		Select an attribute and change the **Datatype** in the Property/Value list. Depending on the datatype you choose you can also provide additional datatype options. For details on option availability, see Structured Type Wizard for IBM DB2 (IBM DB2 LUW).
		Add an attribute by clicking the **New** button and typing a name for the attribute.
		Select an attribute and click **Drop** to delete the attribute.

Tab	Settings and tasks
Methods	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:
	Click **Add** to begin the process of adding a new method specification. Similarly, select a method and click **Edit** to modify the method specification. For more information, see Adding or editing structured type methods.
	Select a method specification and click **Drop** to delete that method specification from the structured type.
	The **Create Body** 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: **External name** - lets you provide a 'string" or SQL identifier for a method specified with a **Language** of C, JAVA, or OLE. **Transform group** - lets you specify the transform group used when invoking the method. This setting is available for methods specified with a **Language** of C, JAVA, or OLE. **As identifier** - This setting is available for methods specified with a **Language** 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 **Body** tab. **Inherit isolation level with lock request** - 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. After specifying options, click **Create Body** to work with the CREATE METHOD statement you generated on the editor's **Body** tab.
Body	Lets you view and modify the CREATE METHOD statement generated with your choices on the **Methods** tab.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 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:

1 Open an editor on a structured type. For details, see Opening an object editor.

2 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**.

3 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 **Name**, **Specific Name**, **Language** (OLE, SQL, JAVA, C), and **SQL Access Level** (CONTAINS SQL, READS SQL DATA, NO SQL). When editing a method, these properties are for display only.

Tab		Settings and tasks	
Advanced (available after clicking the **Advanced** button on the **Properties** pane)	Lets you select or enable the following ADD METHOD options: **Deterministic**, **External Action**, **Called on NULL Input**, **DBINFO**, **Fenced**, **Allow Parallel**, **Scratchpad**, **Scratchpad Length**, **Parameter Style** (SQL, DB2GENERAL), and **Final Call**.		
Parameters	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 **Datatype** in the Property/Value list. Depending on the datatype you choose you can also provide **Precision**, **Size**, and **Scale** options.	
		Delete a selected parameter by clicking the DELETE button.	
		Order the parameters using the arrow buttons.	
Return	Lets you provide details of the return value of the method:		
	Return Datatype	Lets you specify the base **Type** of the return value. Depending on the datatype you choose you can also provide **Precision**, **Size**, and **Scale** options.	
	Cast Datatype	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 **Return Datatype** group. This feature must be **Enabled** and provide the same type options as the **Return Datatype** group.	
	An **As Locator** 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. **As Locator** applies to a LOB type in the **Cast Datatype** group, if specified. Otherwise it applies to a LOB type in the **Return Datatype** group		

4 When ready click **Finish**.

5

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	
Columns	Displays the currently defined columns in the table. Optionally, you can:	
		Select a column, and in the **Property**/**Value** list modify property values for that column.
		Click **Add Column**, provide a name for the new column, and set property values for the column.
		Select a column and click **Delete** to remove the column from the table.

Tab	Settings and tasks
Properties	Displays view **Name, Created**, **Last RunStats**, **Invalidate Time**, and **Defined By** properties. Lets you set **Percent Free**, **Lock Size**, **Append Data**, **Volatile**, **Compress**, **Row Compression**, **Security Policy**, **RestrictDrop**, **LogIndexBuild**, **CCSID**, **Do not initially log**, and **Data Capture** properties.
Status	Lets you view the following statistics: **Total Number of Rows**, **Number of Overflow Rows**, **Number of Pages With Rows**, **Number of Page**s, **Table Status**, **Row Organization**, **RowTypeSchema**, **RowTypeName**, **AccessMode**, **ActiveBlocks**, **AvgCompressedRowSize**, **AvgRowCompressionRatio**, **CodePage**, **PercentOfPagesSaved**, **LastRegeneratedTime**, and **ProtectionGranularity**.
Partition	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 **Commands** menu to attach or detach a data partition.
Tablespaces	**Data Tablespace**, **Long Tablespace**, and **Index Tablespace** — Lets you view the **Database partition Group**, **Managed By**, **Page Size**, and **Extent Size** properties. Lets you choose the **Name** of a tablespace.
Dimensions	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.
Distribution Key Columns	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.
Indexes	Lets you manage indexes for a table:
	Click **Add** to open a dialog that lets you add a new index to the table.
	Select an index and click **Edit** to open a dialog that lets you edit index properties.
	Select an index and click **Drop** to open a dialog that lets you remove the index from the table.
Constraints	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 **Add** to open a dialog that lets you add a constraint of that type.
	Select a constraint and click **Edit** to open a dialog that lets you modify the constraint details.
	Select a constraint and click **Drop** to remove the constraint.
Comment	For details on using this tab, see Adding a comment to an object.
Dependencies	For details on using this tab, see Working with object dependencies.
Permissions	For details on using this tab, see Working with privileges and permissions.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 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, DBArtisan 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, DBArtisan 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	
Properties	Lets you work with settings in the following categories:	
	Tablespace	Lets you view the **Type** (REGULAR, LARGE, TEMPORARY, or USER TEMPORARY), **Use Automatic Storage**, **Managed By**, and **Database Partition Group** properties. Lets you modify the **Buffer Pool** and **Drop Recovery** properties.
	Performance	Lets you view the **Page SIze** and **Extent Size** properties. Lets you specify or select the **Prefetch Automatic**, **Overhead**, **Transfer Rate**, and **File System Caching** properties.
	Automatic Storage	Lets you work with the **AutoResize** (enables following size settings), **Initial Size**, **Increase Size**, **Max Size Unlimited**, and **Max Size** (enabled if **Max Size Unlimited** is not set) attributes. The **Initial Size** and **Max Size** values can be provided as kB, mB, or gB values.
Container	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 **Size** controls to RESIZE, EXTEND, or REDUCE the container.
		Add a container by clicking **New**, and then in the Property/Value list, provide **Name**, number of **Database Partitions**, **Type** (DEVICE, FILE) and original **Size**.
		Click **Delete** to remove a container.
Performance	Lets you manage settings in the following categories:	
	Page Setting	Lets you view **Page SIze** and **Extent Size** settings and modify the **Prefetch Size** setting.
	I/O Setting	Let you modify the **Overhead** and **Transfer Rate** settings.
	Dropped Table Settings	Lets you view the **Recovery Status**.
	Defaults	Lets you view the **Nodegroup** and specify a **Bufferpool**.
Comment	For details on using this tab, see Adding a comment to an object.	
Space	Lets you view the table usage and the distribution of space for a tablespace. Specific statistics include **Free Pages**, **Used Pages**, **Reserved Pages**, and **Total Pages**.	
Objects	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 **Edit** to open an object editor on the selected object.
		Select an object and click **Drop** to initiate dropping the selected object.
Privileges	For details on using this tab, see Previewing and submitting object editor changes.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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:

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
Properties	Lets you view the **Trigger Timing**, **Trigger Events**, **Trigger Type**, **Object Status**, **Definer**, and **Function Path** properties.
Comment	For details on using this tab, see Adding a comment to an object.
Dependencies	For details on using this tab, see Working with object dependencies.
Definition	Lets you view and modify the CREATE TRIGGER DDL that will implement your changes. To modify a trigger, DBArtisan 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 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:

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	
Columns	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.
Properties	Displays the owning **Table Schema** and **Table Name**, the Name of the primary key, a **Defined By** property, and **Percent Free** and **Minimun Percent Used** properties.	
Space	Lets you view settings in the following categories:	
	Attributes	Lets you view **Make Unique**, **System Required**, **Total Keys**, **Distinct Keys,** and **Page Fetch Pairs** settings.
	Statistics	Lets you view **Index Level**, **Cluster Ration**, **Cluster Factor**, **Leaf Pages**, **Sequential Pages**, and **Density** settings.
	Cards	Lets you view **First Key**, **First 2 Keys**, **First 3 Keys**, **First 4 Keys** and **Full Keys** settings.
Comment	For details on using this tab, see Adding a comment to an object.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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
Properties	Displays the user **Datatype** name. Lets you select a new **Owner** or **Type**, and depending on the type specified, offers additional type-specific properties such as **Size**. Also lets you modify the **Allow Bit Data** property, if appropriate.
Comment	For details on using this tab, see Adding a comment to an object.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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	
Properties	Displays the user **Name**.	
Objects	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 **Edit** to open an object editor on the selected object.
		Select an object and click **Drop** to initiate dropping the selected object.
Object Permissions and **System Permissions**	For details on using these tabs, see Working with privileges and permissions.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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 **Column Name**, the **Datatype** (and if applicable, with the precision in parentheses), and whether or not **Nulls** are allowed for that column. You can also set the **EnableQueryOptimization** 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, DBArtisan must drop then create the view.
Comment	For details on using this tab, see Adding a comment to an object.
Dependencies	For details on using this tab, see Working with object dependencies.
Permissions	For details on using this tab, see Working with privileges and permissions.

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

IBM DB2 for z/OS Object Editors

DBArtisan includes an Object Editor for all supported IBM DB2 for z/OS objects:

- Aliases Editor (IBM DB2 Z/OS)
- Buffer Pools 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
Properties	In addition to the owning **Schema** and **Name** of the alias, this tab lets you view the **Target Owner**, **Target Type**, and **Target Name**, of the referenced object.
Comment	For details on using this tab, see Adding a comment to an object.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Buffer Pools Editor (IBM DB2 Z/OS)

The buffer pool editor lets you manage basic properties of a buffer pool and view usage details.

To edit a buffer pool:

1 Open an editor on the buffer pool. 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	
Definition	Lets you work with settings in the following areas:	
	Storage	Lets you set the **Size (VPSIZE)**, **Type (VPTYPE)**, **Hiperpool Size (HPSIZE)**, **Castout**, **Page Stealing Algorithm (PGSTEAL)** properties.
	Threshold	Lets you set the **Sequential Steal (VPSEQT)**, **Parallel Sequential (VPPSEQT)**, **Assisting Parallel Sequential (VPXSEQT)**, **Hiperpool Sequential Steal Threshold (HPSEQT)**, **Deferred Write Threshold (DWQT)**, and **Vertical Deferred Write Threshold (DDWQT)** properties
Usage Details	Displays space utilization details for each associated database.	
Privileges	For details on using this tab, see Working with privileges and permissions.	

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
Definition	Lets you view basic identification information on the check constraint: **Table Schema** and **Table Name**, the **Name** of the constraint and the date it was **Created**. You can also edit the condition in the **Check Condition** box. The **Table Columns** 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.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Properties	The **Database Creation** group lets you view the **Name** and **Group Member** properties and set the **Type** property. The **Database Attributes** group lets you view the **Encoding Scheme**, **Create Date**, and **Last Altered** properties and lets you set the **Tablespace Buffer Pool**, **Index Buffer Pool**, **Storage Group**, and **CCSID** properties.
Permissions	For details on using this tab, see Working with privileges and permissions.
Dependencies	For details on using this tab, see Working with object dependencies.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

DBRMs Editor (IBM DB2 Z/OS)

The DBRM Editor lets you view into the metadata of each DBRM, view and cross-reference all associated Plans and Packages as well as view and Explain all Package SQL Statements.

To edit a DBRM:

1 Open an editor on the DBRM. 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
Description	Displays the main contents of the DBRM header information, including the **Precompile Timestamp**, **Language**, and **Consistency Token**, as well as **String Delimiter**, **Katakana Charset**, **Mixed Character Set**, and **Decimal Point** specifiers.

Tab	Settings and tasks
Plans	Lists details of the plans containing the DBRM. Optionally, you can:
	Select a plan and click **Edit** to open the Plans editor on the plan. For more information, see Plans Editor (IBM DB2 Z/OS).
	Select a plan and click **Rebind** to initiate rebinding the plan. For more information, see Rebind Plans..
	Select a plan and click **Free** to initiate deleting the plan. For more information, see Free Plans.
Packages	Displays a cross-reference to the collections, packages, and plans where the DBRM was bound. Optionally, you can:
	Select a package and click **Edit** to open the Packages editor on the package. For more information, see Packages Editor (IBM DB2 Z/OS).
	Select a package and click **Rebind** to initiate rebinding the package. For more information, see Rebind Packages.
	Select a package and click **Free** to initiate deleting the package. For more information, see Free Packages.
Statements	Displays the information about the SQL statements in the DBRM. The information is for programmers reference only. Optionally, click **Explain** to copy the selected statement(s) to an ISQL window. The Explain output displays in the report pane. For more information, see ISQL Editor.

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
Column Mapping	The existing column mapping for the foreign key is represented by selected columns in the **Main Table** and **Referenced Table** lists. Additional candidates in the **Referenced Table** list are indicated by enabled column check boxes. If necessary, use the **Table** dropdown in the **Referenced Table** group to choose a new table for this foreign key. Select or deselect columns in the **Main Table** list and **Referenced Table** list to form the referential constraint between the two tables.
Properties	Lets you specify a **Delete Rule** (CASCADE, NO ACTION, RESTRICT, SET NULL).
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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	
Properties	Lets you work with properties in the following categories:	
	Identification	Lets you view the **Create Timestamp**, **Schema**, **Routine ID**, and **Origin** properties. Lets you set the **External Name** and **Collection ID** properties.
	Run Time	Lets you view the **Result Sets** property. Lets you set the **WLM Environment**, **WLM For Nested**, **ASU Time**, and **Run Options** properties.
	Structure	Lets you view the **Language**, **Parameter Style**, **Number of LOB Columns**, **Number of Parameters**, and **Allow Parallel** properties. Lets you set the **Program Type**, **Security Type**, **SQL Access Level**, **Inherit Special Registers**, **Fenced**, **DBINFO**, **Deterministic**, **Called On Null Input**, **External Action**, **Stay Resident**, **Final Call**, **Scratchpad**, and **Scratchpad Length** properties.
	Run Estimates	Lets you view the **Initial I/Os**, **I/Os Per Invocation**, **Initial Instructions**, and **Instructions Per Invocation** properties.
	Java Structure	Lets you view the **Java Class**, **Jar ID**, **Package ID**, **Method ID**, **Jar Schema**, and **Java Signature** properties.
Parameters, **Return Columns**, and **Return Scalar**	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 **Attributes** area, modify values, as permissible. Attributes differ by the tab you chose but typically include items such as **Type**, **Precision**, **Scale**, **As Locator**, and **Size**.
		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 **Drop** to delete that parameter.
Body	Lets you modify the code in the text area of the tab.	
Comment	For details on using this tab, see Adding a comment to an object.	
Permissions	For details on using this tab, see Working with privileges and permissions.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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:

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	
Columns	Lets you manage the columns or key-expressions that make up the index. On opening, this tab shows the existing columns/key-expressions. **NOTE:** Key-expression functionality is only enabled if the **Index on Expression** setting is enabled on the **Properties** tab. Optionally you can:	
		Change the **Sort** order of a column or key-expression.
		If **Index on Expression** is disabled, click the New button and select a column name from the dropdown to add a column to the index.
		If **Index on Expression** is enabled, click the New button and type a valid key-expression in the **Expression** 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.
Properties	Lets you work with the **Buffer Pool**, **Piece Size**, **Close**, **Copy**, **Compress**, **Index on Expression**, and **Padded** properties. For more information on these properties, see Indexes - Properties.	
Storage	Lets you select a dataset management scheme and provide associated attribute values. For details, see Indexes - Storage.	
Partitions	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 delate settings. Optionally, you can select a partition and click the Edit button top open an editor that lets you edit that partition.	
Space	Lets you view the following property groups:	
	Attributes	**Make Unique**, **System Required**, **Total Keys**, **Page Fetch Pairs**, and **Distinct Keys**
	Statistics	**Index Level**, **Cluster Ratio**, **Cluster Factor**, **Leaf Pages**, **Sequential Pages**, **Density**, **DASD Storage**, and **Data Blocks/Key**.
	Cards	**First Key** and **Full Keys**
Comment	For details on using this tab, see Adding a comment to an object.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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:

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
Properties	Lets you work with the following properties: **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**.
Bind Parameters	Lets you set the **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** properties.
Plan/Packlists	Displays the plans contained in the package if the package was bound individually, or as part of a complete collection ID.
Statements	Shows the contents of any package statement on the datasource. Optionally, select a statement and click **Explain** to copy the statement to an ISQL Editor window. NOTE: For more information, see ISQL Editor.
Dependencies	For details on using this tab, see Working with object dependencies.
Privileges	For details on using this tab, see Working with privileges and permissions.
Environments	Displays run-time environments information for a package. Use the arrow buttons to move environments between the **Enabled Environments** and **Disabled Environments** lists. To modify a connection, select the environment and click Edit Connections. For more information, see Connection Editor.
Command	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
Connections	Displays the connections for the plan or package.
Add	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	
Definition	Lets you work with the following attribute groups:	
	Properties	Lets you view the **Timestamp**, **Group Member**, **Plan Size**, **Average Size**, **Valid, Operative**, **Pkg. List Entries**, **SysEntries**, and **Release Bound** properties. Lets you set the **Owner**, **Qualifier**, and **Current Server** properties.
	Bind Parameters	Lets you set the **Schema Path**, **Explain**, **Validate**, **Degree**, **Dynamic Rules**, **Lock Isolation**, **Resource Acquire**, **Resource Release**, **Disconnect**, **Current Data**, **DB Protocol**, **Reoptvar**, **SQL Rules**, **Keep Dynamic**, **Defer Prepare**, **Encoding**, **Page Writes**, **Flag**, **Optimization Hint**, and **Cache Size** properties.
DBRMs	Lists details for each DBRM associated with a plan. Optionally, select a DBRM from the list and click **Edit** to open the DBRM editor on that DBRM. For more information, see DBRMs Editor (IBM DB2 Z/OS).	
Packages	Displays details for each package associated with a plan. Optionally, you can:	
		Select a package and click **Edit** to open the Package editor on that package. For details, see Packages Editor (IBM DB2 Z/OS).
		Select a package and click **Rebind** to open the Rebind Packages dialog.
		Select a package and click **Free** to open the Free Packages dialog.
DBRM/Packages	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 **Edit** to open the Package editor on that package. For details, see Packages Editor (Oracle).
		Select a package and click **Rebind** to open the Rebind Packages dialog.
		Select a package and click **Free** to open the Free Packages dialog.
Dependencies	For details on using this tab, see Working with object dependencies.	
Privileges	For details on using this tab, see Working with privileges and permissions.	
Environments	Displays run-time environments information for a plan. Optionally, you can:	
		Use the arrow buttons to move environments between the **Enabled Environments** and **Disabled Environments** lists.
		Select an environment and click **Edit Connections** to open the Connection editor.
Command	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	
Column	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 **Sort** 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.
Properties	Lets you set **Buffer Pool**, **Close**, **Copy** and **Piece Size** properties.	
Storage	Lets you view **Storage Group** and **VCAT Catalog** properties. Lets you set **Primary Space Allocation**, **Secondary Space Allocation**, **Erase**, **Free Page, Percent Free,** and **GBP Cache** properties.	
Partitions	Lets you work with partitions for the primary key.	
Space	Lets you view values in the following property groups:	
	Attributes	**Make Unique**, **System Required**, **Total Keys**, **Page Fetch Pairs**, and **Distinct Keys**.
	Statistics	**Index Level**, **Cluster Ratio**, **Cluster Factor**, **Leaf Pages**, **Sequential Pages**, and **Density**.
	Cards	**First Key** and **Full Keys**.
Comment	For details on using this tab, see Adding a comment to an object.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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	
Properties	Lets you work with properties in the following categories:	
	Identification	Lets you view the **Schema** and **Routine ID** properties. Lets you set the **External Name** and **Collection ID** properties.
	Run Time	Lets you set the **Result Sets**, **WLM Environment**, **WLM For Nested**, **ASU Time**, and **Run Options** properties.
	Structure	Lets you view the **Language**, **Number of LOB Columns**, **Number of Parameters**, and **DBINFO** properties. Lets you set the **Program Type**, **Security Type**, **SQL Access Level**, **Parameter Style**, **Inherit Special Registers**, **Fenced**, **Commit on Return, Deterministic**, and **Stay Resident** properties.
	Run Estimates	Lets you view the **Initial I/Os**, **I/Os Per Invocation**, **Initial Instructions**, and **Instructions Per Invocation** properties.
	Java Structure	Lets you view the **Java Class**, **Jar ID**, **Package ID**, **Method ID**, **Jar Schema**, and **Java Signature** properties.
Parameters	Displays a listing of the existing input/output parameters for the procedure. For the selected parameter, the **Datatype** list shows details for that parameter, including the **Type** and **Parameter Mode** (INPUT, OUTPUT, INPUT_OUTPUT). Depending on the type other parameters such as **Size**, **Precision**, or **Scale** may be available for viewing or modification. Optionally you can:	
		Select a parameter and in the **Datatype** 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 **Datatype** list.
		Click the Delete button to drop a selected parameter.
		Use the arrow buttons to change the order location of a selected parameter.
Comment	For details on using this tab, see Adding a comment to an object.	
Body	Lets you modify the SQL code for procedures on the current datasource	
Permissions	For details on using this tab, see Working with privileges and permissions.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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
Volume Devices	Shows details on volumes associated with the stogroup. Optionally you can:
	Click **Add** to open a dialog that lets you add volumes.
	Select a volume and click **Remove** to delete the volume.
Privileges	For details on using this tab, see Working with privileges and permissions.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Properties	Lets you view the owning **Schema**, the **Name** of the synonym, and the synonym's **Referenced Object Owner**, **Referenced Object Type**, and the **Referenced Object Name**.
Dependencies	For details on using this tab, see Working with object dependencies.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Columns	Displays the currently defined columns in the table. For each column, the Property/Value list displays the **Name** and **Type** for the column. In addition, depending on the **Type** selected, the list also displays **Size**, **Scale**, **Identity Column**, **Allow Nulls**, **Default Value**, **Comment**, and **For Data** property values, as appropriate. The **Statistics** group displays a **Number of Distinct Values in the Column** value. Optionally, you can:	
		Select a column, and in the **Property**/**Value** list modify property values for that column.
		Click **Add Column**, provide a name for the new column, and set property values for the column.
		Select a column and click **Delete** to remove the column from the table.
Properties	Lets you work with settings in the following categories:	
	Table	Lets you view **EditProcedure**, **Table Type**, **Table Status**, **Check Flag**, **DBID**, and **OBID** properties. Lets you set **Volatile**, **Audit**, **RestrictDrop**, **Label**, and **ValidProc** properties.
	Tablespace Placement	Lets you select a **Tablespace**.
	Log Options	Lets you select a **Data Capture** option of DATA CAPTURE NONE or DATA CAPTURE CHANGES.
	Statistics	Lets you view the following statistics: **Last Runstats**, **Total number of rows**, **Average Row Length**, **Number of Pages**, **Percent Compressed Rows**, **Max Record Length**, and **DASD storage**.
Partitions	Displays existing partition columns and data partitions. Optionally you can add, edit, or delete partition columns and data partitions.	
Indexes	Displays the list of indexes for the table. Optionally, take one of the following actions:	
		Click **Add** to open a dialog that lets you add a new index to the table.
		Select an index and click **Edit** to open a dialog that lets you edit index properties.
		Select an index and click **Drop** to open a dialog that lets you remove the index from the table.
Constraints	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 **Add** to open a dialog that lets you add a constraint of that type.
		Select a constraint and click **Edit** to open a dialog that lets you modify the constraint details.
		Select a constraint and click **Drop** to remove the constraint.
Comment	For details on using this tab, see Adding a comment to an object.	
Dependencies	For details on using this tab, see Working with object dependencies.	
Permissions	For details on using this tab, see Working with privileges and permissions.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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, DBArtisan 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, DBArtisan 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:

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	
Properties	Lets you work with properties in the following categories:	
	Bufferpool	**Buffer Pool**
	Partitions and size	**Number of partitions**, **Partition size (DSSIZE)**, **Segment size**, and **Max rows per page**
	Other parameters	**GBPCACHE**, **Compress**, **Track modified pages**, **Encoding scheme**, **Log**, **Define**, **Member Cluster**, **Close rule**, **Lock Size**, and **Maximum locks**
	For more information on these properties, see Tablespaces - Properties.	
Partitions	Displays a list of partitions for the tablespace. Details for each partition include the stoage group, VCAT, primary and secondary space allocations, the free space portion of each page and free spage frequency, the group buffer cache scheme, and whether modifications are tracked. Optionally, you can:	
		Select a partition and click **Edit** to open a dialog that lets you modify properties for that partition.
		Select a partition and click **Clone** to open a dialog that lets you apply the attributes of the selected partition to another partition.
Status	Lets you display CLAIMERS, LOCKS, LPL, USE, or WEPR status details	
Space	Lets you view space usage and allocation details for the tablespace.	
Objects	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 **Edit** to open an object editor on that object.
		Select an object and click **Drop** to initiate dropping that object.
Permissions	For details on using this tab, see Working with privileges and permissions.	
DDL	For details on using this tab, see Viewing the SQL/DDL for an object.	

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
Properties	Lets you view the **Trigger Timing**, **Trigger Events**, **Trigger Type**, **Object Status**, **Definer**, and **Function Path** properties.

Tab	Settings and tasks
Comment	For details on using this tab, see Adding a comment to an object.
Dependencies	For details on using this tab, see Working with object dependencies.
Definition	Lets you modify the CREATE TRIGGER body for a trigger. To modify a trigger, DBArtisan 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 uniique 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	
Column	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 **Sort** 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.
Properties	Lets you set the **Buffer Pool**, **Piece Size**, **Close**, and **Copy** properties.	
Storage	Lets you view the **Storage Group** and **VCAT Catalog** properties. Lets you set the **Primary Space Allocation**, **Secondary Space Allocation**, **Erase**, **Free Page**, **Percent Free**, and **GBP Cache** properties.	
Partition	Lets you work with partitions for the index.	
Space	Lets you view values in the following property groups:	
	Attributes	**Make Unique**, **System Required**, **Total Keys**, **Page Fetch Pairs**, and **Distinct Keys**
	Statistics	**Index Level**, **Cluster Ratio**, **Cluster Factor**, **Leaf Pages**, **Sequential Pages**, and **Density**
	Cards	**First Key** and **Full Keys**
Comment	For details on using this tab, see Adding a comment to an object.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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	
Properties	Lets you work with properties in the following categories:	
	Base Datatype	Lets you set **Type** and **Size** settings.
	Character Options	Lets you set the **For Data** and **CCSID** properties.
Comment	For details on using this tab, see Adding a comment to an object.	
Permissions	For details on using this tab, see Working with privileges and permissions.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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	
Properties	Displays the user **Name**.	
Objects	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 **Edit** to open an object editor on the selected object.
		Select an object and click **Drop** to initiate dropping the selected object.
Object Permissions and **System Permissions**	For details on using these tabs, see Working with privileges and permissions.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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
Properties	Displays all columns for the view. Details for each column include the **Column Name**, the **Datatype** (and if applicable, with the precision in parentheses), and whether or not **Nulls** are allowed for that column.
Definition	Lets you view and modify the CREATE VIEW DDL that will implement any changes you make in this editor. To modify a view, DBArtisan must drop then create the view.
Comment	For details on using this tab, see Adding a comment to an object.
Dependencies	For details on using this tab, see Working with object dependencies.
Permissions	For details on using this tab, see Working with privileges and permissions.

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Microsoft SQL Server Object Editors

DBArtisan includes an Object Editor for all supported Microsoft SQL Server objects. To see an Editor for a specific object, click the corresponding link below:

- Backup Devices Editor (SQL Server)
- Check Constraints Editor (SQL Server)
- Databases Editor (SQL Server)
- Defaults Editor (SQL Server)
- Extended Procedures Editor (SQL Server)
- Foreign Keys Editor (SQL Server)
- Functions Editor (SQL Server)
- Indexes Editor (SQL Server)
- Linked Servers Editor (SQL Server)
- Logins Editor (SQL Server)
- Primary Keys Editor (SQL Server)
- Procedures Editor (SQL Server)
- Remote Servers Editor (SQL Server)
- Roles 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)

Backup Devices Editor (SQL Server)

The Backup Device editor lets you view key properties for a backup device.

To edit a backup device:

1 Open an editor on the backup device. 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
Definition	Displays the follwing property values: **Name** - the logical name. **Device Type** - DISK or TAPE for Microsoft SQL Server 2005 or DISK, TAPE, or PIPE for Microsoft SQL Server 2000. **Skip ANSI Labels** - Available only with a **Device Type** of TAPE, indicates whether ANSI tape labels are ignored. **Physical Name** - A file name (sqltable.dat for example) or full path (a:\sqltable.dat for example) identifying the device.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Check Constraints Editor (SQL Server)

The Check Constraints Editor lets yo 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
Definition	Use the **Enabled** check box to enable or disable the check constraint. Use the **Not For Replication** check box to enable or disable the check constraint for replication. You can also edit the condition in the **Check Condition** box. The **Table Column** button opens a dialog that lets you select and paste column names into the check condition expression.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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	Lets you choose values for the **Compatible Level** and **Owner** settings. **NOTE:** 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. **NOTE:** 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.
	Properties	Lets you set the following properties: **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/pllsort**, **single user**, **subscribed**, **torn page detection**, and **trunc log on chkpt**.
Placement	Displays the currently defined data files for the database. Optionally, you can:	
		Select a file from the list on the left and in the **Device File Properties** group, modify the **Size**, **File Growth Rate**, **Max Size**, and **Unlimited Max Size** settings.
		Click the **New** button, provide a **Device File Name** for the new file, and use the other settings in the **Device File Properties** group to provide additional details for the file.
		Select a file from the list on the left and click **Delete** to remove the file.
Transaction Log	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 **Log Device Properties** group, modify the **Size**, **File Growth Rate**, **Max Size**, and **Unlimited Max Size** settings.
		Click the **New** button, provide a **Device File Name** for the new transaction log, and use the other settings in the **Log Device Properties** group to provide additional details for the transaction log.
		Select a transaction log from the list on the left and click **Delete** to remove the transaction log.
Space	Lets you view pie charts showing the data space usage and the transaction log (if available) space usage for the database.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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:

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
Properties	Lets you modify the **Value** of the Default and the owning **Schema**.
Dependencies	For details on using this tab, see Working with object dependencies.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 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
Properties	Lets you provide a dynamic-link library (DLL) **Library Name** for the extended procedure.
Dependencies	For details on using this tab, see Working with object dependencies.
Permissions	For details on using this tab, see Working with privileges and permissions.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 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
Column Mapping	The existing column mapping for the foreign key is represented by selected columns in the **Main Table** and **Referenced Table** lists. Additional candidates in the **Referenced Table** list are indicated by enabled column check boxes. If necessary, use the **Table** dropdown in the **Referenced Table** group to choose a new table for this foreign key. Select or deselect columns in the **Main Table** list and **Referenced Table** list to form the referential constraint between the two tables.
Properties	Lets you specify **Enabled** and **Not For Replication** properties. This tab also lets you select a **Delete rule** and an **Update rule** (NONE, SET NULL, SET DEFAULT, CASCADE).
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Properties	Lets you modify the owning **Schema** of the function.
Definition	Lets you view and modify the dynamic-link library (DLL) or view the data definition language (DDL).
Dependencies	For details on using this tab, see Working with object dependencies.
Permissions	For details on using this tab, see Working with privileges and permissions.

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	
Columns	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 **Sort** 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.
Properties	Lets you work with settings in the following categories:	
	Creation	In addition to displaying identifying inforrmation, you can modify the **Build Online** and **Max degree of parallelism** properties. For details on these properties, see Index Wizard (SQL Server).
	Attributes	Lets you set **Index Type**, **Clustered**, **Ignore Duplicate Key** (for **Index Type** of UNIQUE), **Statistics Recompute**, **Allow Row Locks**, and **Allow Page Locks** properties. For details on these properties, see Index Wizard (SQL Server). **NOTE:** You cannot reorganize an index (primary key, or unique key) that has an **Allow Page Locks** property set to FALSE. For information on reorganizing indexes, see Reorganizing SQL Server indexes, primary keys, and unique keys.
	Storage	Lets you set **Fill Factor**, **File Group**, **Pad Index**, and **Sort in Tempdb** settings.
Statistics	Lets you view statistics in the following categories:	
	Page Statistics	**Data Pages**, **Reserved Pages**, **Used Pages**, and **Total Pages Modified**.
	Row Statistics	**Maximum Row Size**, **Minimum Row Size**, **Max Size of Non-Leaf Index Row**, and **Total Rows Modified**.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Linked Servers Editor (SQL Server)

The Linked Servers Editor lets you manage basic properties and logins for a linked server.

>

To edit a linked server:

1 Open an editor on the linked server. 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	This panel lets you work with settings in the following categories:	
	General	**Is SQL Server** - Setting this check box indicates a linked server defined through Microsoft SQL Server. Any tables retrieved from the server are from the default database defined for the login on the linked server. Setting this check box disables the remaining OLE DB Provider settings. **Product Name** - The product name of the OLE DB datasource to which you are linking. **Datasource** - The name of the OLE DB provider managing access to the linked server. **OLE DB Provider Name** - Select the name of the OLE DB provider managing access to the specified linked server. **Provider Location** - The OLE DB location property corresponding to the linked server. **Provider Connection String** - The OLD DB provider-specific connection string that identifies a unique datasource. **Catalog** - The name of the catalog to be used when making a connection to the OLE DB provider.
	Server Options	**Data Access Server** - Indicates whether the target linked server is enabled for distributed query access. **RPC Out Server** - Lets you select the server with RPC encryption. **RPC** - Enables RPC encryption from the server. **Collation Compatible Server** - Indicates Change DIstributed Query execution against linked servers. If you select this option, Microsoft SQL Server assumes that all characters in the linked server are compatible with the local server, with regard to character set and collation sequence (or sort order). This option is for Microsoft SQL Server 8.0 or later. **Lazy Schema Validation** - This option lets you optimize performance by ensuring the query processor does not request metadata for any of the linked tables until data is actually needed from the remote member table. This option is for Microsoft SQL Server 8.0 or later. **Use Remote Collation** and **Collation Name** - Let you indicate that the linked server should use remote collation and select the name of the collation. This option is for Microsoft SQL Server 8.0 or later. **Query Timeout** - Type the query timeout value. This option is for Microsoft SQL Server 8.0 or later. **Connection Timeout** - Type the connection timeout value. This option is for Microsoft SQL Server 8.0 or later.

Tab	Settings and tasks	
Logins	Includes a table listing every user who has a login on the target linked server and their local login names. Optionally, you can:	
		Create a new login by clicking **Add linked login** button and providing the following information in the dialog that opens: **SQL Server Authentication** - Use SQL Server authenticated login credentials to login to linked server. If you choose this option, you must provide a **Local Login Name**. **No Remote Login** - Indicates that no remote login or password are required login to the linked server. **All Users** - All users will use the same login when connecting to the linked server. If you choose this option, you must provide a **Linked Login Name** and a **Password**. **Local Login Name** - Lets you select a local login name. **Linked Login Name** - Lets you type in a linked login name. **Password** - Lets you set a password.
		Select a login, click the **Modify** button, and change the login names and password.
		Select a login and click the **Drop** button to delete that login.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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:

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	
Properties	Lets you work with settings in the following categories:	
	General	Select a new **Default Database** or **Default Language**. Master should not be specified as the default database in order to prevent users from creating objects there by mistake.
	Authentication	**Password** - For STANDARD account types, you can modify the password. **Check Policy**, **Check expiration**, and **Must change** - For STANDARD account types, use these properties to modify the corresponding SQL Server Enforce password policy and Enforce password expiration properties. **Asymmetric Key** - For ASYMMETRIC KEY account types, you can select a new asymmetric key.
Users	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 **Add** button to add a selected database from the **Databases where the login does NOT have a user account** to the **Databases where the login HAS a user account** list. Conversely, remove a login from a database by using the **Remove** button to move a selected database from the **Databases where the login HAS a user account** list.	
Server Roles	Lets you select the defined roles that are to be assigned to this login.	
DDL	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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:

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	
Columns	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 **Sort** 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.
Properties	Lets you work with settings in the following categories:	
	Creation	In addition to displaying identifying inforrmation, you can modify the **Build Online** and **Max degree of parallelism** properties. For details on these properties, see Primary Key Wizard (SQL Server).
	Attributes	**Clustered** - Indicates whether the target primary key is clustered. **Ignore Duplicate Key** - Indicates whether the target primary key ignores duplicate key values. If you select his option, the transaction that generated the duplicate key values can continue. **Statistics Recompute** - Indicates that index statistics are automatically recomputed as the index is updated. Microsoft does not recommend this.
	Storage	**File Group** - Lets you select a file group. **Fill Factor -** 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. **Pad Index** - 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.
Statistics	Lets you view statistics in the following categories:	
	Page Statistics	**Data Pages**, **Reserved Pages**, **Used Pages**, and **Total Pages Modified**.
	Row Statistics	**Maximum Row Size**, **Minimum Row Size**, **Max Size of Non-Leaf Index Row**, and **Total Rows Modified**.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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:

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
Properties	Lets you change the owning **Schema**.
Definition	Lets you view the SQL code for the procedure.
Dependencies	For details on using this tab, see Working with object dependencies.
Permissions	For details on using this tab, see Working with privileges and permissions.

3. When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Remote Servers Editor (SQL Server)

The Remote Server Editor lets you manage the connection and security properties for a remote server, and manage the associated logins.

>

To edit a remote server:

1 Open an editor on the remote server. 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	Lets you work with the following settings:	
	Server Type	Displays the server type, LOCAL or REMOTE. The local server routes all commands to the remote server. To communicate with a remote server, you log into a local server and submit your remote procedure calls for remote servers. The remote server receives remote procedure calls from the local server and processes them accordingly.
	Publication Server	Specifies database publishing. When replication is installed, this option permits the tables of a database to be published for replication. When selected, this option enables publication. When not selected, this option disables publishing, drops all publications, and unmarks all transactions that were marked for replication in the transaction log.
	Distribution Server	Enables the remote server to be a distribution server.
	Subscription Server	Specifies database subscriptions. When selected, the database can be subscribed for publication. When not selected, the database cannot be subscribed for publication.
	Publisher/Subscriber	Enables the remote server to be a publisher/subscriber server.
	Collation Compatible Server	Indicates Change DIstributed Query execution against remote servers. If you select this option, Microsoft SQL Server assumes that all characters in the remote server are compatible with the local server, with regard to character set and collation sequence (or sort order).
	Data Access Server	Indicates whether the target remote server is enabled for distributed query access.
	RPC Out Server	Enables the remote server to execute stored procedures on the local server using a remote procedure call (RPC).
Logins	Includes a table listing the remote name, local name, and trusted status of logins for a remote server.Optionally you can:	
		Click **Add** to open a dialog that lets you add a new login. The dialog displays or lets you provide: **Remote Server Name -** Displays the remote server name. **Remote Login Name** - Lets you specify a remote login name. **All remote users** - Select this box if you want to use this name for all remote logins. **Local Login Name** - Lets you select a local user name. **Is Trusted -** Lets you specify this remote login as a trusted login.
		Select a login, click **Modify**, and use the table above as a guide to modifying login options.
		Select a login, and click **Drop** to delete the login.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Roles Editor (SQL Server)

The Roles Editor lets you manage the users for a role.

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
Properties	Displays the **Name**, and either a **Role Type** of STANDARD and an **Authorization Owner** or a **Role Type** of APPLICATION and a **Password**.
Users	Let you manage users for the role. A user becomes associated with an application role after running the target application. Use the **Join Rule** and **Leave Role** buttons to move users between the **Users In Role** and **Users Not In Role** lists.
Object Privileges and **System Privileges**	For details on using these tabs, see Working with privileges and permissions.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Properties	Lets you specify whether the rule is **Enabled** and view the **Fire On Insert**, **Fire On Update**, **Fire On Delete**, and **Encrypted** settings.
Definition	Lets you view and modify the SQL code that will implement any changes you make using this editor.
Dependencies	For details on using this tab, see Working with object dependencies.

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	
Properties	Lets you modify the **Owner** of the schema.	
Permissions	For details on using this tab, see Working with privileges and permissions.	
Objects	Displays the specific objects owned by the schema, grouped under object type folders. Optionally, you can:	
		Select an object and click **Edit** to open an object editor on that object.
		Select an object and click **Drop** to initiate dropping the object.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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	
Columns	Displays the currently defined columns in the table. Optionally, you can:	
		Select a column, and in the **Property**/**Value** list modify property values for that column.
		Click **Add Column**, provide a name for the new column, and set property values for the column.
		Select a column and click **Delete** to remove the column from the table.
Properties	Lets you work with settings in the following categories:	
	Physical Storage	Lets you specify a **Filegroup** and **Text Image Filegroup**. This is only available for Microsoft SQL Server 7.0 and later.
	Full-Text Index	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.
	Text In Row	Lets you enable Text In Row Data and specify a limit.

Tab	Settings and tasks	
Indexes	Displays any indexes currently defined for the table. Optionally, you can:	
		Click **Add** to open an Index editor, letting you create a new index for this table. For more information, see Index Wizard (SQL Server).
		Select an index and click **Edit** to open an wizard that lets you modify that index. For more information, see Indexes Editor (SQL Server).
		Select an index and click **Drop** 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 **Online** option that lets you specify an online drop (ONLINE-ON clause).
Constraints	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 **Add** to open a dialog that lets you add a constraint of that type.
		Select a constraint and click **Edit** to open a dialog that lets you modify the constraint details.
		Select a constraint and click **Drop** to 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 **Online** option that lets you specify an online drop (ONLINE-ON clause).
Space	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.	
Dependencies	For details on using this tab, see Working with object dependencies.	
Permissions	For details on using this tab, see Working with privileges and permissions.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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
Properties	Lets you set the **Enabled** property. It also lets you view the **Trigger Timing** (AFTER, INSTEAD OF), **Fire On Insert**, **Fire On Update**, **Fire On Delete**, and **Encrypted** properties.
Definition	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. DBArtisan must drop then create the trigger.
Dependencies	For details on using this tab, see Working with object dependencies.

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	
Columns	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 **Sort** 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.
Properties	Lets you work with settings in the following categories:	
	Creation	In addition to displaying identifying inforrmation, you can modify the **Build Online** and **Max degree of parallelism** properties. For details on these properties, see Unique Key Wizard (SQL Server).
	Attributes	**Clustered** - Indicates whether the target index is clustered. **Ignore Duplicate Key** - Indicates whether the target primary key ignores duplicate key values. If you select his option, the transaction that generated the duplicate key values can continue. **Statistics Recompute** - Indicates that index statistics are automatically recomputed as the index is updated. Microsoft does not recommend this.
	Storage	**File Group** - Lets you specify the filegroup on which to place the index. This is for Microsoft SQL Server 7.0 or later. **Fill Factor -** 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. **Pad Index** - 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.
Statistics	Displays statistics in the following categories:	
	Page Statistics	**Data Pages**, **Pages Reserved**, **Used Pages**, and **Total Pages Modified**.
	Row Statistics	**Maximum Row Size**, **Minimum Row Size**, **Max Size of Non-Leaf Index Row**, and **Total Rows Modified Since Last**.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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:

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
Properties	Lets you view the **Without Login**, **Login Name**, **Name**, **User Type**, **Certificate**, **Asymmetric Key**, and **Default Schema** properties.
Roles	Lets you assign roles to a user by selecting the check boxes associated with the roles to be assigned.
Objects	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 **Edit** to open an object editor on the selected object. For information on each editor, use the links in Microsoft SQL Server Object Editors.
	Select an object and click **Drop** to initiate dropping the selected object.
Object Permissions and **System Permissions**	For details on using these tabs, see Working with privileges and permissions.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 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	
Properties	In addition to letting you change the **Schema**, this tab lets you work with settings in the following categories:	
	Base Datatype	Lets you modify the **Type** and **Allow Nulls** properties. Depending on the **Type** selected, **Precision**, **Scale**, or **Size** properties are available as appropriate.
	Bindings	**Default Binding** - Lets you select a default. For information on creating defaults, see Default Wizard (SQL Server). **Rule Binding** - Lets you select a rule. For information on creating rules, see Rule Wizard (SQL Server).

Tab	Settings and tasks	
Usage	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 **Edit** to open an object editor on the selected object.
		Select an object and click **Drop** to initiate dropping the selected object.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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	
Properties	Lets you work with the following properties:	
	Message Number	Displays the message number chosen when the user message was created.
	Severity	Lets you modify the SQL Server error severity level (001-025).
	Write to NT Event Log	Lets you specify that the message is always written to the Windows NT Event Log.
Information	Displays the currently-defined language versions for this message number. Optionally you can:	
		Create a new language version for this message number. Click the **Add new text for the user message** button, and in the dialog box that opens, select a **Language** and provide the **Message Text**.
		Select a language version of the message and click the **Modify user message text** button. This opens a dialog that lets you modify the **Language** or **Message Text**.
		Select a language version of the message and click the **Remove user message text** 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.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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
Properties	Displays a list of the columns that make up the view. Details include the column **Name**, its **Datatype**, and whether the definition allows **Null**s. This tab also lets you change the **Schema** owning the view.
Definition	Lets you view the SQL for the View.
Dependencies	For details on using this tab, see Working with object dependencies.
Permissions	For details on using this tab, see Working with privileges and permissions.

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Oracle Object Editors

DBArtisan 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)
- Directories Editor (Oracle)
- Foreign Keys Editor (Oracle)
- Functions Editor (Oracle)
- Groups 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
Definition	Use the **Enabled** control to enable/disable the check constraint. You can also edit the condition in the **Check Condition** box. The **Table Columns** 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.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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	
Columns	Use the controls in the **General Properties** group to specify a hash or index **Cluster Type** and to specify a **Key Size** for the cluster. If you specify a hash cluster type, use the controls in the **Hash Specifications** 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.	
Storage	Lets you manage storage for a cluster:	
	Data Block Storage	Each transaction that updates a data block requires a transaction entry. **Percent Free** - minimum percentage of free space in a block **Percent Used** - minimum percentage of used space in a block. **Initial transactions** - 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. **Maximum transactions** - The maximum parameter limits concurrency on a data block.
	Extents	The unit of space allocated to an object whenever the object needs more space. **Initial Extent** - The initial space extent (in bytes) allocated to the object. **Next Extent** - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required. **Minimum Extents** - The appropriate minimum extents value for the object. **Maximum Extents** - The appropriate maximum extents value for the object. **Percent Increase** - Magnifies how an object grows and, can materially affect available free space in a tablespace. Select a value in the corresponding box.
	Optionally, you can modify the **Percent Free**, **Percent Used**, or **Max Transactions** values.	
Performance	Lets you modify the following performance settings:	
	Parallel Query Option group	Lets you specify **Degrees** and **Instances** 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.
	Cache group	The **Cache** setting keeps the blocks in memory by placing it at the most recently used end. This option is useful for small lookup tables.
Space	Lets you view the following usage and the space distribution details:	
	Space Utilization group	Displays the percent of space reserved for future updates.
	FreeLists 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.
	Extents group	The unit of space allocated to an object whenever the object needs more space.
Tables	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.	
DDL	For details on using this tab, see Viewing the SQL/DDL for an object.	

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
Definition	Displays connection string information for a database link. Details include the **Public** setting, the **User** name and **Password** associated with the database link, and a **Connect String**.
DDL	For details on using this tab, see Viewing the SQL/DDL for an object.

3. When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Directories Editor (Oracle)

The Directories Editor lets you change the path for a directory and modify associated privileges.

To edit a directory object:

1. Open an editor on the directory. 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
Definition	Lets you view or modify the SQL code for the directory.
Privileges	For details on using this tab, see Working with privileges and permissions.

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
Column Mapping	The existing column mapping for the foreign key is represented by selected columns in the **Main Table** and **Referenced Table** lists. Additional candidates in the **Referenced Table** list are indicated by enabled column check boxes. If necessary, use the **Table** dropdown in the **Referenced Table** group to choose a new table for this foreign key. Select or deselect columns in the **Main Table** list and **Referenced Table** list to form the referential constraint between the two tables.
Properties	Lets you set the foreign key as **Enabled** and specify a **Delete Rule** of CASCADE or NO ACTION.

Tab	Settings and tasks
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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	
Properties	Lets you view properties in the following categories:	
	Properties	Lets you view **Status** and **Last Modified** details for the function
	Size Information	Lets you view **Source SIze**, **Parsed Size**, **Code Size**, and **Error Size** for the function
Definition	Lets you modify the CREATE OR REPLACE FUNCTION DDL for the function.	
Dependencies	For details on using this tab, see Working with object dependencies.	
Permissions	For details on using this tab, see Working with privileges and permissions.	

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Groups Editor (Oracle)

The Groups Editor lets you modify privileges for a group.

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
Object Privileges and **System Privileges**	For details on using these tabs, see Working with privileges and permissions.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Columns	Displays a listing of the columns making up the index. Optionally, you can:	
		Use the **Sort** 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.
Properties	Lets you work with settings in the following categories:	
	Attributes	Lets you view **IsValid** and **Function-Based** properties. Lets you set **Index Type** (UNIQUE, NONUNIQUE, BITMAP), **No Sort**, **Logging**, and **Reverse** properties.
	Parallel Query Option	Lets you view the **No Parallel Execution** property. Lets you set the **Parallel Degree** and **Parallel Instances** values.
Storage	Lets you work with settings in the following categories:	
	Data Block Storage	Lets you choose the **Tablespace**, and specify **Percent Free**, **Initial Transactions**, and **Max Transactions values**.
	Extents	Lets you specify **Initial Extent**, **Next Extent**, **Minimum Extents**, **Maximum Extents**, and **Percent Increase** values.
	Freelists	Lets you specify **Freelists** and **Freelist Groups** values.
	Buffer Pool	Lets you specify a buffer pool.
Space	Lets you view settings in the following categories:	
	Space Utilization	Lets you view the **Size** and number of **Blocks** settings.
	Statistics	Lets you view **Index Level**, **Distinct Keys**, **Cluster Factor**, **Leaf Blocks**, **Leaf Blks/Key** and **Data Blks/Key** properties.

Tab		Settings and tasks
Partition	If the index is not currently partitioned, you can click the **Convert To Partitioned** button to partition the index. For more information, see Partitioning Oracle indexes, primary keys, and unique keys. If the index is currently partitioned, this tab displays the following partition details:	
	Properties	Lets you view the **Locality** (Global/Local), **Alignment** (Prefixed/Non-Prefixed), **Partition Type** (RANGE or HASH), and **Subpartition type** properties.
		Click the Edit Partition button to edit partition details. or more information, see Partitioning Oracle indexes, primary keys, and unique keys.
		Click the Drop Partition button to revert to an unpartitioned index.
		Use the following **Partition Commands** menu items: Allocate Extent Analyze Coalesce Deallocate Unused Space **Mark Unusable** - opens a dialog that lets you select one or more partitions/subpartitions to be marked as unusable. **Rebuild** - opens a dialog that lets you select one or more unusable subpartitions to be rebuilt. **Split** - 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. **NOTES:** 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. You cannot split a local index partition defined on a hash or composite table. 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.
	Columns	Displays partitioning columns.
	Partition Definitions	Displays details for each partition.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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
Definition	Lets you modify the statement, function call, or procedure call in the **Job Definition** area.

Tab	Settings and tasks
Properties	Use the settings in the **Schedule** 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 **Status** group to **Enable** or **Disable** the job.
DDL	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Definition	The **Definition** 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.
Dependencies	For details on using this tab, see Working with object dependencies.
Privileges	For details on using this tab, see Working with privileges and permissions.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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	
Information	Lets you work with properties in the following categories:	
	Master	Lets you work with the **Table Name** and **Master View** properties.
	Last Refresh	Lets you work with the **Last Date** and **Errors** properties.
	Refresh Configuration	Lets you work with the **Type**, **Refresh Method**, and **Mode** properties.
	Rollback Usage	Lets you work with the Local and Master properties.
	Options	Lets you work with the **Updatable** and **Enable Query Rewrite** properties.

Tab	Settings and tasks	
Storage	Lets you work with properties in the following categories:	
	Placement	Lets you work with **Tablespace** and **Cluster** values.
	Data Block Storage	Lets you work with **Percent Free**, **Percent Used**, **Initial transactions**, and **Maximum transactions** values.
	Extents	Lets you work with **Initial Extent**, **Next Extent, Minimum Extents**, **Maximum Extents**, and **Percent Increase** values.
Performance	Lets you work with settings in the following categories:	
	Parallel Query Option	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. **Degrees** - Lets you type a value indicating the number of query server processes that should be used in the operation. **Instances** - Lets you type a value indicating how you want the parallel query partitioned between the Parallel Servers.
	Logging	Select **Logging** to create a log for all Materialized View updates.
	Cache	Select **Cache** 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.
Query	Displays the associated query.	
Partitions	**Partitioning Method**	Displays the partitioning method, including Range-Hash Composite or Range-List Composite. 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. List partitioning lets you control how rows map to partitions. You can specify a list of discrete values for the partitioning column in the description for each partition.
	Row Movement	If its key is updated, migrates the row to a new partition.
	Partitioning Columns	Displays partitioning columns.
	Subpartitioning Columns	Displays subpartitioning columns.
	Partitions	Click Add or Edit to open the Partition dialog box. Click Drop to drop a partition.
	Subpartition Template	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.
Dependencies	For details on using this tab, see Working with object dependencies.	
Privileges	For details on using this tab, see Working with privileges and permissions.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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:

1 Open an editor on the materialized view log. 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	
Information	Displays information on the log table used for replication.	
Storage	Lets you work with settings in the following categories:	
	Data Block Storage	Lets you view the associated **Tablespace** and **Initial Transactions** value, and modify the **Percent Free**, **Percent Used**, and **Max Transactions** values.
	Extents	Lets you view **Initial Extent**, **Next Extent**, **Minimum Extents**, **Maximum Extents**, and **Percent Increase** values.
	Column Filtering	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.
Performance	Lets you work with settings in the following categories:	
	Parallel Query Option	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. **Degrees** - Lets you type a value indicating the number of query server processes that should be used in the operation. **Instances** - Lets you type a value indicating how you want the parallel query partitioned between the Parallel Servers.
	Logging	Select **Logging** to create a log for all Materialized View updates.
	Cache	Select **Cache** 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.
Dependencies	For details on using this tab, see Working with object dependencies.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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:

1 Open an editor on the outline. 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	In addition to displaying basic creation and identification properties, this tab lets you select a new **Category** for the outline.
Definition	Lets you view and modify the SQL Statement associated with the outline.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Package Bodies Editor (Oracle)

The Package Bodies Editor lets you: view package body DDL.

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	
Header	Lets you modify the package header specifications.	
Body	Lets you modify the package body specifications.	
Information	Lets you work with status and size properties in the following categories:	
	Header and **Body**	Lets you view **Status**, **Created**, and **Last Modified Source SIze**, **Parsed Size**, **Code Size**, and **Error Size** details for the header and body of the package.
Dependencies	For details on using this tab, see Working with object dependencies.	
Permissions	For details on using this tab, see Working with privileges and permissions.	

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
Columns	Displays a listing of the columns making up the index. 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.
Properties	Lets you work with settings in the following categories:	
	Enforcing Index	Lets you view **User Defined**, **Index Owner**, and **Index Name** properties.
	Attributes	Lets you work with **No Sort** (only available if **Enabled** is set), **Logging** (YES, NO, or NONE), **Reverse** (disabled if **No Sort** is enabled), **Validate**, **Deferrable**, **Deferred** (IMMEDIATE or DEFERRED and only enabled if **Deferrable** is enabled), **Enabled**, **Cascade** (disabled if **Enabled** is set), **Rely** and **Update Date** properties.
Storage	Lets you work with settings in the following categories:	
	Data Block Storage	Lets you specify **Tablespace**, **Percent Free**, **Initial Transactions**, and **Max Transactions** values. **NOTE:** You should never place primary keys on the SYSTEM tablespace.
	Extents	Displays **Initial Extent**, **Next Extent**, **Percent Increase**, **Minimum Extents** and **Maximum Extents** values.
	Freelists	Lets you specify **Freelists** and **Freelist Groups** values.
	Buffer Pool	Lets you specify a buffer pool.
Space	Lets you work with settings in the following categories:	
	Space Utilization	Lets you view **Size** and **Blocks** properties.
	Statistics	Lets you view **Index Level**, **Distinct Keys**, **Cluster Factor**, **Leaf Blocks**, **Leaf Blks/Key**, and **Data Blks/Key** properties.

Tab		Settings and tasks
Partition	If the primary key is currently partitioned, this tab displays the following partition details:	
	Properties	Lets you view the **Locality** (Global/Local), **Alignment** (Prefixed/Non-Prefixed), **Partition Type** (RANGE or HASH), and **Subpartition type** properties.
		Click the Edit Partition button to edit partition details. or more information, see Partitioning Oracle indexes, primary keys, and unique keys.
		Click the Drop Partition button to revert to an unpartitioned primary key.
		Use the following **Partition Commands** menu items: Allocate Extent Analyze Coalesce Deallocate Unused Space **Mark Unusable** - opens a dialog that lets you select one or more partitions to be marked as unusable. **Rebuild** - opens a dialog that lets you select a partition to be rebuilt. **Split** - 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.
	Columns	Displays partitioning columns.
	Partition Definitions	Displays details for each partition.
	If the primary key is not currently partitioned, you can click the **Convert To Partitioned** button to partition the primary key. For more information, see Partitioning Oracle indexes, primary keys, and unique keys.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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:

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	
Properties	Lets you work with properties in the following categories:	
	Properties	Lets you view **Status**, and **Last Modified** properties.
	Size Information	Lets you view **Source Size**, **Parsed Size**, **Code Size**, and **Error Size**
Definition	Lets you modify the SQL code for a procedure.	
Dependencies	For details on using this tab, see Working with object dependencies.	
Permissions	For details on using this tab, see Working with privileges and permissions.	

3 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	
Resources	Lets you work with settings in the following categories:	
	General Limits	Lets you specify **Composite Limit** and **Private SGA** settings.
	Session Limits	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 **Sessions Per User**, **CPU Per Session**, **Logical Reads**.
	Time Limits	Lets you specify the limit on total connection time per session. Specific settings are **Connect Time** and **Idle Time**.
	Call Limits	Lets you specify the CPU time limit for a call (a parse, execute, or fetch), expressed in hundredths of seconds. Specific settings are **CPU Per Call** and **Logical Reads**.
	Login Limits	Lets you specify the number of **Failed Login Attempts** on the user account before the account is locked and the **Account Lock Time**.
	Password Limits	Lets you specify the **Lifetime**, **Reuse Time**, **Reuse Max**, **Grace Period** and a **Verify Function** for passwords.
Users	Use the **Assign** button to open a dialog that lets you assign a user to this profile or select a user from the list, and click the **Unassign** button to open a dialog prompting you to confirm that the user is to be unassigned.	
DDL	For details on using this tab, see Viewing the SQL/DDL for an object.	

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
Redo Log Members	Lets you add new members to the redo log group, edit existing members, and delete members from the redo log group.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Authentication	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 **Password** and specify whether the role is to be authenticated **Globally** or **Externally**.
User/Roles	Displays permissions for this role to logins or users. Click **Grant** to open a dialog that lets you grant this role to a login or another role. Select a role or login and click **Revoke** to revoke the role or login.
Object Privileges and **System Privileges**	For details on using this tab, see Working with privileges and permissions.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Status	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 **Tablespace**, **Size**, and **No. of Extents**.

Tab	Settings and tasks	
Storage	Lets you work with settings in the following categories:	
	Extents	The unit of space allocated to an object whenever the object needs more space. **Initial Extent** - The initial space extent (in bytes) allocated to the object. **Next Extent** - The next extent (in bytes) that the object will attempt to allocate when more space for the object is required. **Optimal Size** - optimal extent size **Minimum Extents** - The appropriate minimum extents value for the object. **Maximum Extents** - The appropriate maximum extents value for the object.
	Extent Detail	Displays extent details.
Activity	Lets you work with settings in the following categories:	
	Activity Levels	Displays **Active Transactions**, **Writes**, **Gets** and **Waits** values.
	Dynamic Sizing	Displays **High Watermark**, **Extends**, **Shrink**s, and **Wraps** values.
DDL	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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:

1 Open an editor on the sequence. 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	
Definition	Lets you work with settings in the following categories:	
	Parameters	Lets you specify increment, minimum value and maximum value settings.
	Current/Next Sequence Numbers	Lets you work with sequence cycle numbers.
	Options	Lets you specify **Cache Size**, **Cycle When Reach Max/Min**, and **Generate Numbers in Order** (useful when you are using the sequence number as a timestamp) values.
Dependencies	For details on using this tab, see Working with object dependencies.	
Privileges	For details on using this tab, see Working with privileges and permissions.	
DDL	For details on using this tab, see Viewing the SQL/DDL for an object.	

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
Properties	Displays the type, owner, name, and other details of the object referenced by the synonym.
Dependencies	For details on using this tab, see Working with object dependencies.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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:

<table>
<tr><th>Tab</th><th colspan="2">Settings and tasks</th></tr>
<tr><td rowspan="5">Columns</td><td colspan="2">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:
Datatype properties- Lets you select a type and depending on your selection, additional properties such as Scale, Width, and Unused may be available).
Allow Nulls - Select this check box to allow nulls in this column.
Encryption properties include Password, Salted, and EncryptionAlgorithm.
Default Value lets you choose among CURRENT_TIMESTAMP, USER, SYSDATE, and UID.
Comment lets you add a comment to the column.
LOB Storage settings are available for bfile, blob, clob, and nclob types. They include Segment Name, Configuration properties (Tablespace, Chunk, Percent Version, Enable Storage In Row, Cache, and Logging), and Storage properties (Initial Extent, Next Extent, Percent Increase, Minimum Extents, Maximum Extents, Free Lists, and Free List Groups).
Optionally, you can:</td></tr>
<tr><td></td><td>Select a column, and in the Property/Value list modify property values for that column.</td></tr>
<tr><td></td><td>Click Add Column, provide a name for the new column, and set property values for the column.</td></tr>
<tr><td></td><td>Select a column and click Delete to remove the column from the table.</td></tr>
</table>

Tab	Settings and tasks	
Properties	Lets you work with settings in the following categories:	
	Table	Lets you work with the following properties, corresponding to CREATE TABLE clauses: **Cache**, **Row Movement**, **Parallel Degree**, and **Parallel Instances**.
	Physical	Lets you view the **Row Organization** property. Lets you set the **Logging** and **Table Compression** properties.
Indexes	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 **Add** to open a dialog that lets you add a new index to the table.
		Select an index and click **Edit** to open a dialog that lets you edit index properties.
		Select an index and click **Drop** to open a dialog that lets you remove the index from the table.
Constraints	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 **Add** to open a dialog that lets you add a constraint of that type.
		Select a constraint and click **Edit** to open a dialog that lets you modify the constraint details.
		Select a constraint and click **Drop** to remove the constraint.
Storage	Lets you work with settings in the following categories:	
	Data Block Storage	Lets you choose the **Tablespace Name**, and specify **Percent Free**, **Initial Transactions**, and **Max Transactions** values.
	Extents	Lets you view **Initial Extent**, **Next Extent**, **Minimum Extents**, **Maximum Extents**, and **Percent Increase** values.
	Freelists	Lets you specify **Freelists** and **Freelist Groups** values.
	Bufferpool	Lets you specify a **Buffer Pool**.
IOT Properties	Lets you work with settings in the following categories:	
	Ungrouped	Lets you specify an overflow segment.
	Percent Threshold	Lets you specify the percentage of space reserved for an index-organized table.
	Key Compression	Lets you enable or disable compression and provide a compression value.
Space	Lets you work with settings in the following categories:	
	Space Utilization	Lets you view **Size** and number of **Blocks**.
	Row Information	Lets you view the **Number of Rows**, **Average Row Length**, and **Chain Rows** properties.
	Extents	Let you view the **Number of Extents** and **Maximum Extents** values.
LOB columns	Lets you work with settings in the following categories:	
	LOB Column Segment and **LOB Column Index**	Lets you view the **Segment Name**, **No. of Extents**, **Max Extents**, **Size**, and **Blocks** properties.

Tab	Settings and tasks	
	Segment Extents and **Index Extents**	Lets you view the **Extent ID**, **File ID**, **Block ID** and number of **Blocks** properties for segment and index extents.
Partition	Lets you work with table partitions	
Comment	For details on using this tab, see Adding a comment to an object.	
Dependencies	For details on using this tab, see Working with object dependencies.	
Permissions	For details on using this tab, see Working with privileges and permissions.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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	
Datafile	The tab lists details for each datafile on the tablespace. Optionally you can **Add** a datafile or **Edit** a selected datafile. For more information, see Adding or editing tablespace datafiles.	
Storage	Lets you work with settings in the following categories:	
	Extents	Lets you view **Initial Extent**, **Next Extent**, **Minimum Extent Size**, **Minimum Extents**, **Maximum Extents**, and **Percent Increase** values.
	Segment Management	Lets you view whether segment management is **Automatic** or **Manual**, and if manual, the **TableSpace Block Size**.
	Extent Details	Lists details for each extent.
	Extent Management	Lets you view whether the **Type** of extent management is LOCAL or DICTIONARY, whether the **Allocation Type** is AUTO, UNIFORM, or USER, and the **Extent Size.**
	For background information, see About Oracle tablespace storage.	
Space	Displays a **Free Space** pie chart and free space, used space and total space statistics. This tab also displays the current **Fragmentation Index** and offers the option to coalesce the tablespace	
Map	The **Map** 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.	
Objects	Displays the objects currently stored on the tablespace, group under object type folders. Optionally you can:	
		Specify a logging option using the **Log Changes When Scheme/Data is Modified?** radio set.
		Select an object under one of the object folders and click **Edit** to open an object editor on that object.

Tab	Settings and tasks	
Quotas	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 **Quotas** tab of the Tablespace editor lets you manage user space quotas for tablespaces on the current datasource. Optionally, you can:	
		Click **Add** or select a user and click **Edit** to assign a user unlimited or a specific space usage quota on the tablespace. For details, see Adding or editing user tablespace quotas.
		Select an existing user and click **Drop** to delete the quota for that user
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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
User selection list (Add only)	Lets you select one or more users to assign a quota.
Quota	Lets you set a quotafor 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

DBArtisan 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**.

DBArtisan 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
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? (available when adding a datafile only)	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**. If you select **Other**, in the **Other** box, type the amount and then click the list to indicate MB or KB.

5 Click **OK**.

DBArtisan 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	
Properties	In addition to displaying basic identification and creation properties, this tab lets you work with settings in the following categories:	
	Attributes	Lets you specify whether the trigger is **Enabled**, select the **Trigger Timing** (BEFORE, AFTER), and the **Trigger Type** (ROW, STATEMENT).
	Correlation Names	Lets you provide an **Old Table Alias** and a **New Table Alias** as well as construct a **When Clause**.
	Status	Lets you view **Object Status**, **Create Date**, **Last Modified**, and **Base Object Type** properties.
	Size Information	Lets you view **Source Size**, **Parsed Size**, **Code Size**, and **Error Size** properties.
Events	Lets you select the INSERT, UPDATE, or DELETE event for the trigger.	
Column Selection	Lets you select the associated columns.	
Action	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.	
Dependencies	For details on using this tab, see Working with object dependencies.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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.

The Type Bodies Editor lets you manage the header and body of a type body and display status and size information.

To edit a type body:

1 Open an editor on the type body. 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
Header	Lets you manage the type header text for a type body.
Body	Lets you manage the type body text.
Information	Displays the status and size information for a type body
Dependencies	For details on using this tab, see Working with object dependencies.

Tab	Settings and tasks
Privileges	For details on using this tab, see Working with privileges and permissions.

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

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
Header	Lets you manage the type header text.
Body	Lets you create and modify type body text for the type. For information on creating type bodies, see Object Type Wizard (Oracle).
Information	Displays the header and body information for a type.
Dependencies	For details on using this tab, see Working with object dependencies.
Privileges	For details on using this tab, see Working with privileges and permissions.

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	
Columns	Displays a listing of the columns making up the index. Details include the **Column** name, the **Datatype**, and whether the definition allows **Nulls**. 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.
Properties	In addition to displaying basic identification information, this tablets you work with settings in the following categories:	
	Enforcing Index	Lets you view **User-Defined**, **Index Owner**, and **Index Name** settings.
	Attributes	Lets you set **No Sort**, **Logging**, **Reverse**, **Validate**, **Deferrable**, **Deferred**, **Enabled**, **Cascade**, **Rely**, and **Update Date** properties.

Tab	Settings and tasks	
Storage	Lets you work with settings in the following categories:	
	Data Block Storage	Lets you choose the **Tablespace Name**, and specify **Percent Free**, **Initial Transactions**, and **Max Transactions values**.
	Extents	Lets you specify **Initial Extent**, **Next Extent**, **Minimum Extents**, **Maximum Extents**, and **Percent Increase** values.
	Freelists	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 **Freelists** and **Freelist Groups** values.
	Buffer Pool	Lets you select a **Buffer Pool**.
Space	Lets you work with settings in the following categories:	
	Space Utilization	Displays the **Size** and number of **Blocks** properties.
	Statistics	Lets you view **Index Level**, **Distinct Keys**, **Cluster Factor**, **Leaf Blocks**, **Leaf Blks/Key** and **Data Blks/Key** settings.
	Extents	The unit of space allocated to an object whenever the object needs more space. This tab lets you view **Number of Extents** and **Maximum Extents** properties.
Partition	If the unique key is currently partitioned, this tab displays the following partition details:	
	Properties	Lets you view the **Locality** (Global/Local), **Alignment** (Prefixed/Non-Prefixed), **Partition Type** (RANGE or HASH), and **Subpartition type** properties.
		Click the Edit Partition button to edit partition details. For more information, see Partitioning Oracle indexes, primary keys, and unique keys.
		Click the Drop Partition button to revert to an unpartitioned unique key.
		Use the following **Partition Commands** menu items: Allocate Extent Analyze Coalesce Deallocate Unused Space **Mark Unusable** **Rebuild** **Split**
	Columns	Displays partitioning columns.
	Partition Definitions	Displays details for each partition.
	If the unique key is not currently partitioned, you can click the **Convert To Partitioned** button to partition the unique key. For more information, see Partitioning Oracle indexes, primary keys, and unique keys.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 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	
Properties	Lets you work with settings in the following categories:	
	Defaults	Lets you select a **Default Tablespace** and Temporary **Tablespace** as well as select the **Profile** used by this user. For more on creating profiles, see Profile Wizard (Oracle).
	Account	**Identified By** - Lets you select among REQUIRED YES, REQUIRED EXTERNAL, and REQUIRED_GLOBAL. **Password** - (available with **Identified By** value of REQUIRED_YES only) lets you specify a password. **External Name** - (available with **Identified By** value of REQUIRED_GLOBAL only) lets you specify the name of the user in the enterprise directory service. **Account Locked** - lets you lock this user account. **Password Expired** - (available with **Identified By** value of REQUIRED_YES only) lets you mark the password as expired, forcing the user to change their password before being allowed to connect.
Role	Lets you select the roles that are to be assigned to this user.	
Quota	Lets you set tablespace quotas for the user. Set an unlimited quota by selecting a tablespace and selecting the **Unlimited** radio button. Set a specific quota by selecting a tablespace, selecting the **Other** radio button, and providing a **Quota** value in kilobytes or megabytes.	
Objects	Lets you manage database objects associated with a user. Objects are grouped within object type folders. Optionally, you can:	
		Select an object and click **Edit** to open an object editor on the selected object.
		Select an object and click **Drop** to initiate dropping the selected object.
Object Privileges and **System Privileges**	For details on using these tabs, see Working with privileges and permissions.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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
Properties	Displays a listing of the columns making up the view.
Definition	Lets you view and modify the DDL that will implement any changes made with the editor.
Comment	For details on using this tab, see Adding a comment to an object.

Tab	Settings and tasks
Dependencies	For details on using this tab, see Working with object dependencies.
Permissions	For details on using this tab, see Working with privileges and permissions.

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
Properties	In addition to displaying basic identification and ownership properties, this pane lets you work with the following settings: **Locality** - GLOBAL or LOCAL. **Alignmen**t - displays whether this partitioned object is prefixed or non-prefixed. **Partition Type** - RANGE or HASH. **Subpartition Type** - displays the subpartition type.
Columns	For each column to be added, click New to add a column, select a **Column** name, and specify a **Sort** order.
Range Definitions (only available for a **Partition Type** of RANGE)	Lets you create an ordered list of partitions. For each partition, click New to open the **Add Partition Definition** wizard. On the **Partition Definition** pane, provide a **Name**, select a **Tablespace**, and enable or disable **Logging**. On the **Storage** pane, provide **Data Block Storage**, **Extents**, and **Freelists** property values. Click **Add**. The **Range Definitions** pane also lets you edit and remove partitions.
Hash Definitions (only available for a **Partition Type** of HASH)	Lets you specify a Partition method: **None** **Number of Partitions** - Specify a **Number of Partitions**, and then for each partition, click the New button and select a tablespace. **By Partition Name** - For each partition, click New, specify a **Name** and select a **Tablespace**.

4 When ready, click **Finish**.

>

Sybase ASE Object Editors

DBArtisan 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)
- Database Devices Editor (Sybase)
- Databases Editor (Sybase)
- Data Caches Editor (Sybase)
- Defaults Editor (Sybase)
- Dump Devices 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)
- Remote Servers 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)
- User Messages 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
Properties	Lets you change the user referenced by the alias.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Definition	Lets you edit the condition in the **Check Condition** box. The **Table Columns** 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.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Database Devices Editor (Sybase)

The Database Devices Editor lets you view or modify basic properties of a database device and view space usage details.

>

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	
Properties	Lets you work with settings in the following categories:	
	Attributes	Displays values for the **Physical Name**, **Device Number**, **Controller Number**, **Virtual Address**, and **Description** settings. **Default Device** - lets you specify whether the database device is the default device. CREATE and ALTER DATABASE commands that do not specify a location use the default database device. **Dsync** and **DirectIO** - let you modify DISK INIT and DISK REINIT parameters.
	Autogrowth	Lets you enable/disable **Auto Growth** and if enabled, to specify **Growth rate** and **Maximum Size** values.
	Page range	The **First Virtual Page** and **Last Virtual Page** numbers for the target database device.
	Mirror	These controls are only enabled if mirroring is enabled on the server. **Mirrored** - enables/disables mirroring for this device. **Mirror Enabled** - enables/disables the existing mirror. **Mirror Name** - name of the existing mirror.
Space	Details include the **Device Size**, **SpaceUsed**, **FreeSpace**, and the size and usage for all **Database fragments**.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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	
Properties	Lets you work with the following categories:	
	Creation Properties	Has the **Database Name**, **Owner**, **Transfer aliases and permissions**, and **Type** properties.
	Options	Lets you work with the **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**, and **trunc log on chkpt** properties.

Tab	Settings and tasks	
Placement	Displays device fragments. Optionally you can:	
		Select a fragment to view the **Device Name**, **Size**, and **Device Type** (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.
Space	Lets you view pie charts showing the data space usage and the transaction log (if available) space usage for the database.	
Bindings	This tab is only available for databases created with a **Type** 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 **Add**, 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 **Drop**.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Data Caches Editor (Sybase)

The Data Caches Editor lets you work with a cache's basic properties, object binding, and buffer pools. After creating or deleting a cache, Sybase ASE requires that you restart the target server to activate the new cache. Until you restart the server the status of the new cache remains Pending Active/Pending Delete.

To edit a data cache:

1 Open an editor on the data cache. 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	Lets you wok with settings in the following categories:	
	General	**Type** - lets you specify a cache status of DEFAULT, MIXED, or LOGONLY. **Policy** - lets you specify a replacement policy of STRICT or RELAXED. **Status** - displays the current status of the cache: ACTIVE, PENDING ACTIVE, or PENDING DELETE. You can only change a mixed data cache if it does not have any non-log objects bound to it. **Partitions** - lets you specify the number of partitions in the cache.
	Size	**IO Buffer** - displays the IO buffer size. **Configured Size** - lets you specify the size of the cache after the server is restarted. **Run Size** - displays the size that the cache is currently using. For the default cache, this is the amount of all cache space not explicitly configured for another cache

Tab	Settings and tasks	
Objects	Lets you manage database object binding for caches on the current datasource. You can bind entire databases, indexes, tables, as well as text and image columns to a mixed cache: 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. **NOTE:** You can bind the syslogs table to a logonly cache. On opening, the currently bound objects are displayed, stored under the folders displayed on the tab. Optionally you can take ne of the following options:	
		Initiate binding another object to the cache by clicking the **Bind** button. For detailed information, see Binding objects to a data cache.
		Select a bound object and unbind the object by clicking the **Unbind** button.
		Create a caching strategy for a selected, bound index or table using the **Strategy** properties: Select **Apply to text only** or leave unselected to specify the entire table. Enable or disable the **Prefetch Strategy** feature. Select **MRU Replacement Strategy** or leave unselected to specify LRU replacement.
Pools	Lets you manage buffer pools for a cache. Sybase ASE requires that you restart the server to activate any changes that you make to the cache configuration. Information for each bufferpool on the tab includes the name and type of the **Pool**, **Configured Size** of the cache after server restart, current **Run Size** of the pool, and current **Wash Size**. **NOTE:** Wash size must be a minimum of 10 pages and a maximum of 80 percent of the number of buffers in the pool. The buffers used by the sort cannot be used by another other task on the server until the sort completes. Therefore, to avoid seriously affecting query processing if you are creating indexes while other transactions are active, you are limited to setting the number of sort buffers to 80 percent of the pool size. Optionally, take one of the following actions:	
		Use the **Add** button, provide a Configured Size and I/O Size and press RETURN to add a buffer pool
		Select a buffer pool, modify Memory Pool Attributes settings and press RETURN.
		Use the **Delete** button to delete a buffer pool.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Binding objects to a data cache

NOTE: 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.

>

To bind databases, indexes, tables, or text/image columns to a data cache:

1 Open an editor on the data cache. For details, see Opening an object editor.

2 Click the **Objects** tab, and then click the **Bind** button. The **Bind** Wizard opens.

3 Use the following tables as a guide to understanding and completing the tabs of this Wizard:

Tab	Options and descriptions	
Properties	**Objects**	Lets you select the object type (**Entire Database**, **Tables**, or **Indexes**) to be bound.
	Databases	If binding tables or indexes, lets you select the owning database.
	Bind Text or image from table objects	If binding tables, lets you specify whether text or image columns are to be bound.
Objects	Lets you select the specific objects that are to be bound.	

4 When finished click **Finish**.

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
Properties	Lets you change the **Value** of the default.
Dependencies	For details on using this tab, see Working with object dependencies.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Dump Devices Editor (Sybase)

The Dump Devices Editor lets you view basic properties of a dump device.

To edit a dump device:

1 Open an editor on the dump device. 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
Definition	Displays the **Device Type**, **Physical Name**, and **Size** for the dump device.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Properties	Lets you specify the **Library Name** for the extended procedure.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.
Dependencies	For details on using this tab, see Working with object dependencies.
Permissions	For details on using this tab, see Working with privileges and permissions.

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
Column Mappings	The existing column mapping for the foreign key is represented by selected columns in the **Main Table** and **Referenced Table** lists. Additional candidates in the **Referenced Table** list are indicated by enabled column check boxes. If necessary, use the **Table** dropdown in the **Referenced Table** group to choose a new table for this foreign key. Select or deselect columns in the **Main Table** list and **Referenced Table** list to form the referential constraint between the two tables.
Properties	Lets you specify a **Match Full** value of TRUE or FALSE.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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
Users	Lets you add users to, or remove users from the group. To move a user from one of the **Users Not In Group** or **Users In Group** lists to the other list, select the user in the list and click the **Join Group** or **Leave Group** button.
Object Privileges	For details on using this tab, see Working with privileges and permissions.
System Privileges	For details on using this tab, see Working with privileges and permissions.
DDL	For details on using this tab, see Viewing the SQL/DDL for an object.

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	
Columns	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 **Sort** 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.
Properties	Lets you work with settings in the following categories:	
	Attributes	Has the **Index Type** (UNIQUE or NONUNIQUE), **Clustered**, **Ignore Duplicate Key** (available with **Index Type** of UNIQUE), **Ignore Duplicate Rows** (available with **Clustered** enabled), **Allow Duplicate Rows** (available with **Clustered** enabled), and **Maximum Rows Per Page** properties.
	Storage	Lets you set **Reserve Page Gap**, **Segment Name**, **Fill Factor**, **Prefetch Strategy**, **MRU Replacement Strategy** settings

Tab	Settings and tasks	
Partition	Lets you work with settings in the following categories:	
	Properties	Lets you view the **Locality** (Global/Local), **Alignment** (Prefixed/Non-Prefixed), **Partition Type** (including Range-Hash Composite or Range-List Composite), and **Subpartition type** properties.
	Columns	Displays partitioning columns.
	Partition Definitions	Displays details for each partition.
Statistics	Lets you view statistics in the following categories:	
	Page Statistics	**Data Pages**, **Reserved Pages**, **Used Pages**, **OAM Page Ratio**, and **Index Page Ratio**.
	Row Statistics	**Maximum Row Size**, **Minimum Row Size**, **Max Size of Non-Leaf Index Row**, and **Maximum Rows Per Page**.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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	
Definitions	Lets you work with the following settings:	
	Full Name	**OPTIONAL:** A user-defined descriptive name for the target login. **TIP:** To identify the login, in the Full Name box, type the actual name of the user who uses the target login.
	Default Database	The database that Sybase ASE places the target user into when they log in to the target server.
	Default Language	The language that Sybase ASE uses to displays the target login's system messages.
	Password Status	The status, current, expired, of the password.
	Last Password Change	The last date and time that the target login's password changed.
	Password Expiration	Duration, in days, before passwords expire.
	Minimum Password Length	Minimum number of characters in a password.
	Maximum Login Attempts	Maximum number of login attempts before the login (account) is suspended.
	Currently Logged In	The login status of the target login.
	Account Status	Whether the login is locked or unlocked.

Tab	Settings and tasks
Users	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. 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. Use the **Add User** button to add a selected database from the **Databases where the login does NOT have a user account** to the **Databases where the login HAS a user account** list. Conversely, remove a login from a database by using the **Drop User** button to move a selected database from the **Databases where the login HAS a user account** list.
Roles	Displays roles currently assigned to this login. Optionally, you can Copying Images or Runstats Index as needed. Also, you can use the **Def. Roles** button to open a dialog that lets you add or delete roles from the set of default roles assigned to this login.
Accounting	The **Accounting** 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 **Clear Statistics** button to start a new accounting interval.
DDL	For details on using this tab, see Viewing the SQL/DDL for an object.

3 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

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	
Columns	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 **Sort** 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.
Properties	Lets you work with settings in the following categories:	
	Attributes	Lets you set the **Clustered** and **Maximum Rows Per Page** properties.
	Storage	Lets you set **Reserve Page Gap**, **Segment** (DEFAULT, LOGSEGMENT or SYSTEM), and **Fill Factor** properties.
Statistics	Lets you view statistics in the following categories:	

>

Tab	Settings and tasks	
	Page Statistics	**Data Pages**, **Reserved Pages**, **Used Pages**, **OAM Page Ratio**, and **Index Page Ratio**.
	Row Statistics	**Maximum Row Size**, **Minimum Row Size**, **Max Size of Non-Leaf Index Row**, and **Maximum Rows Per Page**.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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
Definition	Lets you modify the SQL code for a procedure.
Dependencies	For details on using this tab, see Working with object dependencies.
Permissions	For details on using this tab, see Working with privileges and permissions.

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Remote Servers Editor (Sybase)

The Remote Server Editor lets you manage the connection and security properties for a remote server, and manage the associated logins.

>

To edit a remote server:

1 Open an editor on the remote server. 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	
Information	Instead of logging directly into a remote server, you can log into your local server and execute a remote procedure call to it.	
	Network Name	The target remote server's network name.
	Local Server	The unique server name that users must supply when executing remote procedure calls.
	Time out inactive connections	Indicates if the target server uses the timeouts option to disable and enable the normal timeout code used by the local server.
	Network Password Encryptions	Indicates whether connections with a remote server are to be initiated with a client-side password encryption handshake or with the usual unencrypted password handshake sequence.
	Security Model A	Indicates if the target remote server uses the Remote Procedure Call (RPC) Security Model A. With this model, Sybase ASE does not support security services such as message confidentiality via encryption between the local and remote servers. For more information on server security, see the Sybase System Administration Guide.
	Security Model B	Indicates if the target remote server uses the RPC Security Model B. With this model, Sybase ASE gets a credential from the security mechanism and uses the credential to establish a secure physical connection with the remote server. With this model, you can choose one or more of the following security services: mutual authentication, message confidentiality, and message integrity. For more information on server security, see the Sybase System Administration Guide.
	Security Mechanism	The name for the security mechanism.
	Mutual Authentication	This option makes the local server authenticate the remote server by retrieving the credential of the remote server and verifying it with the security mechanism. With this option, the credentials of both servers are authenticated and verified. You must select Security Model B to use this option.
	Message Confidentiality	This option means that messages are encrypted when sent to the remote server, and results from the remote server are encrypted. You must select Security Model B to use this option.
	Message Integrity	This option means that messages between the servers are checked for tampering. You must select Security Model B to use this option.
Logins	Includes a table listing the remote and local names and password authentication of logins for every remote server.Optionally you can:	
		Click **Add** to open a dialog that lets you add a new login. The dialog displays or lets you provide: **Remote Server Name -** Displays the remote server name. **Remote Login Name** - Lets you specify a remote login name. Select the All box if you want to use this name for all remote logins. **Local Login Name** - Lets you select a local user name. **Trusted Remote Login -** Lets you specify this remote login as a trusted login.
		Select a login, click **Modify**, and use the table above as a guide to modifying login options.
		Select a login, and click **Drop** to delete the login.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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	
Authentication	Lets you manage role identification for a role. To identify the role, supply a password. Property categories are as follows:	
	Identification	Lets you specify a **Password.**
	Security	Lets you specify **Role Locked**.
	Options	Lets you specify **Password Expiration**, **Minimum Password Length**, and **Maximum Failed Logins**.
Logins/Roles	Displays the current permissions for this role to logins or other roles. Optionally, you can:	
		Click **Grant** to open a dialog that lets you grant this role to a login or another role.
		Select a role or login and click **Revoke**.
Object Privileges and **System Privileges**	For details on using these tabs, see Working with privileges and permissions.	
DDL	For details on using this tab, see Viewing the SQL/DDL for an object.	

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
Properties	Lets you specify a **Restriction** and select a **Type** (STANDARD_RULE, AND_ACCESS_RULE, OR_ACCESS_RULE)
Dependencies	For details on using this tab, see Working with object dependencies.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

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	
Location	Lets you **Drop** or **Extend** selected segments.	
Objects	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	
Space	Lets you view **Segment Usage**, **Distribution of Segment Space**, **AutoGrowth Properties**, and **Object Space Usage** for the segment.	
	Shows the **Hysteresis** value for the segment and for each threshold for the segment, the number of **Free Pages** and an associated **Procedure Name**. Manage thresholds as follows:	
		Click the **Add** button to specify the number of free pages and procedure name for a new segment threshold
		Select an existing threshold and click the **Edit** button to modify the threshold's procedure name value
		Select an existing threshold and click the **Drop** button to delete that threshold
Threshold	Shows the **Hysteresis** value for the segment and for each threshold for the segment, the number of **Free Pages** and an associated **Procedure Name**. Manage thresholds as follows:	
		Click the **Add** button to specify the number of free pages and procedure name for a new segment threshold
		Select an existing threshold and click the **Edit** button to modify the threshold's procedure name value
		Select an existing threshold and click the **Drop** button to delete that threshold
DDL	For details on using this tab, see Viewing the SQL/DDL for an object.	

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	
Columns	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: **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. Optionally, you can:	
		Select a column, and in the **Property**/**Value** list modify property values for that column.
		Click **Add Column**, provide a name for the new column, and set property values for the column.
		Select a column and click **Delete** to remove the column from the table.
Properties	Lets you work with properties in the following categories:	
	Physical Storage	Lets you set the **Segment Name** (DEFAULT, SYSTEM), **Maximum Rows Per Page**, **Reserve Page Gap**, and **Identity Gap** properties.
	Cache Strategy	Lets you set **MRU Replacement Strategy** and **Prefetch Strategy** properties.
	Row Locking Strategy	Lets you view the **Expected Row Size** property and set the **Lock Scheme** (ALLPAGES, DATAPAGES, ROWPAGES) property.
Indexes	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 **Add** to open a dialog that lets you add a new index to the table.
		Select an index and click **Edit** to open a dialog that lets you edit index properties.
		Select an index and click **Drop** to open a dialog that lets you remove the index from the table.
Constraints	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 **Add** to open a dialog that lets you add a constraint of that type.
		Select a constraint and click **Edit** to open a dialog that lets you modify the constraint details.
		Select a constraint and click **Drop** to remove the constraint.
Space	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.	

Tab	Settings and tasks
Partitions	Displays partition details for the table including the **Partition Type** (ROUNDROBIN, HASH, LIST, RANGE), partitioning **Columns** (name and datatype) and **Partition Definitions** (Name, values, segments). You can also add or modify partition definitions. For more information, see Partitioning Sybase tables.
Dependencies	For details on using this tab, see Working with object dependencies.
Permissions	For details on using this tab, see Working with privileges and permissions.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.

3 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:

1 Open an editor on a table. For details, see Opening an object editor.

2 On the **Partition** tab, click the **Create Partition** or **Edit Partition** button.

3 Use the following table as a guide to understanding and modifying the settings on the tabs of this wizard:

Tab	Settings and tasks
Properties	In addition to displaying basic identification and ownership properties, this pane lets you select a **Partition Type** (ROUNDROBIN, HASH, LIST, RANGE).
Columns (not available for a **Partition Type** of ROUNDROBIN)	For each column to be added, click New to add a column and select a **Column** name.
Range Definitions (only available for a **Partition Type** of RANGE)	Lets you create an ordered list of partitions. For each partition, click New to open the **Add Partition Definition** wizard. Provide a **Name**, select a **Segment**, (DEFAULT, SYSTEM, LOGSEGMENT) and click **Add**. The **Range Definitions** pane also lets you edit and remove partitions.
Partitions Definitions (only available for a **Partition Type** of HASH)	Lets you specify a Partition method: **None** **Number of Partitions** - Specify a **Number of Partitions**, and then for each partition, click the New button and select a segment. **By Partition Name** - For each partition, click New, specify a **Name** and select a **Segment**.
List Definitions (only available for a **Partition Type** of LIST)	Lets you create a list of partition definitions. For each partition, click New to open the **Add Partition Definition** wizard. Provide a **Name**, select a **Segment**, (DEFAULT, SYSTEM, LOGSEGMENT) and click **Add**. The **List Definitions** pane also lets you edit and remove partitions.

4 When ready, click **Finish**.

Triggers Editor (Sybase)

The Triggers Editor lets you work 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
Properties	Lets you specify whether the trigger is **Enabled** and lets you view the **Fire On Insert**, **Fire On Update**, and **Fire On Delete** properties.
Definition	Lets you modify the CREATE TRIGGER body for a trigger. To modify a trigger, DBArtisan must drop then create the trigger.
Dependencies	For details on using this tab, see Working with object dependencies.

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	
Columns	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 **Sort** 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.
Properties	Lets you work with properties in the following categories:	
	Attributes	Lets you set **Clustered** and **Maximum Rows Per Page** properties.
	Storage	Lets you set the **Reserve Page Gap**, and **Segment** (DEFAULT, LOGSEGMENT, SYSTEM) properties.
Statistics	Lets you view statistics in the following categories:	
	Page Statistics	**Data Pages**, **Reserved Pages**, **Used Pages**, **OAM Page Ratio**, and **Index Page Ratio**.
	Row Statistics	**Maximum Row Size**, **Minimum Row Size**, **Max Size of Non-Leaf Index Row**, and **Maximum Rows Per Page**.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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	
Properties	Lets you work with properties in the following categories:	
	Base datatype	Lets you set **Type**, **Size**, **Allow Nulls**, and **Identity** settings.
	Bindings	Lets you set **Default Binding** and **Rule Binding** properties.
Usage	Displays all objects referring to or referenced by a user datatype.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

User Messages Editor (Sybase)

The User Messages Editor lets you manage user messages and their respective object bindings.

NOTE: On Sybase, user messages are only available under the **Master** database.

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	
Information	Displays the text and language for the message and offers the option to **Edit** or **Drop** the message.	
Bindings	Displays any constraints to which the user message is bound. Bindings are grouped under folders corresponding to **Check Constraints** and **Foreign Key Constraints**. Optionally, you can:	
		Click **Bind** to open a dialog that lets you bind the message to another constraint.
		Remove existing bindings by selecting one or more bindings and clicking the **Unbind** button.
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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	
Properties	Lets you view the user **Name** and **Login Name** for the user. Lets you specify a **Group Name** for the user.	
Logins	Lets you assign logins to a user by selecting associated check boxes.	
Objects	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 **Edit** to open an object editor on the selected object.
		Select an object and click **Drop** to initiate dropping the selected object.
Object Permissions and **System Permissions**	For details on using these tab, see Working with privileges and permissions.	
Comment	For details on using this tab, see Adding a comment to an object.	
DDL View	For details on using this tab, see Viewing the SQL/DDL for an object.	

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
Properties	Displays all columns for the view. Details for each column include the **Column Name**, the **Datatype** (and if applicable, with the precision in parentheses), and whether or not **Nulls** are allowed for that column.
Definition	Displays the SQL (CREATE VIEW) for the view.
Dependencies	For details on using this tab, see Working with object dependencies.
Permissions	For details on using this tab, see Working with privileges and permissions.

3 When finished, you can submit your changes. For details, see Previewing and submitting object editor changes.

Common functions performed against objects

The table below lists common tasks performed against databse objects:

Functionality	IBM DB2 for Windows/ Unix	IBM DB2 for OS/390 and z/OS	Microsoft SQL Server	Oracle	Sybase ASE
Activate Logging	X				
Add or Modify Constraint	X	X	X	X	X
Add, Insert, or Modify Partition	X	X	X	X	X
Add or Modify Datafile	X	X	X	X	X
Add or Modify Cluster Column				X	
Allocate Extent			X	X	
Analyze			X	X	
Backup (see Data Unload Wizard for Microsoft SQL Server and Sybase ASE)			X		X
Bind Data Cache					X
Bind Package		X			
Bind Plan		X			
Build Query	X	X	X	X	X
Change Category				X	
Change Container Size	X	X	X	X	X
Change Database Comment	X				
Change Password			X	X	
Change Status		X		X	
Check Data		X			
Check Index		X			
Checkpoint			X		X
Clone Partition		X			
Clone Table		X			
Coalesce				X	
Compile				X	
Convert to Partitioned				X	
Copy Object Names	X	X	X	X	X
Copy Schema	X (version 9)				
Copy to Copy		X			
Create Alias	X	X			
Create Clone		X			
Create Insert Statements	X	X	X	X	X

>

Functionality	IBM DB2 for Windows/ Unix	IBM DB2 for OS/390 and z/OS	Microsoft SQL Server	Oracle	Sybase ASE
Create Like	X	X	X	X	X
Create or Edit User Accounts	X	X	X	X	X
Create Synonym	X	X		X	
DataLink Options		X			
DBCC			X		X
Deallocate Unused Space				X	
Detach/Attach			X		
Disable Index			X		
Disable/Enable Job Queues				X	
Disable/Enable Trigger			X	X	X
Disk Resize					X
Drop	X	X	X	X	X
Drop Clone		X			
Edit Data	X	X	X	X	X
Error	X	X	X	X	X
Estimate Size					X
Exchange Data With Clone		X			
Execute				X	X
Extract	X	X	X	X	X
Extract Data as XML			X	X	
Flush Cache	X				
Free Packages		X			
Free Plans		X			
Impact Analysis	X	X	X	X	X
Load Java				X	
Lock	X	X			
Merge Copy		X			
Migrate	X	X	X	X	X
Modify Constraint		X			
Modify Recovery		X			
Modify Statistics			X		X
Open	X	X	X	X	X
Open Control Files				X	
Partitioned Columns				X	
Partition Upper Bound				X	
Place			X		X

>

Functionality	IBM DB2 for Windows/ Unix	IBM DB2 for OS/390 and z/OS	Microsoft SQL Server	Oracle	Sybase ASE
Quiesce Database	X				
Quiesce Instance	X				
Quiesce Tablespaces	X				
Reassign by Category				X	
Rebind Packages	X	X			
Rebind Plans	X	X			
Rebuild Index		X	X	X	
Rebuild Outlines				X	
Recompile			X		X
Redistribute	X				
Refit					X
Refresh Table	X				
Reinitialize					X
Rename	X	X	X	X	X
Reorganize	X		X	X	X
Reorganize Tablespace		X			
Report	X	X	X	X	X
Restart	X				
Runstats Index		X			
Runstats Tablespace		X			
Schema	X	X	X	X	X
Select * From	X	X	X	X	X
Set Default				X	
Set Integrity	X	X			
Set Online/Offline					X
Set UNDO Tablespace				X	
Shrink			X	X	
Start Database		X			
Status			X	X	X
Stop Database		X			
Summary Definition	X				
Switch Online	X				
Transfer Ownership	X				
Truncate	X	X	X	X	X
Update Statistics	X		X		X
Unquiesce Database	X				

Functionality	IBM DB2 for Windows/ Unix	IBM DB2 for OS/390 and z/OS	Microsoft SQL Server	Oracle	Sybase ASE
Unquiesce Instance	X				
View Contents			X		

Overview of object actions/operations execution

Most commonly, initiating an object action opens a wizard with panels that guide you through initiating the action. The wizard typically opens on a first panel named either **Action Options** or using a name specific to that object action. While the remaining panels of the wizard differ by action type, the following are the most common:

- Dependencies
- Preview

Dependencies

For those object action wizards that include a **Dependencies** panel, it lists the objects potentially impacted by this object action. Within the main categories of **Referencing Objects** and **Referenced Objects**, objects are grouped within object type folders.

CAUTION: A **Dependencies** panel, if present, is provided for information only. DBArtisan does not attempt to resolve the current action in referenced or referring objects.

Preview

The Preview panel lets you inspect the SQL generated by wizards or editors when you create, edit, or perform other actions against a database object. Depending on the action being taken, this panel may contain the following controls:

Button	Description
	Lets you save the SQL.
	Lets you print the SQL.
	Lets you send the SQL via e-mail.
	Opens an SQL Editor on the SQL contained in the Preview dialog. For more information, see ISQL Editor.
	**IBM DB2 for OS/390 and z/OS ONLY:** Opens the **Job Cards for Batch Execution** dialog box. It lets you specify the job cards that should be used to execute the utility in batch. You can check the status of this utility on the **Utilities** tab of the Database Monitor. For more information, see Database Monitor.
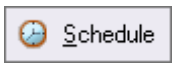	Opens a dialog that lets you select job scheduling options. Once you have selected your options, a third-party job scheduler opens. For more information, see Scheduling.

Button	Description
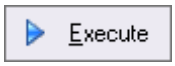	Executes the SQL.
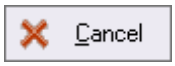	Closes the Preview dialog without executing the SQL and returns you to the wizard, editor, or dialog box from which you invoked the Preview dialog. Since you cannot edit the SQL in the Preview dialog, this lets you make changes before once again previewing and then executing the SQL.

Activate Logging

NOTE: This functionality is available for IBM DB2 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.

Completing the Activate Logging Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the target object node.

   DBArtisan displays the target objects in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Activate Logging**.

   DBArtisan opens the Active Logging dialog box.

3. Click Execute. For information on the other options, see Preview, and Scheduling.

Add or Modify Constraint

NOTE: This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Create/Modify Constraint Dialog Box lets you create or edit a constraint without knowing the underlying commands.

For more information, see Creating objects.

Step	Settings and tasks	
Definition	**Table Owner** or **Table Schema**	Select the owner of the table for which the check constraint is to be created.
	Table	Lets you select where you want to place the check 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.
	Enabled (IBM DB2 for Windows, Unix, and Linux, Microsoft SQL Server, and Oracle, only)	Enables or disables the check constraint.
	Not for Replication (Microsoft SQL Server only)	Enabling this feature ensures that the constraint is not enforced if the table is replicated for another database.

Step	Settings and tasks	
	Check Condition	Type the check constraint condition. As a convenience, use the **Table Columns** button to paste in column names as part of the condition.
Comment (IBM DB2 for Windows, Unix, and Linux only)	Lets you type a comment.	
DDL View	Preview the DDL generated by your choices. Finally, use the **Execute** button to create the check constraint.	

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 **Partition Upper Bound** dialog box.The partitioning columns you chose in Step 1 appear in the column grid, with a default upper bound of MAXVALUE. For more information, see Partition Upper Bound.
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.
Subpartition definition	None - Oracle creates one subpartition. Subpartition by name - Click Add, Insert, or Edit to open a dialog box that lets you work with subpartition properties. Click Drop or to delete the subpartition. Use the tablespace(s) specified at the partition level - Select to use the tablespace(s) specified at the partition level. Create a list of two subpartitions - Displays a list of two subpartitions. Add - Click to open a dialog box that lets you work with subpartition properties. Remove - Click to remove a selected subpartition.

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.

DBArtisan opens the **Partition Upper Bound** dialog box. For details, see Partition Upper Bound.

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.
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.**

DBArtisan establishes the upper bound.

6 Click **OK.**

DBArtisan closes the DBArtisan message.

7 When you finish setting the upper bound for the partitioning columns, click **Close.**

DBArtisan 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.**

DBArtisan 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.**

DBArtisan confirms the drop, closes the DBArtisan 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

Creating objects

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 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**.

DBArtisan closes the Add Cluster Column dialog box.

For more information, see:

Add or Modify Cluster Column

Creating objects

>

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 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.
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.

   DBArtisan 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**.

   DBArtisan adds the datafile to the tablespace definition.
6. Click **Close**.

   DBArtisan closes the Add Datafile dialog box.

For more information, see:

Add or Modify Datafile

Creating objects

Allocate Extent

NOTE: 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	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. **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.

Completing the Allocate Extent Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the target object node.

   DBArtisan displays the target objects in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Allocate Extent**.

   OR

   In the right pane of the application, right-click the target object, and then select **Allocate Extent**.

   DBArtisan 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. Click **Execute**. For information on the other options, see Preview and Scheduling.

>

For more information, see:

Allocate Extent

Creating objects

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 table 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 **Histogram Statistics** dialog box. For more information, see Histogram Statistics Dialog Box.

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.	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.

Analysis Option	Description	Additional Information
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.	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. If you choose Estimate Statistics, you can also choose to generate a histogram in **Histogram Option**. 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.
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
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.	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, DBArtisan will prompt you to create the target table after you click **Execute**.

Completing the Analyze Dialog Box

To complete this dialog box, do the following:

1 On the **Datasource Explorer**, select the target object node.

DBArtisan displays the target objects in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Analyze**.

OR

In the right pane of the application, right-click the target object, and then select **Analyze**.

DBArtisan 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 Click **Execute**. For information on the other options, see Preview and Scheduling.

For more information, see:

Analyze

Creating objects

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.

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

Creating objects

Bind Data Cache

NOTE: 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.

To bind databases, indexes, tables, or text/image columns to a data cache:

1. Open the Data Cache editor, click the **Objects** tab, and then click the **Bind** button. For details, see Data Caches Editor (Sybase). The **Bind** Wizard opens.
2. Use the following tables as a guide to understanding and completing the tabs of this Wizard:

Tab	Options and descriptions	
Properties	**Objects**	Lets you select the object type (**Entire Database**, **Tables**, or **Indexes**) to be bound.
	Databases	If binding tables or indexes, lets you select the owning database.
	Bind Text or image from table objects	If binding tables, lets you specify whether text or image columns are to be bound.
Objects	Lets you select the specific objects that are to be bound.	

3. When finished click **Finish**.

For more information, see Creating objects.

Bind Package

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS only.

The Bind Package Wizard lets you set package parameters, add environments, and set package properties.

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.

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.

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.

Option	Functionality
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.
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, either at each commit point or when the program terminates. Commit - Releases resources at each commit point. Deallocate - Releases resources only 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 package owner.

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.
Path	Lets you select a path that DB2 uses to resolve unqualified user-defined distinct types, functions, and stored procedure names (in CALL statements).

Option	Functionality
Flag	Lets you select a flags, messages to display: - All informational, warning, error, and completion messages - Only warning, error, and completion messages - Only error and completion messages - Only completion messages.

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.

Bind Plan

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS only.

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).

Panel	Settings and tasks	
	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 **Member Name**, **PDS Name** (partitioned data set) and click **Add** to enter each member and PDS name.	
3	Lets you select the **Location** to connect to, the **Collection** (location of the DBMS where the plan binds and where the description of the plan resides.) and a **Package**.	
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 **Commit** point or **Deallocate** 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.

Panel	Settings and tasks	
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.
6	**Enable** or **Disable** system connection types that can use the plan or package, select a **System** or **Cname** option.	

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

DBArtisan lets you change the category of a target Stored Outline.

NOTE: This functionality is available for Oracle only.

To change the category of a stored outline:

1. Connect to the datasource that contains the stored outline. For more information, see Datasource Management.
2. On the **Datasource Explorer**, expand the datasource nodes until the **Stored Outlines** node is visible and then select the **Stored Outlines** node. DBArtisan displays the stored outlines in the right pane of the Explorer window.
3. In the right pane of the datasource explorer, select the stored outline for which you want to change categories.
4. On the **Datasource Explorer** tool bar, click **Command**, and then select **Change Category**. DBArtisan opens the Change Category dialog box.
5. Use the following table as a guide to understanding and modifying settings in the wizard:

Step	Settings and tasks
Action options	From the **Category** dropdown, choose a new, target category.

Step	Settings and tasks
Dependencies	Review the referring and referred objects that will be automatically resolved when you execute this operation.
Preview	Preview the DDL generated for the operation and when ready, click **Execute**. For information on the other options, see Preview and Scheduling.

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.

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.

For more information, see Creating objects.

Change Database Comment

NOTE: 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.

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.

For more information, see Creating objects.

Change Password

NOTE: This functionality is available for Microsoft SQL Server, Oracle, and Sybase ASE only.

The Change Password dialog box lets you change user or login passwords, which you should do on a regular basis to guard against security leaks.

To change the password for a user or login:

1 On the **Datasource Explorer**:

- For Sybase or SQL Server, expand the **Logins** node.
- For Oracle, expand the **Users** node.

DBArtisan displays the Users or Logins in the right pane of the Explorer window.

2 Right-click the user or login, and then select **Change Password**

DBArtisan 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 Click **Execute**. For information on the other options, see Preview and Scheduling.

Change Status

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 only.

To change the status of a tablespace:

1 On the **Datasource Explorer**, expand the**Tablespaces** node.

DBArtisan 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**.

DBArtisan 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
Action options	From the Change Status dropdown, select ONLINE, OFFLINE, or READ ONLY. If you select OFFLINE, from the OfflineMode dropdown, select NORMAL, TEMPORARY, or IMMEDIATE.
Dependencies	Review the referring and referred objects that will be automatically resolved when you execute this operation.
Preview	Preview the DDL generated for the operation and when ready, click **Execute**. For information on the other options, see Preview and Scheduling.

Check Data

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS only.

The Check Data Wizard checks tablespaces for violations of referential and check constraints. The wizard also checks for consistency between a base tablespace and a LOB tablespace.

>

Upon completion, the wizard reports information about violations that are detected and the status of tablespaces that violate any constraints is set to CHECK-PENDING. Also, the CHECK-PENDING status of all tablespaces that do not violate any constraints is reset. Optionally, the utility deletes rows that violate constraints.

Important Notes

- Run the Check Data Wizard after a conditional restart or a point-in-time recovery on all tablespaces where parent and dependent tables might not be synchronized or where base tables and auxiliary tables might not be synchronized.
- Run the Check Data Wizard against a base tablespace only, not a LOB table space.

For more information, see Check Data Wizard - Panel 1.

Check Data Wizard - Panel 1

The table below describes the options and functionality on this panel of the Check Data Wizard:

Option	Description
Select tablespaces:	If Wizard is opened from the Tablespaces Editor, the target tablespaces are initially displayed. To check additional tablespaces, click the Add button.
Add	Opens theTablespace Selector Dialog Box.
Delete	Deletes selected tablespace.

Check Data Wizard - Panel 2

The table below describes the options and functionality on this panel of the Check Data Wizard:

Option	Description
Do you want the utility to be restartable?	Lets you specify to execute Wizard in restartable mode. You can restart it from the Datasource Explorer.
Would you like to specify the scope?	Lets you specify scope options: Pending (default) - DBArtisan keeps track of the data rows that must be checked to ensure referential integrity. Auxonly - Select if you only want to check tables with LOB columns. All - Select whenever the scope information is in doubt. The scope information is recorded in the DB2 catalog. The scope information can become in doubt whenever you start the target table space with ACCESS(FORCE), or the catalog is recovered to a point in time. RefOnly - Select if you want to check all dependent tables in the specified tablespaces except tables with LOB columns.
Would you like to specify the auxiliary error?	Lets you specify the auxiliary error options: Report (default) - DBArtisan sets the base tablespace to the auxiliary CHECK-pending (ACHKP) status. If CHECK DATA encounters only invalid LOB columns and no other LOB column errors, the base table space is set to the auxiliary warning status. Invalidate - DBArtisan sets the base table LOB column to an invalid status, and sets the base table space to the auxiliary warning (AUXW) status. You can use SQL to update a LOB column in the AUXW status, however, any other attempt to access the column will result in a -904 SQL return code.
Exceptions	Lets you specify the number or exceptions, from 0 to 2147483647.

Option	Description
Error ddn	Lets you specify the name of a error data definition (DD) statement that corresponds to a data control block that contains the same name. Lets you select a previously saved dataset template or click Edit to open the Dataset Template Editor.
Work ddn	Lets you specify the name of a work data definition (DD) statement that corresponds to a data control block that contains the same name. Lets you select a previously saved dataset template or click Edit to open the Dataset Template Editor.

Check Data Wizard - Panel 3

The table below describes the options and functionality on this panel of the Check Data Wizard:

Option	Description
Exceptions	Displays the Source and Destination tables for each exception. You can select one or more rows from the grid.
Add Button	Opens the Exceptions dialog box. For details, see Exceptions Dialog Box.
Delete Button	Deletes the selected exception.

Exceptions Dialog Box

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

Option	Description
Base Table Owner	Lets you select the base table owner.
Base Table	Lets you select the base table.
Exception Table Owner	Lets you select the exception table owner.
Exception Table	Lets you select the exception table.

Check Data Wizard - Panel 4

The table below describes the options and functionality on this panel of the Check Data Wizard:

Option	Description
Do you want to specify the device type for temporary data sets by DFSORT?	Lets you specify the device type for temporary data sets to be dynamically allocated by DFSORT.
Do you want to specify the sort number?	Lets you type the number of datasets to be dynamically allocated by the sort program. The range of the text field is 0 to 2147483647.

Check Index

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS only.

The Check Index Wizard lets you test whether indexes are consistent with the data they index. The Wizard issues warning messages when an inconsistency is found.

>

Run the Check Index Wizard after a conditional restart or a point-in-time recovery on all tablespaces where you suspect that the indexes may not be consistent with the data. Also run the wizard before executing the Check Data wizard to ensure that the indexes used by the Check Data Wizard are valid. This is especially important before using the Check Data with DELETE YES.

When checking an auxiliary table index, the Check Index Wizard verifies that each LOB is represented by an index entry, and that an index entry exists for every LOB.

For more information, see Check Data Wizard - Panel 1.

Check Index Wizard - Panel 1

The table below describes the options and functionality on this panel of the Check Index Wizard:

Option	Description
Select Index	If Wizard is opened from the Indexes Editor, the target tablespaces are initially displayed. To check additional Indexes, click the Add button.
Add	Opens the Indexes Selector dialog box.
Delete	Deletes selected index.

Check Index Wizard - Panel 2

The table below describes the options and functionality on this panel of the Check Index Wizard:

Option	Description
Select tablespaces	If Wizard is opened from the Tablespaces Editor, the target tablespaces are initially displayed. To check additional tablespaces, click the Add button.
Add	Opens the Tablespace Selector Dialog Box.
Delete	Deletes selected tablespace.

Check Index Wizard - Panel 3

The table below describes the options and functionality on this panel of the Check Index Wizard:

Option	Description
Do you want the utility to be restartable?	Lets you specify to execute Wizard in restartable mode. You can restart it from the Datasource Explorer.
Do you want to specify the work ddn?	Lets you specify the name of a work data definition (DD) statement that corresponds to a data control block that contains the same name. Lets you select a previously saved dataset template or click Edit to open the Dataset Template Editor.
Do you want to specify the device type for temporary data sets by DFSORT?	Lets you specify the device type for temporary data sets to be dynamically allocated by DFSORT.
Do you want to specify the sort number?	Lets you type the number of datasets to be dynamically allocated by the sort program. The range of the text field is 0 to 2147483647.

Checkpoint

Running a Checkpoint 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.

NOTE: This functionality is available for Microsoft SQL Server, and Sybase ASE only.

Important Notes

The default permission for executing a checkpoint is the db_owner fixed database role.

To run a Checkpoint against one or more databases:

1. Connect to the datasource that contains the database. For more information, see Datasource Management.
2. On the **Datasource Explorer**, select the **Databases** node.

   DBArtisan displays the Databases in the right pane of the Explorer window.
3. In the right pane of the Explorer window, select one or more target databases.
4. On the **Datasource Explorer** tool bar, click **Command**, and then select **Checkpoint**.

   DBArtisan opens the **Checkpoint** dialog box.
5. Use the following table as a guide to working through the panels of the dialog box

Step	Description
Action options	Displays the names of the database(s) you chose.
Preview	Displays the DDL generated to execute the Checkpoint operation.

6. Use one of the **Schedule** or **Execute** buttons to execute the Checkpoint.

Clone Partition

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS only.

The Clone Partition dialog box lets you clone a definition of a partition.

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.

For more information, see Creating objects.

Clone Table

NOTE: This functionality is available for Oracle.

>

The Clone Table dialog box lets you clone a definition of a table.

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.

For more information, see Creating objects.

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. It you coalesce immediately, DBArtisan 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.

DBArtisan 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**.

DBArtisan opens the Coalesce wizard.

3 Use the following table as a guide to understanding and modifying settings in the wizard:

Step	Settings and tasks
Action options	Verify display of the tablespaces to be coalesced.
Dependencies	Review the referring and referred objects that will be automatically resolved when you execute this operation.
Preview	Preview the DDL generated for the operation and when ready, click **Execute**. For information on the other options, see Preview, and Scheduling.

Compile

NOTE: This functionality is available for IBM DB2 for Linux, Unix, and Windows and Oracle only.

DBArtisan 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.

The objects you can compile and the steps in compiling them differ according to the type of datasource you are working against. See the following topics:

- Compiling IBM DB2 for Windows, Unix, and Linux procedures
- Compiling Oracle objects

Compiling IBM DB2 for Windows, Unix, and Linux procedures

The table below describes the options and functionality on the Confirm Compile dialog box for IBM DB2 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.

Completing the Compile Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the target object node.

   DBArtisan displays the target objects in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Compile**.

   OR

   In the right pane of the application, right-click the target object, and then select **Compile**.

   DBArtisan opens the Compile dialog box.

3. In the **Compile Options** section, select the appropriate options for the compilation.

4. Click **Execute**. For information on the other options, see Preview, Scheduling, and Impact Analysis.

Compiling Oracle objects

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. Prior to compiling objects of a particular type, see the relevant topic below:

- Notes on compiling Functions
- Notes on compiling Java Sources
- Notes on compiling Materialized Views
- Notes on compiling Packages and Package Bodies
- Notes on compiling Procedures

- Notes on compiling Types and Type Bodies
- Notes on compiling Views

To compile an Oracle database object:

1 Connect to the datasource that contains the database object. For more information, see Datasource Management.

2 On the left pane of the **Datasource Explorer**, select the node corresponding to the object type you want to compile.

DBArtisan displays objects of that type in the right pane of the Explorer window.

3 In the right pane of the Explorer window, select a specific database object.

4 On the **Datasource Explorer** tool bar, click **Command**, and then select **Compile**.

DBArtisan opens the **Compile** dialog box.

5 Use the following table as a guide to working through the panels of the dialog box

Step	Settings and tasks	
Action options	**Compile dependents**	If enabled, this option 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. If disabled, only the current object is compiled and the object's dependencies are ignored. This is the default setting.
	Compile only invalid dependents	If you enabled the **Compile dependents** option, enabling this option compiles only invalid dependent objects - Creates ALTER COMPILE statements for only those objects that are currently invalid.
	Compile system dependents	If you enabled the **Compile dependents** option, enabling this option compiles dependent system objects - Compiles all of the referenced objects with the debug option.
	Compile with debug	Enabling this option instructs the Oracle PL/SQL compiler to generate and store the code for use in debugging sessions.
Dependencies	Lets you review any dependecies before you proceed.	
Preview	Displays the DDL generated to execute the Checkpoint operation.	

6 Use one of the **Schedule** or **Execute** buttons to execute the Checkpoint.

Notes on compiling Functions

DBArtisan 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. DBArtisan 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.

For more information, see:

Compile

Creating objects

>

Notes on compiling Java Sources

Oracle lets you compile a Java source. Oracle resolves references to other Java sources.

Notes on compiling Materialized Views

DBArtisan 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

Creating objects

Notes on compiling Packages and Package Bodies

DBArtisan lets you recompile a package, and recompiles all package objects together.

Recompiling a package in DBArtisan compiles both the package specification and the package body by issuing two consecutive ALTER statements. However, DBArtisan 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

Creating objects

Notes on compiling Procedures

DBArtisan lets you compile a procedure that is part of a package, by compiling the package itself. DBArtisan 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 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.

>

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

Creating objects

Notes on compiling Types and Type Bodies

DBArtisan lets you recompile a type. DBArtisan recompiles both the type specification and the type body.

For more information, see:

Compile

Creating objects

Notes on compiling Views

DBArtisan lets you recompile a view when you have altered an object referenced by that view.

When you recompile a view, Oracle invalidates all dependant objects.

For more information, see:

Compile

Creating objects

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 **Datasource Explorer**, select the target object node.

DBArtisan displays the target objects in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Convert to Partitioned**.

OR

In the right pane of the application, right-click the target object, and then select **Convert to Partitioned**.

DBArtisan opens the Convert to Partitioned wizard.

3 Complete the wizard panels and then click **Finish**.

DBArtisan opens 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.

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, Insert, and Edit - Open the Add, Insert, or Modify Partition. Drop - Lets you drop a partition.

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.

Copy Object Names

NOTE: 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.

To copy an object name to be pasted into another application:

1. Connect to the datasource that contains the database object. For more information, see Datasource Management.
2. On the left pane of the **Datasource Explorer**, select the node corresponding to the object type whose name you want to copy.

   DBArtisan displays objects of that type in the right pane of the Explorer window.
3. Right-click on the specific database object and select **Copy Object Name(s)**.

Copy Schema

For IBM DB2 for Windows, Unix, and Linux (version 9), the Copy Schema function lets you quickly make copies of a database schema and its corresponding objects.

To copy a schema:

1. In the datasources pane, right-click a datasource and select **Copy Schema** from the context menu.
2. Use the following table as a guide to understanding and choosing options the options in the Copying Schema dialog:

Pane	Options
Action Options	Lets you specify **Copying Schema**, **Target Schema**, specify a **Copy Mode** of COPY, COPYNO, or DDL, specifiy a **Log filename** (COPY and COPYNO modes only), select an **Object Owner**, and **Error table schema**, provide an **Error table name**, and specify a **Drop error table after action execution** options.
Tablespace mapping	Lets you specify source and target tablespace options.
Preview	Displays the DDL generated to execute the operation.

3. **Execute** or **Schedule** the action.

Copy to Copy

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS only.

The Copy to Copy dialog box lets you make image copies. Starting with either the local primary or recovery site primary copy, the Copy to Copy dialog box can make up to three copies of the following types of copies:

- Local primary
- Local backup
- Recovery site primary
- Recovery site backup

The copies are used by the RECOVER utility when recovering a table space or index space to the most recent time or to a previous time. These copies can also be used by MERGECOPY, UNLOAD, and possibly a subsequent Copy to Copy execution.

>

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

Tab	Option	Description
Objects	**Object Grid**	Displays the target tablespaces or indexes. Your can select multiple items in the list. The Partition column is editable. You can add the partition numbers in this column, separated by commas.
	Add	For details, see Object Selector Dialog Box.
	Remove	Deleted the selected objects.
	Utility ID	Lets you type the utility id for the job. DBArtisan identifies a single job by its utility identifier, the value given for the UID parameter when the job was created. Max size is 16 characters.
Options	**From Last Image Copy**	The most recent image copy that was taken for the tablespace or index space to be the input to COPYTOCOPY job.This could be a full or incremental copy.
	From Last Full Image Copy	The most recent full image copy that was taken for the object to be the input to COPYTOCOPY job.
	From Last Incremental Image Copy	The most recent incremental image copy that was taken for the object to be the input to COPYTOCOPY job.
	From Specific Image Copy	The particular image copy data set (dsn) as the input to COPYTOCOPY job.
	Volume	Identifies the image copy data set if it is not a GDG and there is more than one copy with the same name.
	Catalog	Select for an image copy that was created as a cataloged data set.
	Volume serial number	Identifies the data set by an alphanumeric volume serial identifier of its first volume. Use this option only for an image copy that was created as a noncataloged data set.
	From sequence number	Identifies the image copy data set by its file sequence number.

>

Tab	Option	Description
Data Sets	**Local Site Primary Copy**	Lets you specify a dataset name for the local site primary image copy. This relates to the copydsn1 parameter in the stored procedure.
	Local Site Backup Copy	Lets you specify the device for the image copy. This relates to the copydevt1 parameter in the stored procedure.
	Recovery Site Primary Copy	Lets you specify the primary space allocation amount for the image copy. This relates to the copyspace1 parameter in the stored procedure.
	Recovery Site Backup Copy	Lets you specify a dataset name for the local site backup image copy. This relates to the copydsn2 parameter in the stored procedure.
	Local site backup device	Lets you specify the device for the image copy. This relates to the copydevt2 parameter in the stored procedure.
	Local site backup space	Lets you specify the primary space allocation amount for the backup image copy. This relates to the copyspace2 parameter in the stored procedure.
	Recovery site primary copy	Lets you specify a dataset name for the recovery site primary image copy. This relates to the rcpydsn1 parameter in the stored procedure.
	Recovery site primary device	Lets you specify the device for the image copy. This relates to the rcpydevt1 parameter in the stored procedure.
	Recovery site primary space	Lets you specify the primary space allocation amount for the recovery primary image copy. This relates to the rcpyspace1 parameter in the stored procedure.
	Recovery site backup copy	Lets you specify a dataset name for the recovery site backup image copy. This relates to the rcpydsn2 parameter in the stored procedure.
	Recovery site backup device	Lets you specify the device for the image copy. This relates to the rcpydevt2 parameter in the stored procedure.
	Recovery site backup space	Lets you specify the primary space allocation amount for the recovery backup image copy. This relates to the rcpyspace2 parameter in the stored procedure.

Object Selector Dialog Box

The table below describes the options and functionality on the panel of the Object Selector dialog box:

Option	Description
Database Like	Lets you type the pattern for querying the tablespaces. You can type multiple words, separated by spaces.
Tablespace Like	Lets you type the pattern for querying the tablespaces. You can type multiple words, separated by spaces.
Creator Like	Lets you type the pattern for querying the creator. You can type multiple words, separated by spaces.
Index Like	Lets you type the pattern for querying the index. You can type multiple words, separated by spaces.
Match Case	Lets you specify match case.
Object Grid	Lets you select one or more rows in this list.
Query	Repopulates grid based on the object query.
Select All	Selects all rows in Object Grid.
Unselect All	Unselects all rows in Object Grid.

For more information, see Copy to Copy

Create Alias

The Create Alias dialog box lets you create an alternate name for a table or view 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.

> **NOTE:** This functionality is available for IBM DB2 for Linux, Unix, and Windows and IBM DB2 for OS/390 and z/OS.

To create an alias for a table or view:

1. Connect to the datasource that contains the table or view. For more information, see Datasource Management.
2. On the left pane of the **Datasource Explorer**, select the table or view node.

   DBArtisan displays objects of that type in the right pane of the Explorer window.
3. In the right pane of the Explorer window, select a specific table or view.
4. On the **Datasource Explorer** tool bar, click **Command**, and then select **Create Alias**.

   DBArtisan opens a Create Alias wizard.

For information on working with Create Alias wizards, see the following topics:

- Alias Wizard for IBM DB2 (IBM DB2 LUW)
- Alias Wizard (IBM DB2 Z/OS)

Create Clone

DBArtisan 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. For more information, see Tablespace Wizard (IBM DB2 Z/OS).

To create a clone for one or more tables:

1. On the **Datasource Explorer**, expand the **Tables** node.

   DBArtisan 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**.

   DBArtisan 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
Action Options	Lets you provide a **Clone Name**.
Dependencies	Lists referring and referenced objects potentially impacted by the change. For details, see Dependencies.
Preview	Displays the DDL that will execute the object action. For details, see Preview/.

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 datasource. 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.

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.

Completing the Create Insert Statements Dialog Box

To complete this dialog box, do the following:

1. In the right pane of the application, select the **Tables** node.

   DBArtisan displays the Tables in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Create Insert Statements**.

   OR

   In the right pane of the application, right-click the table, and then select **Create Insert Statements**.

   DBArtisan 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.

When you add a script in a new file, you not only choose a name for the file, but choose among create options. Finally, when you complete your script, you can **Preview**, **Schedule** or immediately **Execute** it.

Create Like

The Create Like Editor lets you create a new object based on an existing object. The following table shows availability of the Create Like function for particular object types by DBMS:

IBM DB2 for Windows, Unix, and Linux	IBM DB2 for OS/390 and Z/OS	Microsoft SQL Server	Oracle	Sybase
Materialized Query Tables Tables	Tables	Logins Tables Users	Tables	Logins Tables Users

To create an object based on another object:

1. Connect to the datasource that contains the object you want to copy. For more information, see Datasource Management.
2. On the left pane of the **Datasource Explorer**, expand nodes until the target object type is visible and then click the taget object type.

   DBArtisan displays objects of that type in the right pane of the Explorer window.
3. In the right pane of the Explorer window, select the specific object you want to copy.
4. On the **Datasource Explorer** tool bar, click **Command**, and then select **Create Like**.

   DBArtisan opens a Create Like wizard.
5. Provide a name for the new object.
6. Modify settings on each panel of the wizard as required.

   For particular object tpes, tasks and setttings on the panels of the Create Like wizard are similar to those for the object editors for that object type. For information, see Modifying objects using editors.
7. Click **Execute** to create the new object.

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.

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.

For more information, see Creating objects.

Create Synonym

NOTE: The Create Synonym functionality is available for IBM DB2 for Linux, Unix, and Windows, OS/390, and Oracle.

The Create Synonym wizard 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.

DBArtisan lets you specify options for creating synonyms with the Create Synonym and Create Synonyms wizards. The wizard that DBArtisan returns depends on whether you are creating a synonym for a single object or are creating synonyms for multiple objects at the same time.

To create a synonym for one or more objects:

1. Connect to the datasource where you want to create a synonym. 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**, click the specific object type.

   The right-hand pane is populated with all objects of that type.
5. In the right-hand pane, select one or more objects.

>

6 On the **Datasource Explorer** tool bar, click **Command**, and then select **Create Synonym**.

7 If you selected a single object, use the table below as a guide to completing the steps in this dialog:

Step	Settings and tasks	
Properties	**Owner**	Accept the default owner of the referenced object or select an owner for the synonym.
	Name	Provide a name for the synonym. 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. If you choose to create synonyms for multiple objects in a single operation, DBArtisan uses the object name as the synonym name. After creating synonyms in such a manner, users no longer need to specify the owner name.
	Referenced Object Owner, **Type**, and **Name**	Defaults show the owner, type, and name of the specific object you chose. Optionally, you can at this point select values for a different object to create a synonym for.
	Database Link	If the referenced object resides on a remote database, provide a database link.
DDL View	Preview the DDL and if necessary navigate backward through the steps to make corrections. Finally, use the **Execute** button to Create the new object.	

8 If you selected multiple objects, use the table below as a guide to completing the steps in this dialog:

Step	Settings and tasks
Action Options	For each object in the **Synonyms...** list, select the object and enable or disable **Public** scope for the synonym, and if you leave **Public** disabled, choose the **Schema** in which to create the synoynym.
Preview	Preview the DDL and if necessary navigate backward through the steps to make corrections. Finally, use the **Execute** button to Create the new object.

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.

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

Option	Description
File-Link Options	Lets you select file-link options: Read On Unlink Write Permission Recovery

For more information, see Creating objects.

DBCC

NOTE: This functionality is available for Microsoft SQL Server and Sybase ASE only.

DBArtisan's DBCC (Database Consistency Check) function box lets you:

- Specify single or multiple tables or indexes for validation.
- Perform database-wide validations.
- Perform object-level validations of databases.

DBCC options vary by the type of object and DBMS. Before performing a Database Consistency Check, see the relevant topic:

- DBCC for Microsoft SQL Server Databases
- DBCC for Microsoft SQL Server Tables
- DBCC for Microsoft SQL Server Indexes
- DBCC for Sybase ASE Databases
- DBCC for Sybase ASE Tables
- DBCC for Sybase ASE Indexes

To perform a Database Consistency Check:

1. Connect to the datasource that contains the object. For more information, see Datasource Management.
2. On the left pane of the **Datasource Explorer**, expand the datasource nodes until the target object type is displayed.
3. Select the target object type.

   DBArtisan displays objects of that type in the right pane of the Explorer window.
4. In the right pane of the Explorer window, select the specific object.
5. If you are checking consistency of a table or index, on the **Datasource Explorer** tool bar, click **Command**, and then select **DBCC**.

   OR
6. If you are checking consistency of a database, on the **Datasource Explorer** tool bar, click **Command**, then select **DBCC**, and then select one of the DBCC options listed in DBCC for Microsoft SQL Server Databases or DBCC for Sybase ASE Databases.

   DBArtisan opens a dialog that lets you specfy DBCC options.

7 Select options for the operation and when complete, use the **Schedule** or **Execute** button to complete the operation.

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

For more information, see:

DBCC for Microsoft SQL Server Databases

DBCC for Microsoft SQL Server Tables

DBCC for Microsoft SQL Server 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.
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 panel is reasonable, and that panel 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.

>

DBCC Operation	Description
Update Usage	Executes a DBCC UPDATEUSAGE command. Reports and corrects the rows, used, reserved, and dpanels 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:

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.
Check Fragmentation	Displays the target table's data and index fragmentation information, determining whether the table is heavily fragmented. 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 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/Image Allocation	NOTE: This option is not available Microsoft SQL Server version 7 or later. 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.
Pin Table	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. 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. This option is not available for Microsoft SQL Server 2005.
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.

Option	Description
Unpin Table	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. This option is not available for Microsoft SQL Server 2005.
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:

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.
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:

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 Sybase ASE Databases

DBCC for Sybase ASE Tables

DBCC for Sybase ASE Indexes

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:

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

DBArtisan 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 DBArtisan 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:

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	DBArtisan 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

Option	Description
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:

DBCC

Deallocate Unused Space

NOTE: 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 and 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.

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.	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. 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.

TIP: You can verify that the deallocated space is freed by going to the Space Tab in the appropriate DBArtisan editor.

Completing the Deallocate Unused Space Dialog Box

To complete this dialog box, do the following:

1 On the **Datasource Explorer**, select the Tables node.

DBArtisan displays the tables in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Deallocate Unused Space**.

OR

In the right pane of the application, right-click the table, and then select **Deallocate Unused Space**.

DBArtisan 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 Click **Execute**. For information on the other options, see Preview, Scheduling, and Impact Analysis.

For more information, seeCreating objects.

Detach/Attach

NOTE: 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 DBArtisan commands are:

- Detaching a Database
- Attaching a Database

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.

Completing the Detach Dialog Box

To complete this dialog box, do the following:

1 On the **Datasource Explorer**, select the **Databases** node.

DBArtisan displays the Databases in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Detach**.

OR

In the right pane of the application, right-click the database, and then select **Detach**.

DBArtisan 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 Click **Execute**. For information on the other options, see Preview, Scheduling, and Impact Analysis.

For more information, see Detach/Attach.

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.

Completing the Attach Dialog Box

To complete this dialog box, do the following:

1 On the **Datasource Explorer**, select the **Databases** node.

DBArtisan displays the Databases in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Attach**.

OR

In the right pane of the application, right-click the database, and then select **Attach**.

DBArtisan 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 2000 or 2005

- 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.

DBArtisan automatically adds the appropriate *.ldf file.

5 To drop database file(s), click **Drop** and then select the target file(s).

6 Click **Execute**. For information on the other options, see Preview, Scheduling, and Impact Analysis.

For more information, see Creating objects.

Disable Index

The **Disable Indexes** dialog box lets you mark an index, primary key, or unique key as disabled and unavailable for use.

NOTE: This functionality is available for Microsoft SQL Server 2005 and above.

>

To disable 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 **Disable** 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
Action options	Displays the name of the index, primary key, or unique key being disabled.
Dependencies	Review the referring and referred objects that will be automatically resolved when you execute this operation.
Preview	Preview the DDL generated for the operation and when ready, use the **Schedule** or **Execute** button to perform this action.

4. Click **Execute**. For information on the other options, see Preview and Scheduling.

Disable/Enable Job Queues

NOTE: 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.

Completing the Enable/Disable Dialog Box for Job Queues

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the **Job Queues** node.

   DBArtisan displays the job queues in the right pane of the Explorer window.
2. 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 job queue, and then select **Enable** or **Disable**.

   DBArtisan opens the Enable/Disable dialog box.
3. Click **Execute**. For information on the other options, see Preview, and Scheduling.

For more information, see Creating objects.

Enable/Disable Automatic Stored Outline Creation/Use Dialog Box

The table below describes the options and functionality on the Enable/Disable Automatic Stored Outline Creation/Use dialog box:

Option	Description
Disable automatic stored outline creation	If selected, Oracle does not automatically create a stored outline for every query submitted.

Option	Description
Enable automatic stored outline creation	If selected, Oracle automatically creates a stored outline for every query submitted on the system. Category - Lets you select an outline category. No Override - If selected, outlines will not override.

Completing the Enable/Disable Automatic Stored Outline Creation/Use Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the **Outlines** node.

   DBArtisan displays the Outlines in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Disable/Enable**.

   OR

3. In the right pane of the application, right-click the outline, and then select **Disable/Enable**.

   DBArtisan opens the Enable/Disable Automatic Stored Outline Creation dialog box.

4. Select **Disable automatic stored outline creation** to not automatically create a stored outline for every query submitted on the system.

5. Select **Enable automatic stored outline creation** to automatically create a stored outline for every query submitted on the system.

6. In **Category** select an outline category.

7. In **No Override** select to not allow outline overrides.

8. Click **Execute**. For information on the other options, see Preview, Scheduling, and Impact Analysis.

For more information, see Creating objects.

Disk Resize

NOTE: 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.

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.

Completing the Disk Resize Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the **Database Devices** node.

   DBArtisan displays the Database Devices in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Disk Resize**.

OR

In the right pane of the application, right-click the database device, and then select **Disk Resize**.

DBArtisan 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 Click **Execute**. For information on the other options, see Preview, Scheduling, and Impact Analysis.

For more information, see Creating objects.

Drop

NOTE: This functionality is available for DB2 LUW, DB2 z/OS, 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.

DBArtisan 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**.

DBArtisan 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.

DBArtisan 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**.

DBArtisan 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
Action Options	Lets you review the action you initiated.
Dependencies	Lists referring and referenced objects potentially impacted by the change. For details, see Dependencies.
Preview	Displays the DDL that will execute the object action. For details, see Preview/.

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

The Edit Data function opens the Data Editor. For more information, see Data Editor.

NOTE: This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

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. For more information, see Query Builder.

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 Using Data Editor with Query Builder.

TIP: For information on customizing Data Editor, see Data Editor options.

Error

The Error message displays a warning. After reading the warning, click Close to continue.

Estimate Size

NOTE: This functionality is available for Oracle and Sybase ASE only.

The 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. See the following topics:

- Estimate Table Sizes
- Estimate Index Sizes

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.
Estimated Table Size (KB)	Displays an estimated size when you click the **Estimate Size** button.
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. DBArtisan then calculates the information, and displays it in a browser-ready HTML report that you can print or save.
Update Statistics	For more information, see Update Statistics.
Save As	Lets you save as a result grid (.rsl) file.

TIP: To generate an HTML report, click the Report button.

For more information, see Creating objects.

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.
Estimated Index Size (KB)	Displays an estimated size when you click the **Estimate Size** button.
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. DBArtisan then calculates the information, and displays it in a browser-ready HTML report that you can print or save.
Update Statistics	For more information, see Update Statistics.
Save As	Lets you save as a result grid (.rsl) file.

TIP: To generate an HTML report, click the Report button.

For more information, see Creating objects.

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.

   DBArtisan 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**.

   DBArtisan 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
Action Options	Lets you review the action you initiated.

Step	Settings and tasks
Dependencies	Lists referring and referenced objects potentially impacted by the change. For details, see Dependencies.
Preview	Displays the DDL that will execute the object action. For details, see Preview/.

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

NOTE: This functionality is available for 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.

DBArtisan lets you enter arguments and set execution options in the Execution dialog box.

Completing the Execution Dialog Box

To complete this dialog box, do the following:

1 On the **Datasource Explorer**, select the target object node.

DBArtisan displays the target objects in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Execute**.

OR

3 In the right pane of the application, right-click the target object, and then select **Execute**.

DBArtisan opens the Execute dialog box.

4 To use the calendar, click the **Down** arrow.

5 To use the calculator, click the **Down** arrow.

NOTE: These built-in functions are context-sensitive to the type of parameter required

6 To execute the function with show plan, click **Query Plan**.

7 To modify your results options, click **Results**.

8 Click **Execute**. For information on the other options, see Preview, Scheduling, and Impact Analysis.

For more information, see:

Executing Functions

Executing Procedures

Executing Functions

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

Option	Description
Query Plan	Click the Query Plan button to activate the Oracle EXPLAIN PLAN functionality. For information on where DBArtisan stores the EXPLAIN PLAN output, see Query options DBArtisan's Query Plan window now displays data for the estimated costs, number of rows, and bytes returned by each plan step.
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.

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.
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.

Extract

NOTE: This functionality is available for all platforms.

>

DBArtisan 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

With respect to functions and stored procedures, it can be necessary to have statements executed before and after creation of the procedure or function. This can be useful for example, if you need to create or drop temporary tables used by the function or procedure. DBArtisan supports the use of two tag pairs, ETStart and ETEnd, that let you embed statements in the first comment block of a stored procedure or function. The following shows the expected syntax:

```
create procedure dbo.procname(@a numeric) as
/*
<ETStart>SQL Statement</ETStart>
<ETEnd>SQL Statement</ETEnd>
*/
begin
...
```

For more information, see:

Extracting Data

Extracting Data

To extract data, do the following:

9 In the **Datasource Explorer,** click the target object node.

DBArtisan displays the selected object in the right pane of the Explorer window.

10 In the right pane of the **Explorer** window, click the target object.

11 On the **Datasource Explorer** tool bar, click **Extract**.

OR

On the **Datasource Explorer** tool bar, click **Command,** and then click **Extract**.

OR

In the right pane of the **Explorer** window, right-click the target object, and then click **Extract**.

DBArtisan opens the DDL Editor.

Extract Data as XML

This function allows you to take SQL data, extract it, and make it available as XML data. The **XML Editor Filter** allows you to pick and choose among columns and designate rows and then creates the For XML statement that enables the operation. The resulting XML document is created and presented in an active XML Editor. At this point the document can be saved in XML format.

NOTE: This functionality is available for Oracle 9i, SQL Server 8.0, and Sybase 12.5.1 or later.

CAUTION: To use the Extract Data as XML feature in Sybase, you must first purchase and install an XML Services license from Sybase.

>

To Open the XML Editor Filter

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**.

The table below describes the fields you may encounter as you complete the **XML Editor Filter** dialog box.

Required Information	Description
Columns	Check the columns with data you want to extract.
Where Clause	Write a Where clause to identify how rows should be connected in the result. For example, WHERE doctors.city = hospitals.city.
Select	Displays the SELECT statement generated as you select other options.
SYBASE ONLY: XML Options	For detailed information, refer to theinformation on XML mapping at http://infocenter.sybase.com/.
ORACLE ONLY: Row Set Tag Row Tag Max Rows.	Identify the XML element names you want to use to replace the Row Set and the Row tags. Max rows is the maximum number of rows to fetch.
SQL SERVER ONLY: Optional Settings AUTO RAW XMLDATA ELEMENTS BINARY BASE64	AUTO mode returns query results as nested XML elements. For any table in the From clause with a column in the Select clause, it is represented as an XML element. When you select RAW mode, each row in the query result set is transformed into an XML element with the generic row identifier. The XML attribute name will be the same as the column name for non-null columns. XML DATA specifies that an XML data schema will be returned. ELEMENTS specifies that columns will be returned as subelements--otherwise they are mapped to XML attributes. This is an option only if AUTO is selected. When BINARY BASE64 is selected, any binary data is encoded in base-64 format. You must specify this option in RAW mode. It is the default for AUTO.

NOTE: As you select options, inspect the **Select** box to verify the query you are creating.

Flush Cache

NOTE: This functionality is available for IBM DB2 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 for Linux, Unix, and Windows to recompile the SQL the next time it is called.

Completing the Flush Cache Dialog Box

To complete this dialog box, do the following:

1 On the **Datasource Explorer**, select the target object node.

DBArtisan displays the target objects in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Flush Cache**.

DBArtisan opens the Flush Cache dialog box.

3 Click **Execute**. For information on the other options, see Preview and Scheduling.

Free Packages

NOTE: This functionality is available for IBM DB2 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, click **Execute**. For information on the **Preview** option, see Preview.

Free Plans

NOTE: This functionality is available for IBM DB2 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.

Completing the Free Plans Dialog Box

To complete the Free Plans dialog box, do one of the following:

1. On the **Datasource Explorer**, select the **Plans** node.

   DBArtisan displays the Plans in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Free**.

   OR

   In the right pane of the application, right-click the target object, and then select **Free**.

   DBArtisan opens the Free Plans dialog box.

3. Do one of the following:

4. Click **Execute**. For information on the **Preview** option, see Preview.

Copying Images

The Image Copy Wizard lets you create up to four image copies, exact reproductions, of any of the following objects:

- Tablespace
- Tablespace partition
- Data set of a linear table space
- Index space defined with COPY YES
- Index space partition

The Image Copy Wizard lets you create two types of image copies:

- A full image copy of all pages in a tablespace, partition, data set, or index space.

>

- An incremental image copy of only pages that have been modified since the last use of the Image Copy Wizard.

For more information, Image Copy Wizard - Panel 1.

Image Copy Wizard - Panel 1

The table below describes the options and functionality on this panel of the Image Copy Wizard:

Option	Description
Do you want the utility to be restartable?	Lets you specify to execute Wizard in restartable mode. You can restart it from the Datasource Explorer.
Tape Stacking	Stacks similar output data sets as successive files on one logical tape volume, where a logical tape volume can consist of a multi-volume aggregate, if any of the resulting data sets on the tapes span volumes. Within one utility execution, output data sets are stacked on a logical tape volume of the same usage type. For example, local primary image copies are stacked separately from local backup image copies. Also lets you specify the Copy and Recovery dataset information once as variable names appended with generated values.
Set Properties	Opens the Dataset Properties Dialog Box.
Do you want to execute DFSMS concurrent copy to make the full image copy?	Lets you specify the Concurrent Copy function of Data Facility Storage Management Subsystem (DFSMS). The wizard records the resulting DFSMS concurrent copies in the catalog table SYSIBM.SYSCOPY with ICTYPE=F and STYPE=C. If you select No, to not to execute DFSMS concurrent copy to make a full image copy, you can select Check Page.
Do you want the objects to be processed in parallel?	DBArtisan initiates multiple concurrent requests for a single user query and performs I/O processing concurrently (in parallel) on multiple data partitions. If you select Yes, to process the objects in parallel, you can type the number of objects. If you select to process the tape units in parallel, you can type the number of tape units.
Do you want other programs to access or update the selected objects when the COPY utility is running?	If you select Yes, you can select the Level of Access.

Dataset Properties Dialog Box

The table below describes the options and functionality on the Dataset Properties Dialog Box:

Option	Description
Local Site Primary Copy	Lets you specify the Local Site Primary Copy.
Local Site Backup Copy	Lets you specify the Local Site Backup Copy.
Recovery Site Primary Copy	Lets you specify the Recovery Site Primary Copy.
Recovery Site Backup Copy	Lets you specify the Recovery Site Backup Copy.

Image Copy Wizard - Panel 2

The table below describes the options and functionary on this panel of the Image Copy Wizard:

Option	Description
Select tablespaces	Displays the tablespaces to be copied.

>

Option	Description
Add	Opens the Tablespace Selector Dialog Box.
Set Properties	Opens the Object Properties Dialog Box.
Delete	Deletes the selected tablespace.

Tablespace Selector Dialog Box

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

Option	Description
Database Like	Lets you type the pattern for querying the tablespaces. You can type multiple words, separated by spaces.
Tablespace Like	Lets you type the pattern for querying the tablespaces. You can type multiple words, separated by spaces.
Match Case	Lets you specify match case.
Tablespace Selector	Lets you select one or more rows in this list.
Query	Repopulates Tablespace Selector based on the tablespace query.
Select All	Selects all rows in Tablespace Selector.
Unselect All	Unselects all rows in Tablespace Selector.
Add	Click to open the Select Partition(s) for Tablespace Dialog Box.

Select Partition(s) for Tablespace Dialog Box

The table below describes the options and functionality on the Select Partition(s) for Tablespace Selector dialog box:

Option	Description
Add all partition	Select to add all partitions for the tablespace.
Select Partitions	Lets you select partition(s) or click Select All or Unselect All.

Image Copy Wizard - Panel 3

The table below describes the options and functionary on this panel of the Image Copy Wizard:

Option	Description
Select Indexspaces	Displays the indexspaces to be copied.
Add	Opens the Indexspace Selector Dialog Box.
Set Properties	Opens the Object Properties Dialog Box.
Delete	Deletes the selected indexspace.

>

Indexspace Selector Dialog Box

The table below describes the options and functionality on the panel of the Indexspace Selector dialog box:

Option	Description
Database Like	Lets you type the pattern for querying the indexspaces. You can type multiple words, separated by spaces.
IndexSpace Like	Lets you type the pattern for querying the indexspaces. You can type multiple words, separated by spaces.
Match Case	Lets you specify match case.
Indexspace Selector	Lets you select one or more rows in this list.
Query	Repopulates Indexspace Selector based on the tablespace query.
Select All	Selects all rows in Indexspace Selector.
Unselect All	Unselects all rows in Indexspace Selector.

Image Copy Wizard - Panel 4

The table below describes the options and functionary on this panel of the Image Copy Wizard:

Option	Description
Select index.	Displays the indexes to be copied.
Add	Opens the Index Selector Dialog Box.
Set Properties	Opens the Object Properties Dialog Box.
Delete	Deletes the selected indexspace.

Index Selector Dialog Box

The table below describes the options and functionality on the panel of the Index Selector dialog box:

Option	Description
Database Like	Lets you type the pattern for querying the Indexes. You can type multiple words, separated by spaces.
Index Creator	Lets you type the pattern for querying the Indexes. You can type multiple words, separated by spaces.
Index Like	Lets you type the pattern for querying the Indexes. You can type multiple words, separated by spaces.
Match Case	Lets you specify match case.
Index Selector	Lets you select one or more rows in this list.
Query	Repopulates Index Selector based on the tablespace query.
Select All	Selects all rows in Index Selector.
Unselect All	Unselects all rows in Index Selector.

Object Properties Dialog Box

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

Option	Functionality
Full Copy?	Select to copy the entire tablespace.
Change Limit	Lets you specify the Change Limit option, the changed pages in the tablespace. If the high Change Limit value is met, then a full image copy is taken.
Range	Lets you specify a change limit range between 0.0 and 100.0, the changed pages in the tablespace. If the high Change Limit value is met, then a full image copy is taken.
Report Only	Lets you specify the Report Only option.
Local Site Primary Copy	Lets you specify the Local Site Primary Copy.
Local Site Backup Copy	Lets you specify the Local Site Backup Copy.
Recovery Site Primary Copy	Lets you specify the Recovery Site Primary Copy.
Recovery Site Backup Copy	Lets you specify the Recovery Site Backup Copy.

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.

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.

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.

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.

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	Lets you specify the objects to search within the schemas defined. Add - Click to open the Select a Resolver Option dialog box to add a new resolver option in the list. Edit - Click to open the Resolver Edit dialog box to modify a resolver option. Remove - Select one or more resolver option and click to delete.

Lock

NOTE: This functionality is available for IBM DB2 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. DBArtisan releases locks at the end of a transaction.

The table below describes the options and functionality on the Lock Table dialog box.

Option	Description
Lock Mode	Share - Lets other users view but not modify the table data. Exclusive - Prevents other users from viewing or modifying the table data.

Completing the Lock Table Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the **Tables** node.

   DBArtisan displays the Tables in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Lock**.

   OR

   In the right pane of the application, right-click the table, and then select **Lock**.

   DBArtisan 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 Click **Execute**. For information on the other options, see Preview, Scheduling, and Impact Analysis.

Migrate

Right-clicking an object and selecting **Migrate** from the context menu opens the Migration wizard. This lets you move object schema between DBMS or servers by copying the DDL between the source and target databases. For detailed instructions, see Using the schema migration wizard.

Merge Copy

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS.

The Merge Copy function makes a merged copy of the data sets within a table space. You can merge incremental copies of a tablespace to make one incremental copy, or you can merge incremental copies with a full image copy to make a new full image copy.

To Create a Merge Copy

1 Connect to a DB2 datasource and expand the Explorer tree.

2 Click the **Utility Execution** node and double-click **MERGECOPY**.

The tables below displays all the fields you'll find in the merge copy wizard:

Required Information	Description
Do you want the utility to be restartable?	Yes or No
Tape Stacking Properties	Yes or no. When you select Yes, you then set the Dataset properties.
Do you want to merge all incremental copies with the full image copy to make a new full image copy?	Self-evident
WorkDDN	Here you are setting the dataset information for the working data definition name.
Select tablespaces	Self-evident.

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.

>

Modify Recovery

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS only.

The Modify Recovery Wizard lets you delete records from SYSIBM.SYSCOPY, related log records from SYSIBM.SYSLGRNX, and entries from the DBD. You can delete records based on a specific age or based on the date the record was written. You can also delete records for an entire table space, partition or data set.

You should run the wizard regularly to clear outdated information from SYSIBM.SYSCOPY and SYSIBM.SYSLGRNX. These tables, and particularly SYSIBM.SYSLGRNX, can become very large and take up considerable amounts of space. By deleting outdated information from these tables, performance for processes that access data from these tables can be improved.

For more information, see Modify Recovery Wizard - Panel 1

Modify Recovery Wizard - Panel 1

The table below describes the options and functionality on this panel of the Modify Recovery Wizard:

Option	Description
Select tablespaces	Displays the available tablespaces.
Add	Opens the Tablespace Selector Dialog Box.
Delete	Deletes the selected items.

Modify Recovery Wizard - Panel 2

The table below describes the options and functionality on this panel of the Modify Recovery Wizard:

Option	Description
Delete history record(s) by	Lets you select Age or Date.
Number of Days	Available if you specified to delete history records by age. Lets you type a number between 1 and 32767.
Date	Available if you specified to delete history records by date. Lets you type a date in eight digit format yyyymmdd (e.g. 20000218)
Delete Options	Lets you select delete options.

Modify Statistics

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS only.

The Modify Statistics Wizard lets you delete unwanted statistics history records from the statistics catalog tables. You can delete statistics history records that were written before a specific date, or records of a specific age. You can delete the records for an entire tablespace, index space, or index.

You should run the wizard regularly to clear outdated information from the statistics history catalog tables. Deleting outdated information from the statistics history catalog tables, improves the performance of processes that access data from those tables.

For more information, see Modify Statistics Wizard - Panel 1.

Modify Statistics Wizard - Panel 1

The table below describes the options and functionality on this panel of the Modify Statistics Wizard:

Option	Description
Delete history record(s) by	Lets you select Age or Date.
Number of Days	Available if you specified to delete history records by age. Lets you type a number between 1 and 32767.
Date	Available if you specified to delete history records by date. Lets you type a date in eight digit format yyyymmdd (e.g. 20000218)
Delete Options	Lets you select delete options.

Modify Statistics Wizard - Panel 2

The table below describes the options and functionality on this panel of the Modify Statistics Wizard:

Option	Description
Select tablespaces	Displays the available tablespaces.
Add	Opens theTablespace Selector Dialog Box.
Delete	Deletes the selected items.

Modify Statistics Wizard - Panel 3

The table below describes the options and functionality on this panel of the Modify Statistics Wizard:

Option	Description
Select indexes	Displays the available indexes.
Add	Opens the Index Selector Dialog Box.
Delete	Deletes the selected items.

Modify Statistics Wizard - Panel 4

The table below describes the options and functionality on this panel of the Modify Statistics Wizard:

Option	Description
Select indexspaces	Displays the available indexspaces.
Add	Opens the Indexspaces Selector dialog box.
Delete	Deletes the selected items.

Move Log

NOTE: This functionality is available for Sybase ASE only.

The Move Transaction Log dialog box lets you move a transaction log from one device to another.

>

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.

Completing the Move Log Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the **Databases** node.

   DBArtisan displays the Databases in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Move Log**.

   OR

   In the right pane of the application, right-click the **Database**, and then select **Move Log**.

   DBArtisan opens the Move Log dialog box.

3. In **New Device**, select the device to which you want to move the transaction log.

4. Click **Execute**. For information on the other options, see Preview, Scheduling, and Impact Analysis.

Open

NOTE: This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

DBArtisan 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

Open Control Files

NOTE: This functionality is available for Oracle 8 or later only.

The control file is a small binary file containing a record of the database schema. It is one of the most essential files in the database because it is necessary for the database to start and operate successfully.

The control file contains information about the associated database that is required for the database to be accessed by an instance, both at startup and during normal operation.

NOTE: Only the Oracle database server can modify a control file's information; no user can edit a database's control file directly.

The Control File Information dialog box displays control file information, including:

- Type
- Record Size
- Records Total
- Records Used
- First Index
- Last Index
- Last Recid

For more information, see Control Files.

Package Editor

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS only.

The Package Editor lets you modify a package.

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

NOTE: 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.

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.

>

For more information, see Creating objects.

Partition Editor

The Partition Editor lets you specify partition management and partition options.

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	A DB2-managed data set that resides on a volume of the identified storage group. Stogroup Name - Lets you specify the stogroup. Primary Space Allocation - Lets you specify the minimum primary space allocation for a DB2-managed data set of the partition. Secondary Space Allocation - Lets you specify the minimum secondary space allocation for a DB2-managed data set of the partition.
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.
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.

For more information, see Creating objects.

Partition Upper Bound

NOTE: 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.

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.

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**.

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.

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.

Completing the Placement Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the target object node.

   DBArtisan displays the target objects in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Place**.

   OR

   In the right pane of the application, right-click the target object, and then select **Place**.

   DBArtisan opens the Placement dialog box.

3. In **New Segment Location**, select the segments to place the object.
4. Click **Execute**. For information on the other options, see Preview, Scheduling, and Impact Analysis.

For more information, see Creating objects.

Quiesce Database

The Quiesce Database dialog box lets you restrict user access to a database.

NOTE: This functionality is available for IBM DB2 for Windows/Unix 8.1 only. DBArtisan does not support Quiesce Database for an IBM DB2 8 server from a IBM DB2 for Windows/Unix 7 client or for an IBM DB2 for Windows/Unix 7 server from an IBM DB2 for Windows/Unix 7 client.

The table below describes the options and functionality on the Quiesce Database dialog box.

Option	Description
Connection	Lets you select a connection option.

Completing the Quiesce Database Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the **Datasources** node.

   DBArtisan displays the Databases in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Quiesce**.

   OR

   In the right pane of the application, right-click the target object, and then select **Quiesce**.

   DBArtisan opens the Quiesce dialog box.

3. Select a **Connection** option.
4. Set the **For user** or **For group** radio box and specify a user or group name as appropriate.
5. Click **Execute**.

Quiesce Instance

NOTE: This functionality is available for IBM DB2 for Windows/Unix 8.1 only.

The Quiesce Instance dialog box lets you restrict user access to an instance.

The table below describes the options and functionality on the Quiesce Instance dialog box.

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.

Completing the Quiesce Instance Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the **Instance** node.
2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Quiesce**.

   OR

   In the right pane of the application, right-click the target object, and then select Quiesce.

   DBArtisan opens the Quiesce dialog box.
3. To specify a user, select **For user** option, and then type user name.
4. To specify a group, select **For group** option, and then type group name.
5. Select a **Connections** option.
6. Click **Execute**.

For more information, see Creating objects.

Quiesce Tablespaces

NOTE: This functionality is available for IBM DB2 for Linux, Unix, and Windows and IBM DB2 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.

The table below describes the options and functionality on the Quiesce Tablespaces dialog box.

Option	Description
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.

Completing the Quiesce Tablespaces Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the **Tables** node.

   DBArtisan displays the Tables in the right pane of the Explorer window.
2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Quiesce Tablespaces**.

   OR

   In the right pane of the application, right-click the target object, and then select **Quiesce Tablespaces**.

   DBArtisan 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 Click **Execute**. For information on the other options, see Preview, Scheduling, and Impact Analysis.

For more information, see Creating objects.

Reassign by Category

DBArtisan lets you reassign the category of stored outlines in Oracle.

NOTE: This functionality is available for Oracle 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.

To reassign a stored outline to a different category:

1 Connect to the datasource that contains the stored outline. For more information, see Datasource Management.

2 On the **Datasource Explorer**, expand the datasource nodes until the **Stored Outlines** node is visible and then select the **Stored Outlines** node.

DBArtisan displays the stored outlines in the right pane of the Explorer window.

3 In the right pane of the datasource explorer, select the stored outline for which you want to reassign categories.

4 On the **Datasource Explorer** tool bar, click **Command**, and then select **Reassign**.

DBArtisan opens the **Reassign by Category** dialog box.

5 From the **Target Category** dropdown, choose a new, target category.

6 Review the **Preview** panel and when ready, use the **Schedule** or **Execute** button to complete the task.

Rebind Packages

NOTE: This functionality is available for IBM DB2 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.

TIP: If the physical storage of a package is changed or dropped, rebinding updates the path of the SQL statements.

Completing the Rebind Packages Dialog Box

To complete this dialog box, do the following:

1 On the **Datasource Explorer**, select the **Packages** node.

DBArtisan displays the Packages in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Rebind**.

OR

In the right pane of the application, right-click the target object, and then select **Rebind**.

DBArtisan opens the Rebind dialog box.

3 Click **Execute**. For information on the other options, see Preview and Impact Analysis.

For more information, see Creating objects.

Rebind Plans

NOTE: This functionality is available for IBM DB2 for Linux, Unix, and Windows and OS/390 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.

Completing the Rebind Plans Dialog Box

To complete the Rebind Plans dialog box, do one of the following:

1 On the **Datasource Explorer**, select the **Plans** node.

DBArtisan displays the Plans in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Rebind**.

OR

In the right pane of the application, right-click the target object, and then select **Rebind**.

DBArtisan opens the Rebind Plans dialog box.

3 Click **Execute**. For information on the **Preview** option, see Preview.

For more information, see Creating objects.

Rebuild Index

See the following topics for DBMS-specific instructions on rebuilding indexes:

- Rebuild Index (DB2 z/OS)
- Rebuild Index (Oracle)
- Rebuild Index (SQL Server)

Rebuild Index (DB2 z/OS)

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS.

The Rebuild Index Wizard reconstructs indexes from the table that they reference.

>

The table below describes the options and functionality on the Rebuild Index Wizard.

Panel	Option	Description
1	Select Indexes	Lets you select an index to rebuild.
	Add	For more information, see Index Selector Dialog Box.
	Delete	Deletes the selected index(es).
2	Select tablespaces	Lets you select the target tablespace.
	Add	For more information, see Table Selector Dialog Box.
	Delete	Deletes the selected items.
3	Do you want the utility to be restartable?	Select to make the utility restartable.
	Do you want to reuse the db2 managed dataset?	Default is No.
	WorkDDN	Lets you select a work ddn template from the list or click Edit to specify the dataset information.
	Do you want to specify to sort device type?	Default is No.
	Do you want to specify the sort number?	Default is No.
	Do you want to specify the sort keys?	Default is No.
4	Do you want to specify the stats online?	Lets you specify the stats online.
	Do you want to output message to SYSPrint?	The default is no.
	Do you want to specify the correlation stats option?	Lets you specify correlation stats on the next panel of the wizard.
	Do you want to force aggregation or rollup processing to be done even though some parts do not contain data?	The default is Default.
5	KeyCard	Collects all of the distinct values in all of the 1 to n key column combinations for the specified indexes. n is the number of columns in the index.
	Frequent Item	Lets you add items to the grid.
	Add	For more information, see Add Frequent Item Dialog Box.
6	Update Statistics	Lets you specify update statistics.
	History Statistics	Lets you specify history statistics.

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.

DBArtisan 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**.

DBArtisan 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
Action options	To rebuild a single partition, specify a **Partition number** and optionally, provide **Sort in tempdb** and **MaxDOP** property values used to create the REBUILD WITH clause. To rebuild the entire index, primary key, or unique key, DO NOT provide a **Partition Number**, and optionally, provide **Pad Index**, **Sort in tempdb**, **Ignore Duplicate Key**, **Statistics no recompute**, **Online**, **Allow Row Locks**, **Allow Page Locks**, **MaxDOP**, and **Fill Factor** property values used to create the REBUILD WITH clause. For more information on these properties, see Index Wizard (SQL Server).
Dependencies	Lets you review the objects potentially impacted by this action.
Preview	Preview the DDL generated for the operation and when ready, use the **Schedule** or **Execute** button to perform this action.

4 Click **Execute**. For information on the other options, see Preview and Scheduling.

Rebuild Outlines

NOTE: 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.

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.

For more information, see Creating objects.

Recompile

NOTE: 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

NOTE: This functionality is available for IBM DB2 for Linux, Unix, and Windows only.

The Redistribute dialog box lets you redistribute the data in a nodegroup.

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

For more information, see Creating objects.

Refit

NOTE: This functionality is available for Sybase ASE only.

The Refit dialog box lets you refit a database device. If you are restoring a damaged master database, or have added devices since the last database or transaction log dump, then refit a device after reinitializing. Refitting a device ensures that the information in sysusages and sysdatabases is consistent with the information in sysdevices.

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

Completing the Refit Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the **Database Devices** node.

   DBArtisan displays the Database Devices in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Refit**.

   OR

   In the right pane of the application, right-click the target object, and then select **Refit**.

   DBArtisan opens the Refit dialog box.

3. Click **Execute**. For information on the other options, see Preview, Scheduling, and Impact Analysis.

Refresh Table

The Refresh Summary Table Dialog box lets you reload materialized query tables that have been defined with refresh options.

NOTE: This functionality is available for IBM DB2 for Linux, Unix, and Windows 8.1 and later.

To reload a materialized query table:

1. Connect to the datasource that contains the materialized query table. For more information, see Datasource Management.
2. On the **Datasource Explorer**, expand the datasource nodes until the **Materialized Query Tables** node is visible and then select the **Materialized Query Tables** node.

   DBArtisan displays the materialized query tables in the right pane of the Explorer window.

3. In the right pane of the datasource explorer, select the materialized query table you want to reload.
4. On the **Datasource Explorer** tool bar, click **Command**, and then select **Refresh**.

   DBArtisan opens the **Refresh Materialized Query Table** dialog box.

5 Use the following table as a guide to understanding and modifying setings on this dialog:

Pane	Description
Action Options	Select an **Online Option** of ALLOW NO ACCESS, ALLOW READ ACCESS or ALLOW WRITE ACCESS.. Enable or disable **Query Optimization**. Select an INCREMENTAL or NOT INCREMENTAL **Refresh Option.**
Preview	Preview the DDL generated by the options you chose.

6 Review the **Preview** panel and when ready, use the **Schedule** or **Execute** button to complete the task.

Important Notes

For procedures on restoring a damaged master database, consult the Commands Reference Manual.

Reinitialize

NOTE: This functionality is available for Sybase ASE only.

The reinitialize function allows you to keep a database device on the Sybase server open while resetting the device to its state immediately after it was opened.

Completing the Reinitialize Dialog Box

To complete this dialog box, do the following:

1 On the **Datasource Explorer**, expand the **Storage** node, and click **Database Devices**.

2 In the right pane, click the device you want to reinitialize.

3 Select **Reinit...** from the Command menu or right-click and select it.

4 Click **Execute**. For information on the other options, see Preview, and Scheduling.

Rename

NOTE: This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

DBArtisan lets you rename an object. In general, all referenced or referring objects are updated to reflect the new name. The following table notes the exceptions and provides prerequisite tasks to be performed before renaming an object.

DBMS	Notes and restrictions on renaming
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 will not rename a 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.

DBMS	Notes and restrictions on renaming
IBM DB2 for OS/390 and z/OS	DBArtisan 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.

To rename an object:

1. Connect to the datasource that contains the object. For more information, see Datasource Management.
2. On the left pane of the **Datasource Explorer**, expand the datasource nodes until the target object type is displayed.
3. Select the target object type.

   DBArtisan displays objects of that type in the right pane of the Explorer window.
4. In the right pane of the Explorer window, select the object to be renamed.
5. On the **Datasource Explorer** tool bar, click **Command**, and then select **Rename**.

   DBArtisan opens the **Rename** dialog.
6. Use the following table as a guide to understanding and modifying settings in the dialog:

Step	Settings and tasks	
Rename	**Name**	Provide the new name for the object.
Dependencies	Review the referring and referred objects for which naming will be automatically resolved when you execute the renaming operation.	
Preview	Preview the DDL generated for the operation and when ready, use the **Schedule** or **Execute** button to perform this action.	

Reorganize

The Reorganize dialog box lets you alter a table's physical storage to ensure that it is stored efficiently and to eliminate fragmentation.

Completing the Reorganize Dialog Box

To complete this dialog box, do the following:

1 On the **Datasource Explorer**, select the target object node.

DBArtisan displays the target objects in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Reorganize**.

OR

In the right pane of the application, right-click the target object, and then select **Reorganize**.

DBArtisan opens the Reorganize dialog box.

3 Use the information in the following topics to provide reorganization details, depending on the DBMS that DBArtisan is working with:

- Reorganizing IBM DB2 for Linux, Unix, and Windows Objects
- Reorganizing Oracle Objects
- Reorganizing SQL Server indexes, primary keys, and unique keys
- Reorganizing Sybase ASE Objects

4 Click **Execute**. For information on the other options, see Preview and Scheduling.

Reorganizing IBM DB2 for Linux, Unix, and Windows Objects

Reorganize Dialog Box (One Table)

The table below describes the options and functionality on the Reorganize dialog box.

Step	Description
Action Options	• Choose a **Reorg Type** of table or all indexes of the table. • Optionally, choose an **Index Schema** and specify an **Index name**. • Enable or disable **Inplace Reorg**anization and if you enable, choose an **Inplace reorg**anization **mode** of Start, Stop, Pause, or Resume. • Select an **Access mode** to control read and write access. • Select an **Index reorg**anization **mode** to clean up empty pages, delete after cleaning up empty pages, or convert to a type 2 index. • If you did not enable **Inplace Reorg**anization, select a **Tablespace**. • Select a L**ong Tablespace**. • Enable or disable the following options: **Index scan**, **Reorg**anize **long field and LOB data**, **Reset Dictionary**, **No truncate table**, and **Reorganize all partitions**.
Partitions	If you did not enable **Reorganize all partitions**, select the partitions to reorganize.
Preview	Preview the DDL code generated from your choices.

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.

Reorganizing Oracle Objects

The Reorganize dialog box lets you reduce query processing time against tables.

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

Group	Description
	If you want to move the table(s) to a new tablespace, click the **New Tablespace** list and select the target tablespace.
Data Block Storage	• 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.
Extents	• 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.
Freelists	• In the **Freelists** box, type the new value. • In the **Freelist Group** box, type the new value. • To modify the bufferpool, click the **Bufferpool**, click the list and select **DEFAULT** to use the default bufferpool, **KEEP** to retain the object in memory to avoid I/O conflicts, or **RECYCLE** to save cache space by ridding data blocks from memory as soon as they are no longer in use. This option is only available for Oracle8.

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 aslso lets you specify a LOB_COMPACTION 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 settiing 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
Action options	To reorganize a single partition, specify a **Partition number** and optionally, specify **Lob Compaction** used to create the REORGANIZE WITH clause. To reorganize the entire index, primary key, or unique key, DO NOT provide a **Partition Number**, and optionally,specify **Lob Compaction** used to create the REORGANIZE WITH clause.
Dependencies	Review the referring and referred objects that will be automatically resolved when you execute this operation.
Preview	Preview the DDL generated for the operation and when ready, use the **Schedule** or **Execute** button to perform this action.

4. Click **Execute**. For information on the other options, see Preview and Scheduling.

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

- Reorganizing Sybase ASE Tables
- Reorganize Sybase ASE Indexes

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. For more information, see Incremental Reorganizations.

> **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

> **NOTE:** You can reorganize tables in Sybase ASE versions 12 and 12.5.

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 resume 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. For information on resule options, see Incremental Reorganizations.
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 resume 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. For information on resule options, see Incremental Reorganizations.
Rebuild	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 sp_chgattribute), and drops and re-creates all the target table's (or tables') 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.
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	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.

Incremental Reorganizations

If target tables are too long to reorganize in one session, DBArtisan lets you reorganize them in increments over multiple sessions by specifying a maximum duration for each session. After DBArtisan 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, DBArtisan 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, DBArtisan 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.

Reorganize Tablespace

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS only.

The Reorganize Tablespace Wizard reorganizes a tablespace to improve access performance and reclaim fragmented space. The wizard can also reorganize a single partition or range of partitions of a partitioned table space. You can specify the degree of access to the data during reorganization, and optionally collect inline statistics.

>

Completing the Reorganize Tablespace Wizard

To complete the Reorganize Tablespace Wizard, do the following:

1 On the **Datasource Explorer**, select the tablespaces node.

DBArtisan displays the tablespaces in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Reorganize**.

DBArtisan opens the Reorganize Tablespace Wizard.

3 Complete the wizard panels, and then click **Finish**.

DBArtisan opens the Preview dialog box.

Reorganize Tablespace Wizard - Panel 1

The table below describes the options and functionality on this panel of the Reorganize Tablespace Wizard:

Option	Description
Do you want the utility to be restartable?	Select to make the utility restartable.
Database	Displays the list of databases in the system. You can select the tablespace database or click the filter button to limit the list of databases to a string pattern.
Tablespace	Displays the list of tablespaces.
Partition	Lets you specify a partition or a range of partitions.
Reuse	Select to reuse this tablespace.

Reorganize Tablespace Wizard - Panel 2

The table below describes the options and functionality on this panel of the Reorganize Tablespace Wizard:

Option	Description
Log	List of the sequence of events after you execute the wizard.
Sort Data	Sorts columns alphanumerically in single result sets.
No Sysrec	Select Yes for no sysrec.
Sortkeys	Select Yes for sortkeys. Specifies that index keys are to be sorted and built in parallel during the SORTBLD phase to improve performance.

Reorganize Tablespace Wizard - Panel 3

The table below describes the options and functionality on this panel of the Reorganize Tablespace Wizard:

Option	Description
Share level	Lets you select share level that allows users to view but not modify the table data: None, Reference, or Change.

Option	Description
Deadline	Lets you specify a deadline for the switch phase of reorganization to start. If DBArtisan estimates that the switch phase will not start by the deadline, DBArtisan terminates reorganization. None - Specifies that there is no deadline for the read-only iteration of log processing. Timestamp - Specifies the deadline for the switch phase to start processing. Click the button to open the Deadline Expression Builder dialog box.
Drain Specification	Lets you specify that DB2 drains the write claim class after the delay. The number of log records, and thus the estimated time, for a future iteration of log processing will be 0.
Fast Switch	Keeps the data set name unchanged and updates the catalog to reference the newly reorganized data set.

Deadline Expression Builder

The table below describes the options and functionality on the Deadline Expression Builder dialog box:

Option	Description
Current Date	Lets you select today as the basis of the deadline, click + or -, type the number of years, months, days, hours, minutes, seconds and microseconds.
Current Timestamp	Lets you select the current timestamp as the basis of the deadline, click + or -, type the number of years, months, days, hours, minutes, seconds and microseconds.

Reorganize Tablespace Wizard - Panel 4

The table below describes the options and functionality on this panel of the Reorganize Tablespace Wizard:

Option	Description
Local Site Primary	Lets you specify the Local Site Primary.
Local Site Backup	Lets you specify the Local Site Backup.
Recovery Site Primary	Lets you specify the Recovery Site Primary.
Recovery Site Backup	Lets you specify the Recovery Site Backup.

Reorganize Tablespace Wizard - Panel 5

The table below describes the options and functionality on this panel of the Reorganize Tablespace Wizard:

Option	Description
Off Pos Limit	The specified value is compared to the result of the calculation (NEAROFFPOSF + FAROFFPOSF) × 100 / CARDF for the specified partitions in SYSIBM.SYSINDEXPART for the explicit clustering indexes for every table in the specified table space. Alternatively, the values in SYSINDEXPART are checked for a single non-partitioned table space, or for each partition if you specified an entire partitioned table space as the target object. If at least one calculated value exceeds the OFFPOSLIMIT value, REORG is performed or recommended. **NOTE:** This option is valid for non-LOB table spaces only.

Option	Description
Ind Ref Limit	The specified value is compared to the result of the calculation (NEARINDREF + FARINDREF) × 100 / CARDF for the specified partitions in SYSIBM.SYSTABLEPART for the specified table space. Alternatively, the values in SYSTABLEPART are checked for a single non-partitioned table space, or for each partition if you specified an entire partitioned table space as the target object. If at least one calculated value exceeds the INDREFLIMIT value, REORG is performed or recommended. **NOTE:** This option is valid for non-LOB table spaces only.
Report Only	The reorganization is only recommended, not performed.
Unload	Specifies whether the utility job is to continue processing or end after the data is unloaded. Continue - Specifies that, after the data has been unloaded, the utility continues processing. Pause - Specifies that after the data has been unloaded, processing ends. Only - Specifies that after the data has been unloaded, the utility job ends and the status in SYSIBM.SYSUTIL corresponding to this utility ID is removed. External - Specifies that after the data has been unloaded, the utility job is to end and the status in SYSIBM.SYSUTIL corresponding to this utility ID is removed.

Reorganize Tablespace Wizard - Panel 5

NOTE: This panel is displays only if the share level is Change.

The table below describes the options and functionality on this panel of the Reorganize Tablespace Wizard:

Option	Description
Table Owner	Displays the schema names in the system.
Table	Displays all tables for the selected schema.
Maxrow	Lets you specify the maximum amount of time for the last iteration of log processing. During that iteration, applications have read-only access. Type an integer value or select DEFER. DEFER Specifies that the iterations of log processing with read/write access can continue indefinitely.
Drain	Lets you select Writers or All. Specifies that DBArtisan drains the write claim class after the delay. The number of log records, and thus the estimated time, for a future iteration of log processing will be 0.
Long Log	Lets you specify the action that DBArtisan performs (after sending the LONGLOG message to the console) if the number of log records that are processed during the next iteration is not sufficiently lower than the number of log records that were processed during the previous iterations. Continue - Specifies that DBArtisan continues performing reorganization. Term - Specifies that DBArtisan terminates reorganization after the delay. Drain - Specifies that DBArtisan drains the write claim class after the delay. The number of log records, and thus the estimated time, for a future iteration of log processing will be 0.
Delay	Lets you type a lower bound for the interval between the time when REORG sends the LONGLOG message to the console and the time when REORG performs the action specified by the LONGLOG parameter.
Timeout	Lets you select timeout option: Abend - Abnormal end of task Term - Termination

Reorganize Tablespace Wizard - Panel 6

NOTE: This panel is displays only if the share level is Reference or Change.

The table below describes the options and functionality on this panel of the Reorganize Tablespace Wizard:

Option	Description
Local Site Primary Copy	Lets you specify the Local Site Primary Copy.
Local Site Backup Copy	Lets you specify the Local Site Backup Copy.
Recovery Site Primary Copy	Lets you specify the Recovery Site Primary Copy.
Recovery Site Backup Copy	Lets you specify the Recovery Site Backup Copy.

Reorganize Tablespace Wizard - Panel 7

The table below describes the options and functionality on this panel of the Reorganize Tablespace Wizard:

Option	Description
Off Pos Limit	The specified value is compared to the result of the calculation (NEAROFFPOSF + FAROFFPOSF) × 100 / CARDF for the specified partitions in SYSIBM.SYSINDEXPART for the explicit clustering indexes for every table in the specified table space. Alternatively, the values in SYSINDEXPART are checked for a single non-partitioned table space, or for each partition if you specified an entire partitioned table space as the target object. If at least one calculated value exceeds the OFFPOSLIMIT value, REORG is performed or recommended. **NOTE:** This option is valid for non-LOB table spaces only.
Ind Ref Limit	The specified value is compared to the result of the calculation (NEARINDREF + FARINDREF) × 100 / CARDF for the specified partitions in SYSIBM.SYSTABLEPART for the specified table space. Alternatively, the values in SYSTABLEPART are checked for a single non-partitioned table space, or for each partition if you specified an entire partitioned table space as the target object. If at least one calculated value exceeds the INDREFLIMIT value, REORG is performed or recommended. **NOTE:** This option is valid for non-LOB table spaces only.
Report Only	The reorganization is only recommended, not performed.
Unload	Specifies whether the utility job is to continue processing or end after the data is unloaded. Continue - Specifies that, after the data has been unloaded, the utility continues processing. Pause - Specifies that after the data has been unloaded, processing ends. Only - Specifies that after the data has been unloaded, the utility job ends and the status in SYSIBM.SYSUTIL corresponding to this utility ID is removed. External - Specifies that after the data has been unloaded, the utility job is to end and the status in SYSIBM.SYSUTIL corresponding to this utility ID is removed.

Reorganize Tablespace Wizard - Panel 6

The table below describes the options and functionality on this panel of the Reorganize Tablespace Wizard:

Option	Description
Do you want to keep the dictionary?	A collection of language-related linguistic information used during text analysis, indexing, retrieval, and highlighting of documents in a particular language. **NOTE:** Available only if Continue or Pause option was specified for Unload on the previous panel.
Do you want to specify statistics?	Lets you include statistics for the tablespace in the SQL statement. Enabled only if Continue or Pause option was specified for Unload on the previous panel.

Option	Description
Do you want to output message to SYSPRINT?	Lets you specify to output message to SYSPRINT.
Do you want to force aggregation or rollup processing to be done even though some parts do not contain data?	Lets you specify to process with forced aggregation, gathered into a mass, or rollup.

Reorganize Tablespace Wizard - Statistics Tables

The table below describes the options and functionality on this panel of the Reorganize Tablespace Wizard:

Option	Description
Select Tables	Lets you select tables (and columns) to collect statistics for. The sample column is editable and you can type an integer value into this column.
Add	Opens the Table Selector Dialog Box.
Set Properties	Lets you customize the selection of columns in the selected tables.

Reorganize Tablespace Wizard - Statistics Indexes

The table below describes the options and functionality on this panel of the Reorganize Tablespace Wizard:

Option	Description
Select Index	Lets you specify the indexes and the correlation values for collecting the statistics. The sample column is editable and you can type an integer value into this column.
Add	Opens the Index Selector Dialog Box.
Set Properties	Lets you specify the correlation information for each index.

Reorganize Tablespace Wizard - Panel 7

The table below describes the options and functionality on this panel of the Reorganize Tablespace Wizard:

Option	Description
Punch DDN	Lets you select punch DDN from the list. Contains generated LOAD statements that let you reload the discarded data. Required if Unload External option is chosen on Panel 5.
Discard DDN	Lets you select discard DDN from the list.
Unload DDN	Lets you select DD name of the unload data set. Required unless NOSYSREC on Panel 2 or Share level CHANGE Panel 3 is selected.
Work DDN	Lets you select the DD statement for a temporary data set used for intermediate output.

>

Reorganize Tablespace Wizard - Panel 7

The table below describes the options and functionality on this panel of the Reorganize Tablespace Wizard:

Option	Description
Discard	Lets you select discard DDN from the list.
No Pad	DBArtisan does not add a character to fill empty space.
Add	Opens the Table Selector Dialog Box.
Set Condition	Lets you customize the selection of columns in the selected tables.

Reorganize Tablespace Wizard - Panel 8

The table below describes the options and functionality on this panel of the Reorganize Tablespace Wizard

Option	Description
Would you like to specify the device type?	Lets you specify the device type.
Do you want to specify the sort number?	Lets you type the number of datasets to be dynamically allocated by the sort program. The range of the text field is 0 to 2147483647.
Preformat	Lets you specify that the remaining pages are preformatted up to the high allocated RBA in the table space and index spaces associated with the table specified in table-name. The preformatting occurs after the data has been loaded and the indexes are built.

Reorganize Tablespace Wizard - Panel 9

The table below describes the options and functionality on this panel of the Reorganize Tablespace Wizard:

Option	Description
Punch DDN	Lets you select punch DDN from the list. Contains generated LOAD statements that let you reload the discarded data. Required if Unload External option is chosen on Panel 5.
Discard DDN	Lets you select discard DDN from the list.
Unload DDN	Lets you select DD name of the unload data set. Required unless NOSYSREC on Panel 2 or Share level CHANGE Panel 3 is selected.
Work DDN	Lets you select the DD statement for a temporary data set used for intermediate output.

Condition Dialog Box

The Condition Dialog Box lets you type free form condition text.

NOTE: DBArtisan does not test the correctness of the condition text.

Report

NOTE: This functionality is available for DB2 LUW, DB2 z/OS, 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, DBArtisan generates these reports in HTML. You can view these reports directly in DBArtisan's built-in HTML browser.

Completing the Generate Report Dialog Box

To complete the Generate Report dialog box, do the following:

1. On the **Datasource Explorer**, select the target object node.

   DBArtisan displays the target objects in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Report**.

   OR

   In the right pane of the application, right-click the target object, and then select **Report**.

   DBArtisan opens the Generate Report dialog box.

Generating a Summary Report

1. Select the **Summary Report (Listing all items)** option.

   DBArtisan 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.

   DBArtisan 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.

   DBArtisan displays Object Name and Object Type.

4. Click **Execute**.

   DBArtisan displays the report. You can find data in the report, save, send, and print the report.

5. To exit the report, click **Close**.

Restart

NOTE: This functionality is available for IBM DB2 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.

For more information, see Creating objects.

Runstats Index

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS.

The Runstats Index Wizard gathers statistics on indexes. If no statistics are available, DB2 makes fixed default assumptions which may result in inefficiencies. To ensure the effectiveness of the paths selected, you can run the wizard:

- After a table is loaded.
- After an index is physically created.
- After running Rebuild Index (DB2 z/OS) or Reorganize without collecting inline statistics.
- Before running Reorganize with the OFFPOSLIMIT, INDREFLIMIT, or LEAFDISTLIMIT options.

For more information, see Runstats Index Wizard - Panel 1.

Runstats Index Wizard - Panel 1

The table below describes the options and functionality on this panel of the Runstats Index Wizard:

Option	Description
Select indexes	Displays the available indexes. The Partition Column is editable and can contain one or more numbers separated by commas or, a range of numbers separated by:. (e.g. 1,2,3 or 1:3 or 1 are valid entries). The maximum value for the partition column is 254 and the minimum value is 0.
Add	Opens the Index Selector Dialog Box.
Delete	Deletes the selected items.
Set Properties	Opens the Set Correlation Option Dialog Box for the selected row.

Runstats Index Wizard - Panel 2

The table below describes the options and functionality on this panel of the Runstats Index Wizard:

Option	Description
Select tablespaces	Displays the available tablespaces. The Partition Column is editable and can contain one or more numbers separated by commas or, a range of numbers separated by:. (e.g. 1,2,3 or 1:3 or 1 are valid entries). The maximum value for the partition column is 254 and the minimum value is 0.
Add	Opens the Tablespace Selector Dialog Box.
Delete	Deletes the selected items.
Set Properties	Opens the Correlation Statistics dialog box for the selected row.

Runstats Index Wizard - Panel 3

The table below describes the options and functionality on this panel of the Runstats Index Wizard:

Option	Description
Access level while RUNSTATS is running	Lets you select an access level, either Readonly or Change.

Option	Description
Do you want to output message to SYSPRINT?	Lets you specify to output message to SYSPRINT.
Do you want to force aggregation or rollup processing to be done even though some parts do not contain data?	Lets you specify to process with forced aggregation, gathered into a mass, or rollup.

Runstats Index Wizard - Panel 4

The table below describes the options and functionality on this panel of the Runstats Index Wizard:

Option	Description
Update Statistics	Lets you update the catalog statistics.
History Statistics	Lets you update the history statistics.

Set Correlation Option Dialog Box

The table below describes the options and functionality on the Set Correlation Option Dialog Box:

Option	Description
KeyCard	Collects all of the distinct values in all of the 1 to n key column combinations for the specified indexes. n is the number of columns in the index.
Frequent Item	Lets you add items to the grid.
Add	Opens the Add Frequent Item Dialog Box

Add Frequent Item Dialog Box

The table below describes the options and functionality on the Add Frequent Item Dialog Box:

Option	Description
Number of Columns	Lets you type the number of columns. Default is 1.
Count	Lets you type the count. Default is 10.

Runstats Tablespace

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS.

The Runstats Tablespace Wizard gathers statistics on tablespaces, indexes, and columns. DBArtisan records this information in the DB2 catalog and uses it to select access paths to data during the bind process. The wizard lets you evaluate database design and helps you decide to reorganize tablespaces or indexes.

For more information, see Runstats Tablespace Wizard - Panel 1.

>

Runstats Tablespace Wizard - Panel 1

The table below describes the options and functionality on this panel of the Runstats Tablespace Wizard:

Option	Description
Database Name	Lets you select a database.
Tablespace Name	Lets you select a tablespace.
Partition Number	Lets you specify on which the command needs to be run.
Do you want to runstats on all tables in this tablespace?	Default is Yes.
Do you want to runstats on all indexes in this tablespace?	Default is Yes.

Runstats Tablespace Wizard - Panel 2

NOTE: This panel displays if you selected No for either option on Panel 1.

The table below describes the options and functionality on this panel of the Runstats Tablespace Wizard:

Option	Description
Table Sample	Lets you type a table sample. Default is 25. You can enter any value between 1 and 100.
KeyCard	Collects all of the distinct values in all of the 1 to n key column combinations for the specified indexes. n is the number of columns in the index.
Add	Opens the Add Frequent Item Dialog Box
Delete	Deletes selected item(s.)

Runstats Tablespace Wizard - Panel 3

NOTE: This panel displays if you selected No for the Runstats on All Tables option on panel 1.

The table below describes the options and functionality on this panel of the Runstats Tablespace Wizard:

Option	Description
Select Tables	Lets you select the tables in the table that runstats will be run on. The Sample column is blank initially for each row. It is editable with values between 1 - 100. The Indicator column is blank initially for each new row that is added. However if you set the custom columns for a row (via set properties button), this column shows an asterisk (*).
Add	Opens the Table Selector Dialog Box.
Delete	Deletes the selected table(s.)
Set Properties	Lets you customize the selection of columns in the selected table(s.)

>

Runstats Tablespace Wizard - Panel 3

The table below describes the options and functionality on this panel of the Runstats Tablespace Wizard:

Option	Description
Select Indexes	Lets you select index(es) for the runstats operation.
Add	Opens the Select Column for Table Dialog Box.
Delete	Deletes selected index(es).
Set Properties	Opens the Set Correlation Option Dialog Box.

Select Column for Table Dialog Box

The table below describes the options and functionality on the Select Column for Table Dialog Box:

Option	Description
Column Name	Displays the column name.
Included	Select to include column.
Select All	Click to select all columns.
Unselect All	Click to deselect all columns.

Runstats Tablespace Wizard - Options

The table below describes the options and functionality on this panel of the Runstats Tablespace Wizard:

Option	Description
Access level while RUNSTATS is running	Lets you select an access level, either Readonly or Change.
Do you want to output message to SYSPRINT?	Lets you specify to output message to SYSPRINT.
Do you want to force aggregation or rollup processing to be done even though some parts do not contain data?	Lets you specify to process with forced aggregation, gathered into a mass, or rollup.

Runstats Tablespace Wizard - Statistics

The table below describes the options and functionality on this panel of the Runstats Tablespace Wizard:

Option	Description
Update Statistics	Lets you update the catalog statistics.
History Statistics	Lets you update the history statistics.

Table Selector Dialog Box

The table below describes the options and functionality on the panel of the Table Selector dialog box:

Option	Description
Tablespace Like	Displays the tablespace.
Table Creator	Displays the table(s) creator of the tablespace that was selected on Panel 2.
Table Name	Displays the table(s) of the tablespace that was selected on Panel 2.
All Columns	Selects all columns.
Table Selector	Lets you select one or more rows per table, up to 10.
Query	Repopulates Table Selector based on the tablespace query.
Select All	Selects all rows in Table Selector.
Unselect All	Unselects all rows in Table Selector.
Add	Click to add the table to the wizard.

Index Selector Dialog Box

The table below describes the options and functionality on the panel of the Index Selector dialog box:

Option	Description
Index Creator	Displays the index(es) creator of the tablespace that was selected on Panel 2.
Index	Displays the index(es) of the tablespace that was selected on Panel 2.
Index Selector	Lets you select one or more rows in this list.
Query	Repopulates Index Selector based on the tablespace query.
Select All	Selects all rows in Index Selector.
Unselect All	Unselects all rows in Index Selector.
Add	Click to add the index to the wizard.

Schema

NOTE: This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Schema dialog box lets you view everything that makes up a selected table.

To view schema, do the following:

1 On the **Datasource Explorer**, select the target object node.

DBArtisan displays the target objects in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Schema**.

OR

In the right pane of the application, right-click the target object, and then select **Schema**.

DBArtisan opens Schema window.

>

Select * From

NOTE: This functionality is available for DB2 LUW, DB2 z/OS, Oracle, SQL Server, and Sybase.

The Select * From dialog box lets you retrieve all data from a selected table or view.

Completing the Select * From Dialog Box

To complete this dialog box, do the following:

1 On the **Datasource Explorer**, select the target object node.

DBArtisan displays the target objects in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Select * From**.

OR

In the right pane of the application, right-click the target object, and then select **Select * From**.

3

DBArtisan opens the Select * From dialog box.

For more information, see Creating objects.

Set Default

This function lets you set a tablespace as the default tablespace. Users created without a specified default tablesapce 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
Action options	Verify that the panel displays the tablespace to be set as default.
Dependencies	Review the referring and referred objects that will be automatically resolved when you execute this operation.
Preview	Preview the DDL generated for the operation and when ready, use the **Schedule** or **Execute** button to perform this action.

Set Integrity

NOTE: This functionality is available for IBM DB2 only.

For a complete and in-depth discussion of this topic, refer to IBM's documentation on the subject:
http://www.db2.jp/db2manual/en_US/index.htm?openup=admin/r0000998.htm

To Open the Set Integrity Wizard

1 On the **Explorer**, expand a DB2 database, and click the **Table** node.

2 Select the appropriate table(s), right-click, and choose **Set Integrity** from the menu. You can select more than one table by clicking CTRL + the tables you want.

OR

Select the appropriate table, click **Command** on the toolbar, and then choose **Set Integrity** from the drop-down menu.

The table below lists all fields you can see in the Set Integrity Wizard. Depending on the options you choose, you may not see them all.

Required Field	Description
Tables	
Tables	The list of one or more tables you selected appear.
Integrity Option	
Integrity Option	OFF--When selected, tables have their foreign key and check constraints, and column generation disabled and so are put in a pending state. This also means materialized query or staging tables are not immediately refreshed and are put in a pending state. TO DATALINK RECONCILE PENDING--When selected, DATALINK integrity is disabled and tables are put in a check pending no access state. Dependent and descendant tables are not affected. IMMEDIATE CHECKED--This turns on a table's integrity checking turned on. Any checking that was deferred is carried out. FULL ACCESS--Tables become fully accessible as soon as the SET INTEGRITY statement executes. PRUNE--This is appropriate only for staging tables. The content of the staging table is pruned and set to an inconsistent state. If any table in the table-name list is not a staging table, an error is returned. UNCHECKED--Allows you turn on some or all integrity checking but the table will not be checked for integrity violations. This can affect data integrity.
Table Readability/Cascade/Descendent Types	
Specifies the readability of the table while it is in check pending state:	NO ACCESS--The table(s) are put in a check pending no access state so read/write access to the table is prohibited. READ ACCESS--The table(s) are put in a check pending read state. This allows read access to the non-appended portions of any tables.
Specifies whether to be immediately cascaded to all descendents	CASCADE IMMEDIATE--The check pending state for foreign key constraints is immediately extended to descendant foreign key constraints or to materialized query or staging tables. CASCADE DEFERRED--Only the selected tables are put in the check pending state. Descendant foreign key, materialized query, or staging tables remain unchanged.
Descendent Types	Materialized Query Tables--When selected, the check pending state is immediately cascaded to only descendant materialized query tables. Foreign Key Tables--When selected, the check pending state is cascaded immediately only to descendant foreign key tables. Staging Tables--When selected, the check pending state is cascaded immediately only to descendant staging tables.
Check Appended Portion?	

Required Field	Description
Do you want to check on the appended portion (if any) of the table?	Default/Yes/No Force Generated--If you do not specify this generated column current values will be compared to the computed value of the expression as if an equality check constraint existed. If this is specified, generated columns are computed only for the appended portion. Prune--Possible only for staging tables. When you check this, the contents of the staging table are pruned and the staging table is set to an inconsistent state. Full Access--When selected, tables will become accessible after the SET INTEGRITY statement executes.
Specify Exception Tables	
List of Base Tables	Any row that is in violation of a foreign key or check constraint is copied to the exception table you select.
Integrity Options: IMMEDIATE UNCHECKED	
IMMEDIATE UNCHECKED options	Foreign Key--These constraints are turned on when the table is removed from check pending status. Check--Check constraints are turned on when the table is removed from check pending status. Datalink Reconcile Pending--DATALINK integrity constraints are turned on when the table is removed from check pending status. Materialized Query--Immediate refreshing is turned on for a materialized query table when it is removed from a check pending state. Generated Column--When the table is removed from check pending status, generated columns are turned on. Staging--Immediate propagation is turned on for a staging table.
Do you want tables to become fully accessible after the SET INTEGRITY statement executes?	Yes/No
Preview DDL	
DDL View	The DDL window displays the command code you created by completing the prior panels of the wizard. You have the ability to save the code to disk, to print it, to email it, or to schedule the execution of the task. The code itself is read-only and can only be changed by returning to prior panels of the wizard.

Set Online/Offline

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.

NOTE: This functionality is available for Microsoft SQL Server, Oracle, and Sybase ASE only.

Important Notes

For Sybase, DBArtisan only lets you set databases online.

To set one or more databases online or offline:

1. Connect to the datasource that contains the database. For more information, see Datasource Management.
2. On the **Datasource Explorer**, select the **Databases** node.

   DBArtisan displays the Databases in the right pane of the Explorer window.

3 In the right pane of the Explorer window, select one or more target databases.

4 On the **Datasource Explorer** tool bar, click **Command**, and then select **Set Online/Offline**.

DBArtisan opens the **Checkpoint** dialog box.

5 Use the following table as a guide to working through the panels of the dialog box

Step	Settings and tasks	
Action options	**Set offline**	Set this to TRUE to set the database offline or set it to True to set the database online.
Dependencies		Lets you review any dependecies before you proceed.
Preview		Displays the DDL generated to execute the Checkpoint operation.

6 Use one of the **Schedule** or **Execute** buttons to execute the Checkpoint.

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.

DBArtisan 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**.

DBArtisan 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
Action Options	Displays the tablespace you selected to be set as an Undo tablespace.
Dependencies	View the dependencies on the tablespace.
Preview	Preview the DDL generated from your choices.

3 Click **Execute**. For information on the other options, see Preview and Scheduling.

Shrink

NOTE: The Shrink functionality is available for Microsoft SQL Server and Oracle.

Microsoft SQL Server

The Shrink Database dialog box lets you reclaim space from a database that is too large.

>

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.
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: Shrink Rollback Segments

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 lets you shrink a rollback segment manually by a specific amount or back to its Optimal Size.

Important Notes

For Oracle 9 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.

Oracle: Shrink Tables or Indexes

The Shrink dialog box lets you shrink the size of tables or indexes.

Important Notes

Shrink is only available for tables in which the Row Movement property is set to TRUE.

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

Step	Functionality
Action Options	Enable **Compact** to restrict shrinking to defragmenting the segment space and compact rows. Enable **Cascade** to simultaneously shrink all dependent objects.
Dependencies	View the dependencies on the table or index that you are shrinking.
Preview	Preview the DDL generated from your choices.

For more information, see Creating objects.

Start Database

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS only.

The Start Database dialog box lets you start a database:

- When the database that has been stopped with a **Stop Database** dialog box. For more information, see Stop Database.
- After a tablespace, partition, or index has been placed in group buffer pool RECOVER pending status (GRECP) or if pages have been put on the logical page list (LPL) for that object.

Depending on the specified options, the database can be made available for read-only processing, read-write processing, or utility-only processing. In a data sharing environment, the command can be issued from any DB2 on the group that has access to the database.

Important Notes

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

Option	Functionality
Database Grid	Displays currently stopped databases. The Partition column is editable and can include one or more unique numeric values separated by commas. Initially the Partition column is blank. No validation is made for the correctness of partition numbers so make sure that the partitions exist.
Access	Lets you select access options.

For more information, see Creating objects.

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.

For detailed instructions, 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.

Completing the Set Constraint(s) Status Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the target object node.

   DBArtisan displays the target objects in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Status**.

   OR

   In the right pane of the application, right-click the target object, and then select **Status**.

   DBArtisan opens the Set Constraint(s) Status dialog box.

3. Use the table above to select dialog box options on each panel.

4. Use the **Execute** or **Schedule** button to complete the operation.

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. DBArtisan 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.
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 is the parent table is deleted.

Completing the Set Constraint(s) Status Dialog Box

To complete this dialog box, do the following:

1 On the **Datasource Explorer**, select the target object node.

DBArtisan displays the target objects in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Status**.

OR

In the right pane of the application, right-click the target object, and then select **Status**.

DBArtisan opens the Set Constraint(s) Status dialog box.

3 Use the table above to select dialog box options.

4 Use the **Execute** or **Schedule** button to complete the operation.

Stop Database

NOTE: This functionality is available for IBM DB2 for OS/390 and z/OS only.

The Start Database dialog box lets you stop a database:

Important Notes

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

Option	Functionality
Database Grid	Displays currently running databases. The Partition column is editable and can include one or more unique numeric values separated by commas. Initially the Partition column is blank. No validation is made for the correctness of partition numbers so make sure that the partitions exist.

Summary Definition

NOTE: This functionality is available for IBM DB2 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.

Option	Description
Summary table Options	Lets you set able 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 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.

Completing the Switch Online Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the target object node.

   DBArtisan displays the target objects in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Switch Online**.

   DBArtisan opens the Switch Online dialog box.

3. Click **Execute**. For information on the other options, see Preview and Scheduling.

Transfer Ownership

DBArtisan 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
Indexes	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 datasource 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 datasource nodes until the target object type node is visible and then select that node. DBArtisan 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 datasource explorer, select one or more objects, right-click and select **Transfer Ownership** from the context menu. DBArtisan 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
Action options	If you inititted the ownership transfer against a user, select the specific objects belonging to that user that are to be assigned to a new user. Use the **New Owner** dropdown to select a defined user as the new owner of the selected object or objects.
Preview	Preview the DDL generated for the operation and when ready, click **Execute**. For information on the other options, see Preview and Scheduling.

Truncate

NOTE: The Truncate functionality is available for IBM DB2 for Linux, Unix, and Windows, Microsoft SQL Server, Oracle, IBM DB2 for OS/390 and z/OS and Sybase ASE.

The Truncate dialog box lets you quickly delete the rows of a table.

Important Notes

If you truncate a table, DBArtisan 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.

Completing the Truncate Dialog Box

To complete this dialog box, do the following:

1 On the **Datasource Explorer**, select the target object node.

DBArtisan displays the target objects in the right pane of the Explorer window.

2 On the **Datasource Explorer** tool bar, click **Command**, and then select **Truncate**.

OR

In the right pane of the application, right-click the target object, and then select **Truncate**.

DBArtisan opens the Truncate dialog box.

3 Click **Execute**. For information on the other options, see Preview and Scheduling.

For more information, see:

Truncating IBM DB2 Objects

Truncating Microsoft SQL Server Objects

>

Truncating IBM DB2 Objects

The Truncate dialog box lets you quickly delete the rows of a table.

NOTE: If you truncate a table, DBArtisan 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.

Truncating Microsoft SQL Server Objects

The Truncate dialog box lets you quickly delete the rows of a table.

NOTE: If you truncate a table, DBArtisan 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.

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	Before truncating a cluster containing a parent key, disable all referencing foreign keys existing in different tables. Truncating a cluster does not generate any rollback information and commits immediately. Oracle alters the storage parameter NEXT to the size of the last extent deleted from the segment.
Table	Before truncating a table containing a parent key, disable all referencing foreign keys existing in different tables. Truncating a table does not generate any rollback information and commits immediately. Oracle alters the storage parameter NEXT to the size of the last extent deleted from the segment. Oracle automatically deletes all data in the table's indexes and any materialized view direct-load INSERT information associated with a truncated table. If the table is not empty, all associated nonpartitioned indexes and all partitions of associated global partitioned indexes are marked unusable. You cannot truncate a hash cluster nor can you truncate individual tables in a hash cluster or an index cluster.

>

The table below describes the options and functionality on the Truncate dialog box.

Option	Description
Storage Handling	Drop Storage - Select if you want the freed extents returned to the system where they can be used by other objects. Reuse Storage - Select if you want the space to remain allocated to the table or cluster you have just truncated.

For more information, see Truncate.

Truncating DB2 z/OS Objects

The Truncate dialog box lets you quickly delete the rows of a table.

CAUTION: If you truncate a table, DBArtisan 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
Reuse Storage	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.
Restrict When Delete Triggers	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.
Immediate	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 Rollbackcan undo the Truncate operation.

For more information, see Truncate.

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.

Unload Data Utility Wizard

The Unload Data Utility Wizard lets you export table data. After defining the data export operation, you have the opportunity to save the job definition for future reuse.

The table below describes the options and functionality on the panels of the Unload Data Utility Wizard:

Panel	Option	Description
Panel 1	Database	Lets you specify the database for the unload operation.
	Tablespace	Lets you specify the tablespace for the unload operation.
	Partition	Lets you specify the partition for the unload operation.
	Specify the from copy data set name	Lets you specify the from copy data set.
	Would you like to specify volume information?	Lets you specify volume information.
Panel 2	Select tables	Lets you specify the tables for the unload operation.
Panel 3	Specify the PUNCH DDN	Lets you specify the punch to unload the object list to.
	Specify the UNLOAD DDN	Lets you specify the unload data set templates to unload the object list to.
	Encoding type	Lets you specify the encoding scheme for the output data set and select the appropriate radio button. If you do not specify either EBCDIC, ASCII, UNICODE, or CCSID, the encoding scheme of the source data is preserved.
	CCSID(s)	Lets you specify the up to three coded character set identifiers (CCSIDs) to be used for the data in the output data set, including data unloaded in the external character formats. The CCSID identifies the way in which the characters in the string are encoded.
Panel 4	Perform CCSID code substitution during unload	When a string is converted from one CCSID to another (including EBCDIC, ASCII, and UNICODE), a substitution character is sometimes placed in the output string, for example, when a character that exists in the source CCSID does not exist in the target CCSID.
	Use padding with variable length data type	Select No to ensure that variable length columns in the unloaded records occupy the actual data length without additional padding. As a result, the unloaded or discarded records may have varying lengths.
	Float	Lets you specify the output format for numeric floating point data. You may select either Hexadecimal (s390) or Binary (IEEE). This option applies to the binary output format only.
	Max number or records of error allowed	The unloading process terminates when this value is reached.
	Sharelevel	Lets you specify the type of application access allowed during the unloading process.

Unquiesce Database

NOTE: This functionality is available for IBM DB2 for Windows/Unix 8.1 only. DBArtisan does not support Unquiesce Database for an IBM DB2 for Windows/Unix 8 server from a IBM DB2 for Windows/Unix 7 client or for an IBM DB2 for Windows/Unix 7 server from an IBM DB2 for Windows/Unix 7 client.

The Unquiesce Database dialog box lets you restore user access to a database.

Completing the Unquiesce Database Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the **Datasources** node.

   DBArtisan displays the Databases in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Unquiesce**.

   OR

   In the right pane of the application, right-click the target object, and then select **Unquiesce**.

   DBArtisan opens the Unquiesce Database dialog box.

3. Click **Execute**.

For more information, see Creating objects.

Unquiesce Instance

NOTE: This functionality is available for IBM DB2 for Windows/Unix 8.1 only.

The Unquiesce Instance dialog box lets you restore user access to an instance.

Completing the Unquiesce Instance Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the **Datasources** node.

   DBArtisan displays the Instances in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **Unquiesce**.

   OR

   In the right pane of the application, right-click the target object, and then select **Unquiesce**.

   DBArtisan opens the Unquiesce Instance dialog box.

3. Click the following:

Button	Description
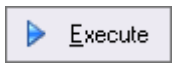	Executes the task.

For more information, see Creating objects.

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.

The following topics provide details on updating statistics by supported DBMS and object type:

- Updating statistics for tables or indexes (IBM DB2 for Linux, Unix, and Windows)
- Updating statistics for views (IBM DB2 for Windows, Unix, and Linux)
- Updating statistics for databases, indexes, or tables (Microsoft SQL Server)
- Updating statistics for indexes or tables (Sybase)

Updating statistics for tables or indexes (IBM DB2 for Linux, Unix, and Windows)

To update statistics for an index or table:

1 Connect to the datasource that contains the index or table. For more information, see Datasource Management.

2 On the **Datasource Explorer**, expand datasource nodes until the **Tables** or **Indexes** node is visible and then select the **Indexes** or **Tables** node.

DBArtisan displays the objects of that type in the right pane of the Explorer window.

3 In the right pane of the Explorer window, select the specific, target table or index.

4 On the **Datasource Explorer** tool bar, click **Command**, and then select **Update Statistics**.

DBArtisan opens the **Update Statistics** dialog box.

5 Use the following table as a guide to understanding and modifying the settings in this dialog.

Tab	Options	Description
Table Options	**Update table statistics**	Updates table statistics.
	Distribution Options	**Do not collect column statistics** - Column statistics provide information that the optimizer uses to choose the best access plans for queries. **Collect column statistics on key columns only** - Collects column statistics on columns that make up all the indexes defined on the table. **Collect column statistics on all columns** - Collects column statistics for all columns. Column statistics provide information that the optimizer uses to choose the best access plans for queries. **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.
	Column Options	**Do not collect distribution statistics** - Does not collect basic statistics or distribution statistics on the columns. **Collect distribution statistics on key columns only** - Collects both basic statistics and distribution statistics on key columns only. **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. **Exclude XML columns** - Lets you collect statistics on non-XML columns only. XML columns are not included in statistics collection.

Tab	Options	Description
Index Options	**Update index statistics**	Lets you enable and disable statatistics updates for indexes and controls the following settings: **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** - DBArtisan 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 DBArtisan updates the statistics.
	Allow read/write access during collection	Allows read and write access while DBArtisan updates the statistics.

6 Click **Execute**. For information on the **Preview** option, see Preview.

Updating statistics for views (IBM DB2 for Windows, Unix, and Linux)

For IBM DB2 for Windows, Unix, and Linux (version 9), DBArtisan lets you update the statistics for a view.

To update statistics for a view:

1 On the **Datasource Explorer**, select a **Views** node.

DBArtisan displays the views in the right pane of the Explorer window.

2 In the right pane, right-click a view and then select **Update Statistics** from the context menu.

The Update Statistics dialog opens.

3 Use the following table as a guide to choosing options:

Pane	Options
Action Options	Lets you enable or disable column statistics update, distribution statistics update as well as frequency and quantiles options, and enable or disable write access.
Dependencies	Lets you view referencing or referenced objects.
Preview	Lets you view the DDL generated for the operation.

4 **Schedule** or **Execute** the statistics update.

Updating statistics for databases, indexes, or tables (Microsoft SQL Server)

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.

To update statistics for a database, index, or table:

1. Connect to the datasource that contains the database, index, or table. For more information, see Datasource Management.
2. On the **Datasource Explorer**, expand datasource nodes until the **Databases**, **Tables**, or **Indexes** node is visible and then select the **Databases**, **Indexes**, or **Tables** node.

   DBArtisan displays the objects of that type in the right pane of the Explorer window.
3. In the right pane of the Explorer window, select the specific, target database, table, or index.
4. On the **Datasource Explorer** tool bar, click **Command**, and then select **Update Statistics**.

   DBArtisan opens the **Update Statistics** dialog box.
5. Use the following table as a guide to understanding and modifying the settings in this dialog:

Step	Settings and tasks	
Action Options	**Scan Range** (tables and indexes only)	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.
	Sample Count (tables and indexes only)	If you specified a **Scan Range** of **Sample Scan**, provide a count.
	Sample Unit (tables and indexes only)	If you specified a **Scan Range** of **Sample Scan**, specify either **%** or **Rows**.
	Statistics Type (tables only)	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.
	Statistics Recompute (tables and indexes only)	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.
Dependencies	Lets you view referencing or referenced objects.	
Preview	Lets you view the DDL generated for the operation.	

6. Use the **Execute** or **Schedule** button to perform the operation.

Updating statistics for indexes or tables (Sybase)

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.

To update statistics for a database, index, or table:

1. Connect to the datasource that contains the index or table. For more information, see Datasource Management.
2. On the **Datasource Explorer**, expand datasource nodes until the **Tables** or **Indexes** node is visible and then select the **Indexes** or **Tables** node.

   DBArtisan displays the objects of that type in the right pane of the Explorer window.
3. In the right pane of the Explorer window, select the specific, target table or index.
4. On the **Datasource Explorer** tool bar, click **Command**, and then select **Update Statistics**.

   DBArtisan opens the **Update Statistics** dialog box.
5. Use the following table as a guide to understanding and modifying the settings in this dialog:

Step	Settings and tasks	
Action Options	Index (tables only)	Enabling this option updates statistics for indexes of the table.
	Partition (tables only, Sybase version 12)	Enabling this option updates statistics for partitions of the table.
	Table (tables only)	Enabling this option updates table-specific statistics. Column statistics stored in sysstatistiics are not affected.
Dependencies	Lets you view referencing or referenced objects.	
Preview	Lets you view the DDL generated for the operation.	

6. Use the **Execute** or **Schedule** button to perform the operation.

View Contents

> **NOTE:** This functionality is available for Microsoft SQL Server only.

The Backup Device Information dialog box lets you view the contents of a Dump or Backup Device which stores backups of databases and transaction logs.

The table below describes the options and functionality on the Backup Device Information dialog box.

Option	Description
Backup Device	Select the tape or disk used in a backup or restore operation.

Completing the Backup Device Information Dialog Box

To complete this dialog box, do the following:

1. On the **Datasource Explorer**, select the target object node.

   DBArtisan displays the target objects in the right pane of the Explorer window.

2. On the **Datasource Explorer** tool bar, click **Command**, and then select **View Contents**.

   OR

   In the right pane of the application, right-click the target object, and then select **View Contents**.

   DBArtisan opens the Backup Device Information dialog box.

3. **Close** the dialog.

SQL Scripting

DBArtisan 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.

DBArtisan's scripting environment is comprised of different windows:

- ISQL Editor
- DDL Editors
- SQL preprocessing: #define and #include
- Automated error detection and coding aid features
- 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, DBArtisan 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.

NOTE: When you modify an ISQL window without performing a SAVE or RENAME operation, you will see an * in the tab, such as SQL 2*. The asterisk disappears as soon as you create a unique name for the script.

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: For Oracle, DBArtisan displays REF CURSOR contents in the ISQL Window and Results Editor.

TIP: To toggle to the next SQL window, press CTRL +T.

TIP: For IBM DB2 for Linux, Unix, and Windows and Oracle, you can access the Code Completion functionality with the CTRL+K shortcut.

Related Topics

- Toolbar Options

- ISQL Windows
- DDL Editors
- Open Files
- Insert File into Current File
- Splitting Windows
- Find and Replace
- Regular Expressions
- Goto
- Scheduling
- Sending SQL Scripts
- Renaming and Closing Query Window Tabs
- Print
- Saving and Closing Scripts

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 datasource connection. When locked, the ISQL Window does not respond to datasource 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 Using the Step Execution Facility.
Schedule	Opens a dilaog that lets you schedule an SQL job. For more information, see Scheduling.
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 **Pre-process Only**, 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 **Pre-process and Execute**, all #include and #define directives are carried out and the script is executed. No new editor window opens. For more information, see SQL preprocessing: #define and #include.
Script Execution Facility	Provides a shotcut to opening the Script Execution Facility. For more information, see Executing Scripts Against Multiple Datasources.
Query Plan	Activates and deactivates Query Plan mode. For more information, see Using the Query Plan Facility.

Option	Description
Query Options	Provides a shortcut to opening the Query Options dialog. For more information, see Query options.
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 Syntax checking.
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. DBArtisan 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**.

DBArtisan 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)
- Keyboard Emulation Mode - Lets you customize your editing environment to an emulation mode.

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 emulation node on the Editor Tab of the Options Editor. You can modify the Oracle Auto Commit status and ISQL Tab of the Options Editor.

DDL Editors

NOTE: This functionality is available for all platforms.

DBArtisan 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.

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.

DBArtisan 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, select **New** > **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**.

   DBArtisan 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.

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. DBArtisan 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**.
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.

   DBArtisan 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.

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.

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**.

   DBArtisan opens the Open File(s) dialog box.
2. In the **Open File(s)** dialog box, select the target script, and then click **Open**.

   DBArtisan opens the What type of file dialog box.
3. Select options, and then click **OK**.

   DBArtisan 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.
Files of Type	Lets you control the files displayed in the file list by changing the selection. DBArtisan displays multiple default file extensions.
Open Button	Click to insert one or more files into the current workspace.

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**.

   DBArtisan 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.

   DBArtisan inserts the file.

For more information, see ISQL Editor.

Splitting Windows

You can split an 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.

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	For more information, see Regular Expressions.
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.

Completing the Find Dialog Box

To complete the Find dialog box, do the following

1. On the **Edit** menu, 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.

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	For more information, see Regular Expressions.
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.
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.

Completing the Replace Dialog Box

To complete the Replace dialog box, do the following:

1. On the **Edit** menu, 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 Sinders, 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] -- Matchesnothing ^[A-Z] -- Match lines beginning with an upper-case letter
\	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.

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**.

   DBArtisan opens the Go To dialog box.

2. In **Line Number**, type or select the target line number.

3. Click **OK**.

   DBArtisan 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**.

   DBArtisan 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**.

   DBArtisan 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**.

DBArtisan 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**.

DBArtisan 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**.

DBArtisan 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.

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.

Completing the Print Dialog Box

To complete the Print dialog box, do the following:

1. On the **File** menu, 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. If you have set the **Auto-Save** feature in the Options Editor, a temporary copy of your scripts is automatically saved periodically as you work. For details, see ISQL options.

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**.

   DBArtisan 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**.

DBArtisan closes the Save As dialog box.

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.

For more information, see Saving a Script.

Closing a Script

To close a script, do the following:

1 On the on the **Main** menu, click **Close**.

DBArtisan starts closing the script.

2 If you have not saved your script, DBArtisan prompts you to save the file. Click **Yes** to save and **No** to close without saving.

For more information, see Saving a Script.

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 About semantic validation.
- **Syntactic** - syntax problems associated with the SQL dialect of the datasource 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 Executing Scripts.

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

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

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**.
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.

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.

Option	Description
Object Type Box	DBArtisan displays a column of objects, given your selections of datasource, owner, and object type. Lets you select the check box next to the target object type, or click All to select all. Paste - Click to paste this object into your Editor window.
Middle Box	DBArtisan displays attributes appropriate to the type of object you selected. For example, if you select Tables, DBArtisan fills the second column with the columns of the table you select in the first column. If you select Procedures, DBArtisan fills the second column with available parameters for the procedure you select in the first column. 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. Paste - Click to paste the object properties into your Editor window, under the Object Type box.
Right Box	Displays the SQL statement. Paste - Click when you are satisfied with the entire SQL statement. DBArtisan pastes the SQL statement into your Editor window.

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**.
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.

    the SQL statement in the right box.
15. In the **Right Box**, click **Paste** when you are satisfied with the entire SQL statement.

    DBArtisan 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

DBArtisan 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.

NOTE: For more keyboard commands, see Keyboard Commands.

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**.

DBArtisan comments code.

Uncommenting Code

To uncomment code, do the following:

1 On the **Edit** menu, click **Undo Comment Out**.

DBArtisan 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**.

DBArtisan 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**.

   DBArtisan copies the text.
3. On the **Edit** menu, click **Paste**.

   DBArtisan 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**.

   DBArtisan cuts the text.
3. On the **Edit** menu, click **Paste**.

   DBArtisan 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**.

   DBArtisan 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**.

   DBArtisan 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**.

   DBArtisan 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**.

DBArtisan 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**.

DBArtisan 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**.

DBArtisan 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**.

DBArtisan 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**.

DBArtisan 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
Script Execution Facility	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.
Step Execution Facility	The Step Execution Facility processes batches from one delimiter to the next.
Using the Query Plan Facility	The Query Plan Facility provides a graphical display that lets you identify the execution path that your SQL follows. DBArtisan's Query Plan window now displays data for the estimated costs, number of rows, and bytes returned by each plan step.
Query options	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**.

DBArtisan starts executing the script.

Executing Partial Scripts

To execute a partial script, select the portion (see ISQL options) 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

DBArtisan 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, DBArtisan 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. For example, you can simultaneously execute an ordinary SELECT statement against Oracle, Sybase ASE, Microsoft SQL Server and IBM DB2 for Linux, Unix, and Windows servers.

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 **Tools** menu, 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/ISQL Editor. To execute a script from a DDL or a PL/ISQL 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	Only Connected Datasources - Select to see only your currently connected datasources. All DBMS Types - Select to see all DBMS types. 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. Database - Lets you type the name of the target database.
Output	Specify how you want the output to be handled Graphical Output - Select for graphical output. File Output - Select for file output. 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. Click the File Type list to specify the file type you want to apply to the output file. Include column titles when saving - Select to include column titles in your saved file. Open files with registered applications - Select to open the files with registered application.

Option	Description
Notify	Job Description - Lets you type a job description to supply a subject in your e-mail message. E-mail addresses - Lets you type the e-mail address(es) separated by semi-colons. Net Send User Names - Lets you type the Net Send user name(s) separated by semi-colons.
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 Scheduling.

For more information, see Script Execution Facility

Compilation Error Tab

DBArtisan 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. For more information, see Results Editor.

You can set options on how you want to utilize this feature in the **Editors** tab of the Options Editor. For details, see ISQL options - Editor tab.

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. DBArtisan 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 DBArtisan 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, DBArtisan 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	DBArtisan 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 DBArtisan 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.

Using the Step Execution Facility

DBArtisan offers three ways to use the step execution facility:

1 Open a script.

2 On the **Query** menu, click **Step-Execute**.

DBArtisan starts the procedure and displays errors at the bottom of the Editor window.

NOTE: DBArtisan 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, DBArtisan 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 tree-based or 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, DBArtisan 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.

DBArtisan'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, DBArtisan 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**.

DBArtisan 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 tree-based or graphical Query Plan (Oracle)

For Oracle, DBArtisan 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.

NOTE: The Options editor has a **Default Query Plan** setting that controls whether the default query plan display is tree-based or graphical. For details, see ISQL 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**.

DBArtisan starts the Show Plan mode.

3 To generate the **Show Plan** in a separate result window, click **Execute**.

Options when working with the graphical Query Plan view 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 rom the same script/ISQL window, using the dropdown at the top of the plan to change the plan displayed

Query options

DBArtisan 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 DBArtisan 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, DBArtisan 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 DBArtisan 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	Batch Delimiter	Check syntax when executing	Create Explain plan tables if required	Create explain plan tables on the SYSTOOLS schema
	Isolation Level	Max Errors Before Aborting	Row Count	Run Script with batch execution
IBM DB2 for z/OS	Batch Delimiter	Check syntax when executing	Max Errors Before Aborting	Row Count
	Run Script with batch execution			
Microsoft SQL Server	Abort On Overflow	Batch Delimiter	Force Plan	Ignore Overflow
	Isolation Level	Max Errors Before Aborting	No Count	No Exec
	Parse Only	Row Count	Run Script with batch execution	SET ... Options
	Show Plan	Statistics I/O	Statistics Time	Text Size
MySQL	Batch Delimiter	Big Tables	Client Character	Insert ID
	Interactive Timeout	Last Insert ID	Max Errors Before Aborting	Query Cache Type
	Row Count	SET Other Variables	SQL Auto IS NULL	SQL Big Selects
	SQL Big Tables	SQL Buffer Result	SQL Log Bin	SQL Log Off
	SQL Log Update	SQL Low Priority Updates	SQL Max Join Size	SQL Quote Show Create
	SQL Safe Updates	SQL Select	SQL Warnings	Transaction Isolation
	Unique Checks			

Oracle	Batch Delimiter	Check syntax when executing	Execution Information	I/O Activity
	Index Activity	LONG Size Bytes	Max Errors Before Aborting	Network Activity
	Parse Activity	Row Count	Run Script with batch execution	Sort Activity
	Table Activity			
Sybase	Abort On Overflow	Abort On Truncation	ANSI NULL	Batch Delimiter
	Chained	Force Plan	Ignore Overflow	Index Selection
	Isolation Level	Join Selection	Last Insert ID	LONG Size Bytes
	Max Errors Before Aborting	Network Activity	No Count	No Exec
	Output to Execution Window	Output to Server Error Log	Parse Activity	Parse Only
	Prefetch	Row Count	Run Script with batch execution	Set Quoted Identifier
	Show Plan	Statistics I/O	Statistics Subduer Cache	Statistics Time

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
Abort On Overflow	If set to TRUE, queries will be aborted on encountering a value greater than the specified size.	SQL Server, Sybase
Abort On Truncation	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, DBArtisan truncates the query results and continues processing.	Sybase
ANSI NULL	When set to TRUE, controls results of logical operations with NULL values.	Sybase
Batch Delimiter	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 DBArtisan 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, MySQL, Oracle, SQL Server, Sybase
Big Tables	When set to TRUE, allows big result sets by saving all temporary sets to file. This can slow queries.	MySQL

Option/Variable	Description	Platform
Chained	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
Check syntax when executing	TRUE/FALSE	DB2 LUW, DB2 z/OS, Oracle
Client Character	Default character set.	MySQL
Create Explain plan tables if required	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
Create explain plan tables on the SYSTOOLS schema	If set to TRUE, Explain Plan tables are created on the SYSTOOLS schema. If the tables already exist in the user's default schema, DBArtisan 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
Execution Information	True/False	Oracle
Force Plan	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
I/O Activity	True/False	Oracle
Ignore Overflow	When set to TRUE, DBArtisan will ignore any overflow caused by a resulting value that is larger than a column's specified size.	SQL Server, Sybase
Index Activity	True/False	Oracle
Index Selection	Valuable when tuning query performance.	Sybase
Insert ID	Choose a value to be used a following INSERT or ALTER TABLE statement when you supply an AUTO_INCREMENT value.	MySQL
Interactive Timeout	28800 default	MySQL
Isolation Level	Lets you set DB2 Isolation Levels of UNCOMMITED READ, RESET, CURSOR STABILITY, REPEATABLE READ, or READ STABILITY.	DB2 LUW

Option/Variable	Description	Platform
Isolation Level	**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. **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. 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. **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. Concurrency is lower than Read Committed. Use this option only when necessary. **Serializable**: The highest level of transaction isolation. Transactions are completely isolated from one another. 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. Concurrency is lower than Repeatable Read. Use this option only when necessary.	SQL Server
Isolation Level	0 **1:** Sybase default isolation level. Prevents dirty reads. **2:** Prevents dirty and non-repeatable reads. **3**: Prevents dirty and non-repeatable reads and phantoms. This level is equivalent to performing all selects with holdlock.	Sybase
Join Selection	Valuable when tuning query performance.	Sybase
Last Insert ID	Set the value to be stored in the binary log when you use LAST_INSERT_ID() in a statement that updates a table.	MySQL
LONG Size Bytes	8,192 is the default	Oracle
Max Errors Before Aborting	Select the maximum number of errors encountered before DBArtisan aborts a script. Setting this value to zero disables the feature.	DB2 LUW, DB2 z/OS, MySQL, Oracle, SQL Server, Sybase
Network Activity	True/False	Oracle
No Count	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
No Exec	When set to TRUE, compiles each query without executing it.	SQL Server, Sybase
Output to Execution Window	TRUE/FALSE	Sybase
Output to Server Error Log	TRUE/FALSE	Sybase
Parse Activity	TRUE/FALSE	Oracle

Option/Variable	Description	Platform
Parse Only	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
Prefetch	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
Query Cache Type	The query is cached for ON or DEMAND.	MySQL
Row Count	When set to TRUE, DBArtisan terminates the query after returning the specified number of rows.	DB2 LUW, DB2 z/OS, MySQL, Oracle, SQL Server, Sybase
Run Script with batch execution	TRUE/FALSE	DB2 LUW, DB2 z/OS, Oracle, SQL Server, Sybase
SET ... Options	The **Send Set Options** 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 **Enable SET query options** setting on the ISQL tab of the Options editor. For details, see ISQL options. If the **Send Set Options setting** is enabled, the remaining settings in this category let you specify the specific ANSI SQL options that are sent to the server: **ansi_defaults**, **ansi_nulls**, **ansi_null_dflt_on**, **QUOTED IDENTIFIER**, **cursor_close_on_commit**, **ANSI_PADDING**, **ANSI WARNINGS**, and **IMPLICIT_TRANSACTIONS**. The initial default values are hard-coded, not obtained from server settings.	SQL Server
SET Other Variables	Variables to be set at runtime.	MySQL
Set Quoted Identifier	TRUE/FALSE	Sybase
Show Plan	When set to TRUE, reports data retrieval methods chosen by the Microsoft SQL Server query optimizer.	SQL Server, Sybase
Sort Activity	TRUE/FALSE	Oracle
SQL Auto IS NULL	When set to TRUE, enables you to find the last inserted row for a table.	MySQL
SQL Big Selects	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
SQL Big Tables	TRUE/FALSE	MySQL
SQL Buffer Result	When set to TRUE, forces results from SELECT statements into temporary tables.	MySQL
SQL Log Bin	When set to TRUE, allows logging to the binary log.	MySQL
SQL Log Off	When set to TRUE, no logging is done to the general query log.	MySQL
SQL Log Update	When set to TRUE, allows logging to the binary log.	MySQL
SQL Low Priority Updates	When set to TRUE, gives table modifying operations lower priority than SELECT operations	MySQL

Option/Variable	Description	Platform
SQL Max Join Size	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
SQL Quote Show Create	When set to TRUE, table and column names will be quoted.	MySQL
SQL Safe Updates	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
SQL Select	The maximum number of records that should be returned from SELECT statements.	MySQL
SQL Warnings	Defines whether or not single row insert statements generate an information string in the event of a warning.	MySQL
Statistics I/O	Lets you display information regarding the amount of disk activity generated by Transact-SQL statements.	SQL Server, Sybase
Statistics Subduer Cache	Displays the number of cache hits, misses, and the number of rows in the subquery cache for each subquery.	Sybase
Statistics Time	Displays the number of milliseconds required to parse, compile, and execute each statement.	SQL Server, Sybase
Table Activity	True/False	Oracle
Table Count	Sets the number of tables that Sybase ASE considers at one time while optimizing a join.	Sybase
Text Size	8,192 is the default	SQL Server, Sybase
Transaction Isolation	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
Unique Checks	Performs uniqueness checks for secondary indexes of MyISAM tables.	MySQL

When query options are sent to the server

DBArtisan 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:

Session unlocked	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. Similarly if you unlock a currently locked session, all query options are sent to the server at each subsequent execution. For information on locking sessions, see Toolbar Options.

Session locked	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. If you open **Query Options** dialog and modify options, on the next execution only the modified settings are sent to the server. For more information, see Setting query options. These rules apply as long as the connection remains locked.

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.

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
	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 DBArtisan 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**.

   DBArtisan jumps the pointer to the next available token.

For more information, see Using the Query Plan Facility.

SQL preprocessing: #define and #include

DBArtisan 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 identiifiers to be replaced with **#define** processing.

Script content before SQL preprocessing

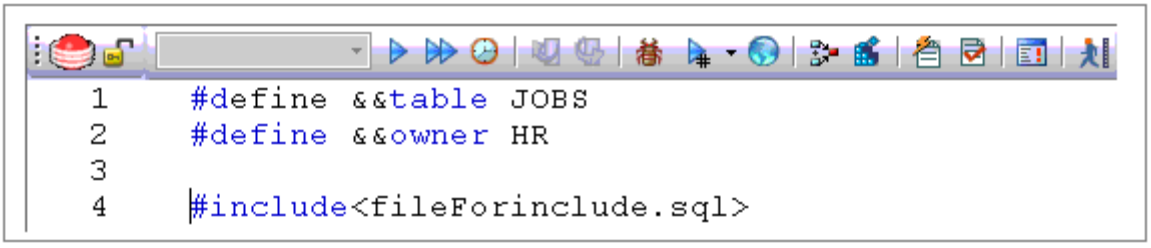

Contents of fileForinclude.sql

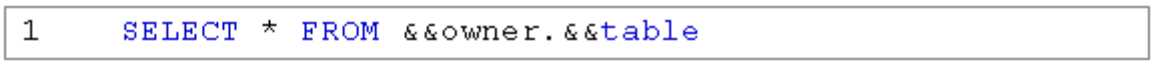

Script content produced by SQL preprocessing

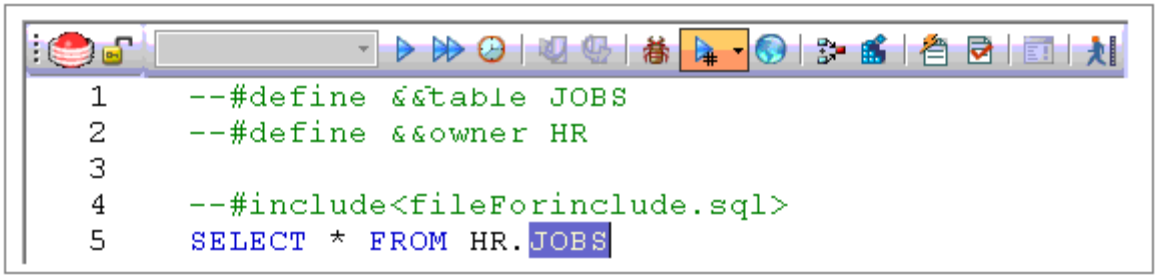

The key steps in working with SQL preprocessing are:

- **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 DBArtisan 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 offer 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 DBArtisan to preprocess #include directives

In ISQL editor processing of a #include directive, DBArtisan 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 Datasource Registration.
2. The location specified on the Directories tab of the Options editor. For details, see Directories options.

NOTE: For detailed information on setting DBArtisan options, see Specifying DBArtisan 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

DBArtisan 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 DBArtisan 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 DBArtisan 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 nor 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:

- DBArtisan 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.
- DBArtisan 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 SQL preprocessing: #define and #include for an introduction to SQL preprocessing in the ISQL editor

Automated error detection and coding aid features

DBArtisan analyzes code as you add content to an ISQL editor session. With real-time parsing of the SQL code, DBArtisan offers the following automated features:

- Syntax checking automatically flags syntactical errors in your scripts
- Code Assist features minimize errors and keystrokes by flagging semantic errors and offering suggestions with regard to object names

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

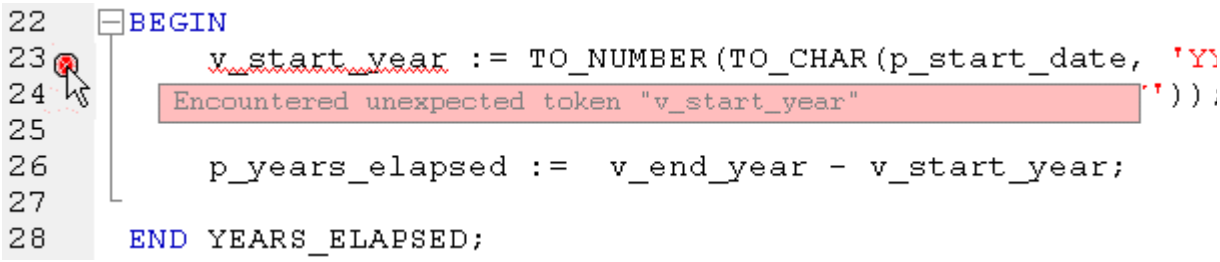

NOTE: A line-by-line listing of individual syntax errors is also available in the ISQL editor's Error pane. For more information, see About the Error pane.

DBArtisan can be configured to perform syntax checks automatically or to require a syntax check to be initiated manually:

- **Automatic** - If SQL Parsing is enabled, 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. For information on enabling and configuring SQL Parsing, see ISQL options - Code Assist tab.
- **Manual** - If SQL Parsing is disabled, a syntax check is only performed when you click the Syntax button on the ISQL editor toolbar:

  

  When SQL Parsing is disabled, syntax error annotations persist until the error is corrected and you explicitly run another syntax check.

Code Assist

Code Assist is a set of ISQL editor features that provide on-the-fly semantic error detection and coding aid with regard to object names:

- **Semantic validation** - detects object name references to objects not present in the datasource index. For more information, see About semantic validation.
- **Code Complete** - lets you insert or replace object names, selected from suggestion lists, as you edit a script. For more information, see Using Code Complete.

About 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 columns, tables, and views present in a script match those currently in DBArtisan's datasource index for the connected datasource. Minimizing errors associated with typographical errors or references to obsolete object definitions, DBArtisan analyzes the names of tables, columns, synonyms, and views as you type, and raises an error condition when it detects a name not present in the datasource index.

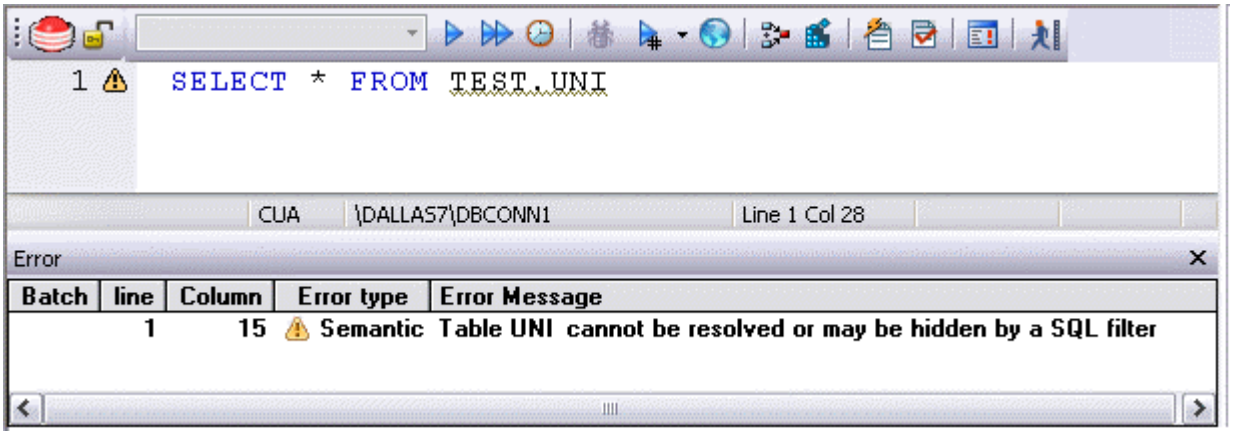

A semantic error can indicate one of the following conditions:

- The object name as specified in the script is not present on the database
- The referenced object is hidden by a SQL filter
- The datasource index is out of date and needs to be refreshed

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.
- To ensure accurate semantic validation, the datasource to which the ISQL editor is connected must be indexed and the index must be up to date. For more information, see ISQL options - Index Manager tab.
- Ensure that the ISQL Editor window is connected to a datasource. The datasource index, used in implementing semantic validation, requires a datasource name in order to resolve an object name.

  NOTE: 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.

Using 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.

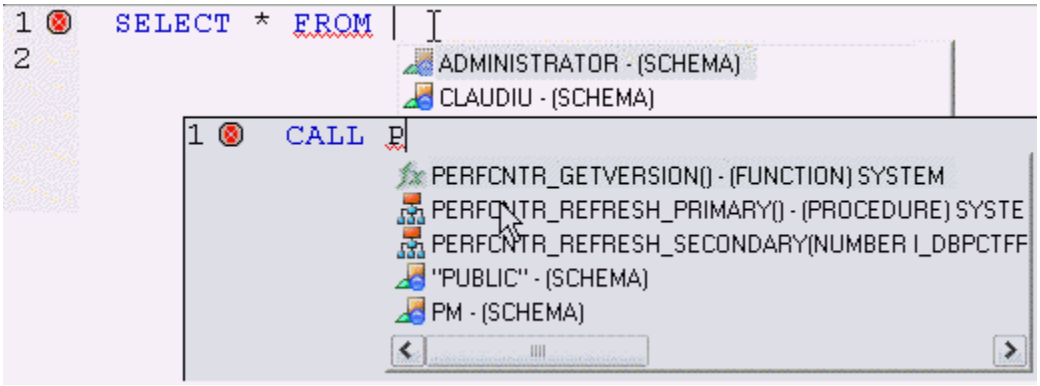

This feature becomes available at all points in INSERT, UPDATE, DELETE, SELECT, CALL, and EXEC statements, including all legal clauses, where a reference to the name of one of the following object types is valid:

- Columns
- Tables
- Views
- Functions
- Procedures
- Packages
- Synonyms

Code complete 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. For more information, see ISQL options - Code Assist tab.
- 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 configured for automatic invocation. For information on disabling auto-activation and other Code Complete preferences, see ISQL options - Code Assist tab. In addition, the datasource to which the ISQL editor is connected must be indexed and for accurate results, the index must be up to date. For more information, see ISQL options - Index Manager 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.

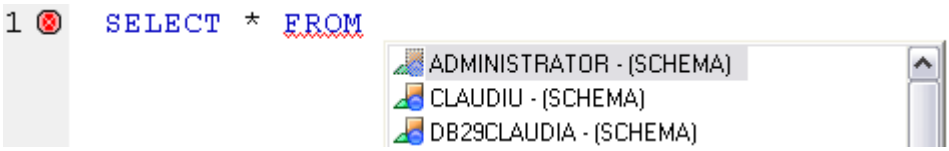

  Suggestions for the column name in a WHERE clause in the same statement would contain the schemas, databases and tables\views referenced in the FROM clause as well as a listing of the applicable columns based on the referenced tables\views.

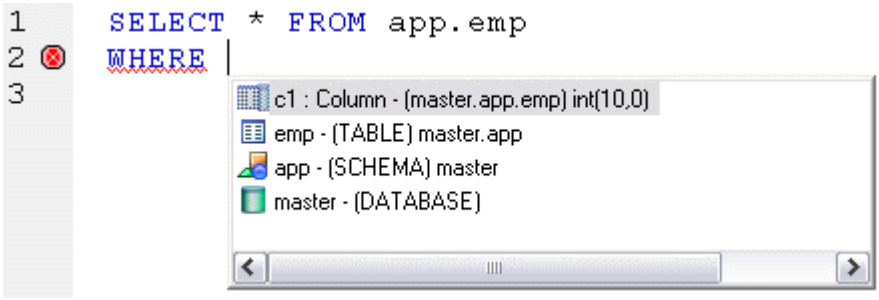

- **After the delimiter within an object name** - Code Complete offers a list that lets you select the next element to build a fully qualified object name.

- **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.

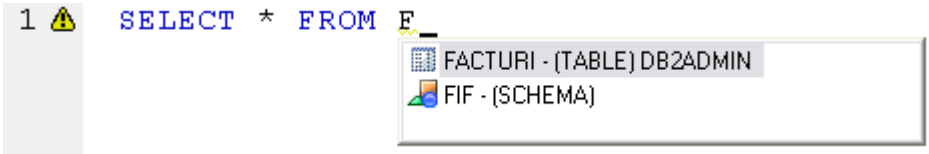

>

- **At any point within a partially-specified or complete object name** - Code Complete offers all selections appropriate to the location of the cursor.

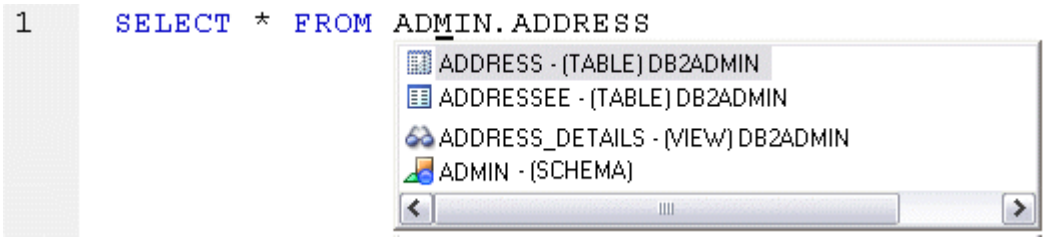

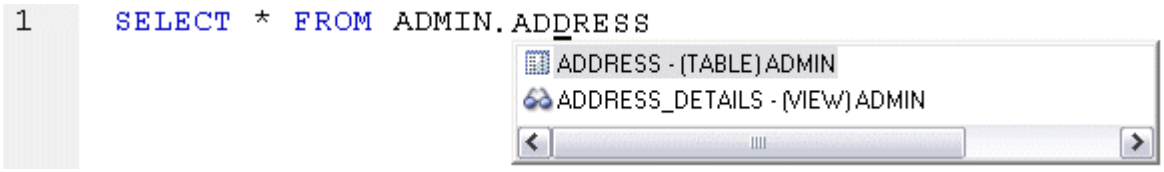

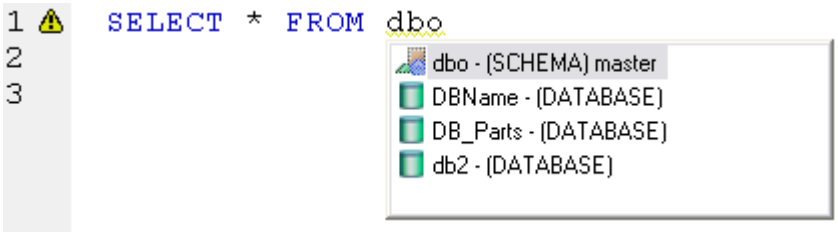

In addition, you can type additional characters while the suggestion list is available to further filter the items appearing in the list.

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.

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, DBArtisan 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.

Results window toolbar options

The table below describes the options of the Results window toolbar:

Option	Description
Lock	Lets you lock an ISQL Window to a specific datasource connection. When locked, the ISQL Window does not respond to datasource change events in other windows, and only uses the current connection for subsequent execution.
Options	Open the Options editor, letting you specify results window preferences. For more information, see Results (ISQL) options.
ReExecute	Lets you execute the script again without switching to the Query tab.
Close	Lets you close the current results window.

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.

To configure result set windows:

1. On the **File** menu, click **Options**.
2. In the **Options Editor**, click the list, and then click **Results** or click the **Results** tab.
3. Refer to Results (ISQL) options for details on how to set result windows options.

Configuring results window grid properties

You can alter the physical appearance in a reults set window. Effects include items such as 3-D buttons, highlighted headings, and enabled or disabled row and column gridlines.

To configure the appearance of result set grids:

1. On the **File** menu, click **Options**.
2. In the **Options Editor**, click the list, and then click **Grid Properties** or click the **Grid Properties** tab.
3. Refer to Grid properties (Results window) options for details on how to set result windows options.

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
- Pipe delimited ('|") files
- HTML
- Excel 2000 or higher
- XML
- User-specified delimited files

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**.

DBArtisan 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 Main Menu.

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**.

DBArtisan 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**.

DBArtisan sends the result set to the specified addressee(s).

For more information, see Results Editor.

Closing Result Window Tabs

DBArtisan lets you close tabbed Result set windows.

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.**

DBArtisan 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**.

DBArtisan 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 save data in Excel 2000 or later .xls, tab-delimited, comma-delimited, pipe-delimited, user-specified delimited .txt, HTML, and XML formats.

4 To include column titles, select **Include column titles when saving**.

5 If you chose User Specified Delimiter, in **User Specified Delimiter** type the delimiter.

6 Click **Save**.

DBArtisan 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**.

DBArtisan closes the Result Set.

2 If you have not saved your result set, DBArtisan 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

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**.

DBArtisan 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**.

DBArtisan 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**.

   DBArtisan 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**.

   DBArtisan 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**.

   DBArtisan 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**.

   DBArtisan 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**.

   DBArtisan 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**.

   DBArtisan 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**.

   DBArtisan 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**.

   DBArtisan 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**.

   DBArtisan 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. Double click the column header 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

- Changing the displayed view
- Format Border Styles
- Format Font Styles
- Format Color Styles

For information on general appearance of results, see Configuring results window grid properties.

Changing the displayed view

You can view results in a standard grid, in HTML format, or as flat ASCII. By default, when a result window is generated, the results are displayed according to the view option currently selected on the Results (ISQL) options. You can change the display in the results window.

To change the result view on the fly:

1 RIght-click in the results window.

2 On the shortcut menu, choose one of the **View as Grid**, **View as HTML**, or **View as ASCII** options.

Format Border Styles

You can use the shortcut menu to alter border properties.

Completing the Format Styles Dialog Box

To complete the Format Styles dialog box, do the following:

1 Right-click the **Result** data grid, and then click **Border**.

2 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.

3 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.

4 In the **Type** box, you can select the type of line you want to border the cell or cells by clicking the corresponding boxes.

5 To select a color, click the **Color** list, and then click the border color.

6 Click **OK**.

DBArtisan saves your changes and closes the Format Styles dialog box.

For more information, see Formatting Result Sets.

Format Font Styles

You can use 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 Right-click the **Result** data grid, and then click **Font**.

`DBArtisan opens` the Format Styles dialog box.

2 In the **Font** box, type or click the font you want to use.

3 In the **Style** box, type or click the font style you want to use.

4 In the **Size** box, type or click the size you want to use.

5 To make a line cross through the length of the text, in the **Effects** box, select the **Strikeout** check box.

6 To underline the text, in the **Effects** box, select the **Underline** check box.

7 To change the script style, click the **Script** list, and then click the new script style.

The Sample box displays the sample text of your selections

For more information, see Formatting Result Sets.

Format Color Styles

You can use the shortcut menu to alter color properties. Selecting the Color command displays the Cell Properties dialog box.

Completing theCell Properties Dialog Box

To complete the Format Styles dialog box, do the following:

1. Right-click the **Result** data grid, and then click **Color**.

   DBArtisan opens the Cell Properties dialog box.

2. Change the Text Color and Background Color options.

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 Uincode 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

MySQL Datasources

As easy as MySQL is to use, DBArtisan makes it even easier. The DBArtisan interface saves you having to write commands for many operations. For example, as you register MySQL datasources, DBArtisan makes it simple to perform multiple flush tasks simultaneously, set up or edit tables, and generate a variety of indexes.

What follows are descriptions of the tools you need to create and manage MySQL datasources. This section includes detailed instructions for creating and editing:

- Flush Options,
- Foreign Keys,
- Functions,
- Indexes/Primary Keys/Unique Keys,
- Users/Privileges,
- MySQL Databases,
- and Tables.

Additionally, from the Utilities menu you can extract and migrate schema, configure your servers, load or unload data, and backup or restore databases:

- Using the schema extraction wizard
- Using the schema migration wizard
- Using the Schema Publication wizard
- Server Configuration
- Database Backups
- Database Restore/Import

Flush Options

In addition to the basic datasource-level menu choices, MySQL gives you the option to flush some internal caches. The Flush options you select do not carry over from one datasource to another, and in fact must be reset every time you want to flush the caches.

A FLUSH command empties internal caches used by MySQL. Only a user with RELOAD privileges can execute FLUSH operations.

To Open the Flush Options Dialog Box:

1 On the **Explorer Tree**, right-click the **datasource** you want to flush, and then select **Flush** from the menu.

The MySQL Flush options are as follows:

Option	Description
DES_KEY_FILE	Reloads the key files for the DES_ENCRYPT and DES_DECRYPT functions
HOSTS	Empties the host cache tables. It's wise to flush the host tables if any hosts changed IP numbers or you get a "host is blocked" error message.

Option	Description
LOGS	Closes and then reopens all log files. If you have created an update log, a new file is created and the number of the file ending increases by 1 (file .07 > file .08).
MASTER	Deletes all binary logs in the index file, resets the binary log index file to empty, and then creates a new binary log file.
PRIVILEGES	Reloads privileges from the grant tables in the database.
QUERY CACHE	Defragments the query cache without removing any queries from the cache.
CLEAR CACHE	This is enabled only when you select clear cache. This is what you should choose if you want to remove queries from the cache.
SLAVE	Makes the slave forget its position in the master relay logs and starts a new relay log.
STATUS	Resets status variable to 0.
TABLES	Closes all open tables and flushes all query results from the query cache.
USER RESOURCES	Resets all hourly-bases user resources to 0, like MAX_CONNECTIONS_PER_HOUR, etc.

If you want to carry out a flush operation not included among the DBArtisan options, you can write your own flush commands using SQL Scripting.

Foreign Keys

MySQL recognizes two types of tables, MyISAM and InnoDB. MyISAM tables access data records using an index, where InnoDB tables allow transactions and foreign keys. Therefore, the discussion of foreign keys is limited to your InnoDB tables.

After you create a foreign key using the Create Foreign Key Wizard, further actions are available to you from the toolbar or by right-clicking an individual foreign key.

- Open - Launches Foreign Key editor for selected index(s).
- Drop - Drops selected Foreign Key (s).
- Report - Launches standard report generation dialog.
- Extract - Reverse engineers selected Foreign Key (s) into DDL and loads into working DDL editor.
- Migrate
- Copy Object Name(s)
- Find Object
- Select All
- UnSelect All

Creating or Editing a Foreign Key

Embarcadero DBArtisan's Create 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. By using the Create Foreign Key Wizard you obviate the need for remembering the code underlying the creation process.

To Open the Create Foreign Key Wizard

1 On the **Explorer**, find the database where you want to add the new foreign key.

2 On the **Foreign Keys** branch, right-click and select **New**.

OR

Click the **Foreign Keys** branch and then select **New** from toolbar.

To Open the Foreign Key Editor

1 On the **Explorer**, find the database where you want to edit the foreign key.

2 Click the **Foreign Keys** branch and select the specific foreign key in the right pane.

3 Right-click the **Foreign Key**, and select **Open**.

OR

Choose **Open** from the toolbar.

For the changes you make to take effect, click **Alter** on the toolbar.

NOTE: Once the Foreign Key Editor is open, you can navigate to any foreign key in the same database.

The tables below displays all the fields you'll find creating or editing a foreign key:

Panel/Tab	Option	Description
Foreign Key Properties	*Constraint Name*: System Generated or User Specified	The constraint name must be unique. You can either rely on MySQL to assign a unique name or you can specify one.
	Constraint State: On Delete and On Update	The default constraint state is NO ACTION. This ensures that referenced values cannot be updated or deleted if to do so would violate referential integrity. CASCADE 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. RESTRICT prevents a command from executing if changes to referential data prompts an error.
Column Mapping	Because MySQL requires matching indexes on the child and parent tables involved in a foreign key relationship, the foreign key columns you select must be found within an already created index in the main table. They must also be selected in the same order.	
	Main Table (child table)	Choose the table to which you want to add the foreign key from the drop-down list. Then select the relevant column(s).
	Referenced Table (parent table)	Select the database and table from the drop down lists that you want this foreign key to refer to. You must also select the appropriate columns.
	Enforce constraints	When checked, it ensures that the constraints you have selected will be enacted when values change in the parent table.
	New Index	The column(s) on the child table must be a part of the same index as the column(s) on the parent table. If the index(es) need to be created, you can open the Create Index Wizard, create an index, and be returned to the Foreign Key Wizard.
DDL Vew	You can see the SQL generated by the selections you made. Click **Execute** to create the new foreign key. Note that this is a read-only panel.	

Functions

As for the other platforms, users can create subroutines that can be reused so you can control, access, and manipulate the data that underlies an object. DBArtisan will construct the CREATE FUNCTION statement based on the information you supply when you generate a new function. Of course you are also able to edit any function after it has been created. It is important to keep in mind that to create a user-defined function, functions must be written in C or C++, your operating system must support dynamic loading, and you must have compiled mysqld dynamically (not statically).

- New - Launches UDF creation wizard.
- Open - Launches UDF editor for selected index(s).
- Drop - Drops selected UDF (s).
- Report - TBD.
- Extract - Reverse engineers selected UDF (s) into DDL and loads into working DDL editor.
- Migrate - Disabled.

Creating or Editing a Function

Embarcadero DBArtisan's Create Function Wizard lets you write a function you without knowing the underlying commands. The wizard gives you the opportunity to bring a function that's stored in an external library into MySQL.

To Open the Create Function Wizard

1. On the **Explorer**, find the **MySQL database** on the server where you want to add the new function.
2. On the **Functions** branch of the **MySQL database**, right-click and select **New**.

   OR

   Click the **Functions** branch of the **MySQL database** and then select **New** from the toolbar.

   NOTE: No functions are installed by default.

To Open the Functions Editor

1. On the **Explorer**, find the **MySQL database** where you want to edit the function.
2. Click the **Functions** branch and select the specific function in the right pane.
3. Right-click the **function**, and select **Open**.

   OR

   Choose **Open** from the toolbar.

   For the changes you make to take effect, click **Alter** on the toolbar.

The table that follows lists all the fields you'll find as you create or edit a function:

Function Properties

Required Information	Description
Function Name:	In 64 characters or less, you can name your function.
Return Value:	STRING, REAL (also known as Double), INTEGER
Shared Object Library:	Identify the file that holds the library where your functions are stored.
Aggregate	Check the box if the function you are creating will collapse a large amount of data into a single output, or aggregate the data.

DDL View

Required Information	Description
DDL View	You can see the SQL generated by the selections you made. Click **Execute** to create the new foreign key. Note that this is a read-only panel.

Indexes/Primary Keys/Unique Keys

Indexes for tables operate in much the same way as an index for a book works. Indexes help you find what you are looking for without altering the data itself. Specialized indexes include primary keys and unique keys.

A **Primary Key** uniquely identifies a row in a table and is used to relate a table to foreign keys in other tables.

A **Unique Key** ensures no duplicate values are entered in table columns where you place a unique constraint.

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. Indexes can improve read but can slow write performance because more information in the system catalog has to be updated. Consequently, extra indexes can actually slow overall performance if data modification occurs frequently on the table.

DBArtisan separates system indexes from user-defined indexes in the Datasource Explorer to ensure that system indexes are not accidentally altered or dropped.

- New - Launches Index creation wizard.
- Open - Launches Index editor for selected index(s).
- Drop - Drops selected index(s).
- Report - Launches standard report generation dialog.
- Extract - Reverse engineers selected table(s) into DDL and loads into working DDL editor.
- Migrate
- Copy Object Name(s)
- Find Object
- Select All
- UnSelect All

Creating or Editing an Index/Primary Key/Unique Key

Whether you are creating an index, a primary key, or a unique key, the same creation wizard opens with defaults in place appropriate to what you want to generate.

To Open the Create Index Wizard:

1. On the **Explorer**, find the database where you want to add the new index/primary key/unique key.
2. On the **Indexes/Primary Key/ Unique Key** branch, right-click and select **New**.

   OR

   Click the **Indexes/Primary Key/ Unique Key** branch and then select **New** from toolbar.

The Create Index Wizard can also be opened from within the Create Foreign Key Wizard.

To Open the Index/Primary Key/Unique Key Editor

1. On the **Explorer**, find the **MySQL database** where you want to edit the index/primary key/unique key.
2. Click the **Index/Primary Key/Unique Key** branch and select the specific index in the right pane.
3. Right-click the **Index/Primary Key/Unique Key**, and select **Open**.

   OR

   Choose **Open** from the toolbar.

   For the changes you make to take effect, click **Alter** on the toolbar.

The table that follows describes all the fields you'll encounter as you create or edit an Index/Primary Key/Unique Key:

Tables and Columns

Required Information	Description
Table Name	Note that the field is populated with a drop-down list of all the tables in the database. The tables appear in alphabetical order, and the field is filled with the first table from the list. Choose the table you want to index.
Index Name	The name you type can be up to 64 alphanumeric characters.
Constraint Type	**Primary**: Each record of a table is identified as uniquely and relates to a foreign key in another table. **Unique:** Ensures that there is no duplication between values in columns where you place this constraint. **Full Text:** Enables the search for several words in arbitrary order in a table. **Spatial:** Allows you to find records that are defined by location, or geometry types.
Index Storage Type	**Hash**: Used for equality comparisons. Only whole keys can be used to search for a row. **BTree**: Tree data structure that keeps data sorted. BTrees grow from the bottom up as elements are inserted. **RTree:** Tree data structure used for spatial indexes and access to multidimensional information.
Specify Columns in Index	Columns are listed by name and datatype and whether or not they are nullable. As you check selections, the sort order is identified.

DDL View

Required Information	Description
DDL View	You can see the SQL generated by the selections you made. Click Execute to create the new foreign key. Note that this is a read-only panel.

Users/Privileges

Adding users to a MySQL database is made simple when you use DBArtisan. Unlike adding new users to other platforms, MySQL doesn't rely solely on a userid/password combination. The MySQL server must not only validate the user and password, but also the host from which the connection is being requested before the new user is admitted to the server's inner sanctum. The Create User Wizard and the User Editor enable these validations.

After you create a foreign key using the Create User Wizard, further actions are available to you from the toolbar or by right-clicking an individual user or from the toolbar:

- Open
- Alter
- New -- Launches the Create New User Wizard
- Drop
- Extract
- Migrate
- Report - N/A
- Command

- Create Like
- Change Password
- Copy Object Name(s)
- Refresh
- Close

Creating or Editing a User

The Create User Wizard lets you give a new user the privileges he or she needs to function on the MySQL database of choice. Here is where you enter not only the User ID and password, but also validate the host from which that user will be connecting to the server. Users can be limited to access from a single host or be given access from multiple hosts.

To Open the Create User Wizard

1. On the **Explorer**, find the database where you want to add the new user and expand the **Security** node.
2. On the **Users** branch, right-click and select **New**.

   OR

   Click the **Users** branch and then select **New** from the toolbar.

To Drop a User

1. On the **Explorer**, find the **database** where you want to drop the user.
2. Click the **Users** branch and select the specific user in the right pane.
3. Right-click the **user**, and select **Drop.**
4. Confirm the action.

To Open the User Editor

1. On the **Explorer**, find the **database** where you want to edit the user.
2. Click the **Users** branch and select the specific user in the right pane.
3. Right-click the **user**, and select **Open**.

   OR

   Choose **Open** from the toolbar.

   For the changes you make to take effect, click **Alter** on the toolbar.

   NOTE: Once the User editor is open, you can navigate to any user in the same datasource.

The table below displays all the fields you'll find when you create or edit a user:

User Information

Required Information	Description
User Name	The name of the new user. At this time, the name cannot exceed 16 characters, although the field accepts 34. DBArtisan can't enforce the discrepancy, so we urge you to show a little restraint. This is the name that is displayed in the **User** branch of the **Explorer** tree. If you leave this field blank and complete the process, you create an <anonymous> user. Only one anonymous user per datasource is allowed.
Full Name	**OPTIONAL**. At this time, the name cannot exceed 16 characters, although the field accepts 34. If you entered a nickname or an alias in the user name field, you can display the user's true identity here.
Description	**OPTIONAL**. Write a brief description of the user if you want.
Email	**OPTIONAL**. The new user's email address.
Contact Information	**OPTIONAL**. The new user's title, address, phone number or whatever contact information you want to include.
User Icon	**OPTIONAL**. If you want to identify your new user by color or pattern, you can use the icon editor to make a distinguishing mark. You can assign groups of users the same icon or assign each user his or her own icon. This tool helps you distinguish users at a glance. **Note**: If you don't want to use the head that's preloaded, simply click the red X and start designing your own with the other tools.

User Hosts

Required Information	Description
New Host	Click the dotted square to open a text box where you can enter the ip address of the computer on which the MySQL server is running (e.g., 10.10.12.114). It's possible to enter more than one address to give the user access to more than one server, but you need to enter them one at a time. % is the symbol you can use to mask user hosts.
Password/Confirm	Enter the password and confirmation the user needs to use to gain access to the MySQL datasource. If you want to apply a password to every datasource user, click Apply To All.

System Privileges

Required Information	Description
Database	Each host you identified for the given user is listed here. You can set the same or different privileges for each host for a given user, but you must set them host by host.
Privileges	System privileges are Reload, Shutdown, Process, File, Show DBs, Execute, Replication Client, Replication Slave, Super (User). To grant system privileges, click the green check mark and the Grant Privileges Dialog Box opens. To revoke system privileges, click the red X and the Revoke Privileges dialog box opens.

Object Privileges

Required Information	Description
Database/Table Name	Each host you identified for the given user is listed here. You can set the same or different privileges for each host for a given user, but you must set them host by host.

Required Information	Description
Privileges	Object privileges include Alter, Create, Create Temporary Tables, Delete, Drop, Index, Insert, Lock Tables, References, Select, Update, with Grant. To grant object privileges, click the green check mark and the Grant/Revoke Privileges Dialog Box opens where you can grant privileges. To revoke object privileges, click the red X and the Grant/Revoke Privileges Dialog Box dialog box opens where you can revoke privileges.

Grant/Revoke Privileges Dialog Box

Required Information	Description
Check Privileges you wish to Grant/Revoke to the grantee and the grantee's host	Check/uncheck the privileges. **NOTE**: The With Grant option, when selected, allows the user you are creating to have the same grant privileges you have. User-only privileges are represented by a single check mark in the system/object privileges grid. The server checks each request the user issues against the granted privileges. User/Grant privileges are represented by a double-check mark in the system/object privileges grid.

DDL View

Required Information	Description
DDL View	The DDL window displays the command code you created by completing the prior panels of the wizard. You have the ability to save the code to disk, to print it, to email it, to open an ISQL window where you can tinker with the code, or to schedule the execution of the task.
Launch object editor after execute	When this option is checked, the script will execute immediately and the User Editor will open so you can grant specific privileges.

MySQL Databases

Each database is a collection of tables. The goal of a database is to manage large amounts of data in a multi-user environment as well as protect data against unauthorized access.

The Create Database Wizard lets you add a new MySQL database using the Embarcadero DBArtisan interface so you don't have to write all the code yourself in a scripting window or from a command prompt.

After you have created a new database, the following additional tasks are available:

- New – Launches the aforementioned Database creation wizard.
- Alter – Enabled when a property of the current database has been modified.
- Drop
- Extract – Reverse engineers DDL for current database and loads into working DDL editor.
- Report – TBD.
- Command Menu

- Refresh – standard repopulation feature.
- Close – Close editor with save validation.

Creating or Editing MySQL Databases

Databases are the shell into which you will pour the tables that contain the data records you want to manage. The Create Database Wizard makes it easy for you to add databases to your registered MySQL servers.

To open the Create Database Wizard

1. On the **Explorer**, find the **Datasource** where you want to add the new database and expand it.
2. On the **Databases** node, right-click and select **New**.

   OR

   Click the **Databases** node and then select **New** from toolbar.

To Open the Database Editor

1. On the **Explorer**, find the **Datasource** where you want to edit the database.
2. Click the **Databases** branch and select the specific database.
3. Choose **Open** from the toolbar.

   For the changes you make to take effect, click **Alter** on the toolbar.

   NOTE: Once the Database editor is open, you can navigate to any database in the same datasource.

The table below displays all the fields you'll find as you create or edit a database:

Database Properties

Database Privileges

DDL View

Database Properties

Required Information	Description
Database Name:	The database name you choose can be up to 64 characters long. Feel free to use all alphabetic characters or mix in some special or numeric characters. But a name cannot be exclusively special or numeric characters.
Default Character Set: (optional)	This is the character set for a language or alphabet. If you do not make a selection from the drop-down list (click the arrow to see the choices), the database will use the server's default character set.
Collation:	The collation encodes the rules governing character used for a language (like Greek) or an alphabet. The database will use the server's default collation unless you specify otherwise. To see the collation choices available to you, click the arrow and make a selection.
NOTE: It's possible to create databases with different character sets and collations on the same MySQL server.	

Database Privileges

Required Information	Description
Scope (in the creation process, database is the only option available)	**Database**: Select user privileges for the specific database you are creating or editing. Privileges are available for all datasource users. **Global:** Displays system-wide privileges for datasource users. To change global privileges you need to edit the user (see Creating or Editing a User). **All**: Displays both system-wide and database privileges.
Grantee	A list of all server users and their hosts is displayed. You can pick and choose to whom you want to grant database privileges and whether each user will have user only or user with grant privileges.
Privileges	Privileges include Alter, Create, Create Temporary Tables, Delete, Drop, Index, Insert, Lock Tables, References, Select, Update, with Grant. To grant privileges, click the green check mark and the Grant/Revoke Privileges Dialog Box opens where you can grant privileges. To revoke object privileges, click the red X and the Grant/Revoke Privileges Dialog Box opens where you can revoke privileges.

DDL View

Required Information	Description
DDL View	The DDL window displays the command code you created by completing the prior panels of the wizard. You have the ability to save the code to disk, to print it, to email it, to open an ISQL window where you can tinker with the code, or to schedule the execution of the code.
Launch object editor after execute	When this option is checked, the script will execute immediately.

Tables

To state the obvious, tables are the building blocks of any database and are the basic unit of data storage. A table's properties are determined by its columns and rows, or data records. MySQL servers can store tables in multiple formats, including MyISAM and InnoDB. MyISAM tables (ISAM is the acronym for indexed sequential access method) are used most often for read operations. The read operation is very fast, but you cannot include any referential integrity, such as a foreign key. Also, MyISAM tables only issue table-level locks. InnoDB tables, on the other hand, do permit transactions and foreign key constraints. InnoDB tables also lock data at the row level, which is appropriate for high transaction tables. Additional table types available are MERGE, MEMORY, FEDERATED, and ARCHIVE among others. For a complete discussion of table types, go to the MySQL documentation of table and engine types: http://dev.mysql.com/doc/mysql/en/storage-engines.html

- New - Launches table creation wizard.
- Open - Launches Table Editor for selected table(s).
- Drop - Drops selected table(s).
- Report - Launches standard report generation dialog.
- Extract - Reverse engineers selected table(s) into DDL and loads into working DDL editor.

- Migrate
- Rename
- Analyze tables
- Create like
- Check Table
- Rebuild table
- Repair table
- Optimize table
- Disable Keys
- Enable keys
- Truncate table
- Flush table
- Convert Table
- Schema
- Select * From
- Copy Object Name
- Find Object
- Select All
- Unselect All

Creating or Editing Tables

DBArtisan lets you create many different table types using the Create Table Wizard. After you have successfully added a table to the database, you can reopen and edit it using the Table Editor.

To Open the Create Table Wizard

1 On the **Explorer**, find the database where you want to add the new table and expand the **Databases** node.

2 On the **Tables** branch, right-click and select **New**.

OR

Click the **Tables** branch and then select **New** from toolbar.

To Open the Table Editor

1 On the **Explorer**, find the **database** where you want to edit the table.

2 Click the **Tables** branch and select the specific table in the right pane.

>

3 Right-click the **table**, and select **Open**.

OR

Choose **Open** from the toolbar.

For the changes you make to take effect, click **Alter** on the toolbar.

NOTE: Once the Table editor is open, you can navigate to any table in the same datasource.

In the following you'll find all the fields you'll encounter as you create or edit a table:

>

Properties

Required Information	Description
Table Name	Enter a meaningful and unique name that is no more than 64 characters. **NOTE**: You must name the table to enable the Create Table Wizard.
Storage Type	From the drop-down list, select how you want the data to be stored. Storage types are described in MySQL documentation. http://dev.mysql.com/doc/mysql/en/storage-engines.html
Row Format	DEFAULT -- returns the default value if there is one FIXED -- each row is stored with a fixed number of bytes. DYNAMIC -- data records have variable length COMPRESSED -- each record is compressed separately.
Create via SELECT	When you create a table using a SELECT command, the individual columns take their data types from the SELECT command, and don't have to be declared explicitly. Attributes, however, are not carried over.
Unique Key Violations	When you try to insert or update a row that causes a violation, you can choose to IGNORE the violation and continue processing the next row. Or, you can choose REPLACE, which is equivalent to an insert statement if a violation occurs. The row will be replaced with the new data.
Default Character Set	This is the character set for a language or alphabet.
Default Collation	The collation encodes the rules governing character used for a language (like Greek) or an alphabet.
Auto-increment	You can specify the initial value used for the auto-increment sequence. This is possible only for MyISAM tables and for InnoDB tables built in MySQL versions 5.0.3 or greater. For InnoDB tables built in versions prior to 5.0.3, you can manually insert a dummy row with a value one less than the desired value after creating the dummy table; you then delete the dummy row.
Comment	
Min. Rows/Max Rows	Note that if you set the value for either parameter at 0, MySQL removes the setting. MySQL will take care of the row settings in this case.
Average Row Length	
Pack Keys	0 1 Default
Check Sum	True False
Delay Key Write	True False

Columns

Required Information	Description
Column Name(s)	Click the little square box above the left panel to open a text field where you name the column. You can create multiple columns by returning to the little square box as many times as you need.

Required Information	Description
Type	MySQL documentation explains the column types available to you. http://dev.mysql.com/doc/mysql/en/column-types.html
Enum/Set Values	
Unsigned	True/False
Zerofill	True/False
Binary	True/False
Width	
Scale	
Auto-increment	True/False--True allows an integer column to behave like a sequence based column. If you insert NULL (recommended) or 0 in the column, the column value is automatically to the next sequence value.
Allow Nulls	True/False
Default Character	
Default Collation	
Comment	
Use Default Value	

Indexes

Required Information	Description
Index	In the upper panel, specify the name of any index you want to use in the table you are creating. Click the little square box to enter the name. You can enter multiple indexes.

Privileges

Required Information	Description
Grantee	Select column by column which users you want to bestow with privileges for the target table. The Grant/Revoke Privileges Dialog Box opens.

Merge Tables

Required Information	Description

Space

Required Information	Description
This is available only as you edit a table.	Table usage and its distribution of table space is graphically displayed.

DDL View

Required Information	Description
DDL View	The DDL window displays the command code you created by completing the prior panels of the wizard. You have the ability to save the code to disk, to print it, to email it, to open an ISQL window where you can tinker with the code, or to schedule the execution of the code.
Launch object editor after execute	When this option is checked, the script will execute immediately.

Schema Publication

DBArtisan features a Schema Publication Wizard that can generate HTML reports documenting the details of your databases, which you can share over an intranet or the Internet. DBArtisan uses HTML Templates to create the reports that you can manually edit to create customized reports.

The Schema Publication Wizard lets you:

- Generate reports for all or any level of your database.
- View, find, and print the HTML reports directly in DBArtisan with the built-in HTML browser.

- Create a new schema publication operation.
- Reuse a previously defined schema publication operation.
- Select a datasource for which you want to generate an HTML report.
- Select the server object type and database object type for which you want to publish a schema report and lets you customize the report.
- Select specific objects and to set certain script options for each object type.
- Specify the report directory, HTML file name, and report title.
- Save the schema publication operation for re-use.
- Preview the definition of your schema publication operation.

Publishing a Schema

Generating a report for any level of your MySQL database is made easy by DBArtisan.

To Open the Schema Publication Wizard:

1 On the **Datasource Explorer**, select the MySQL server.

2 On the **Utilities** menu, click **Schema Publication**.

The sections below describe the fields you may encounter as you create the publication operation:

Playback or New Operation

Required Information	Description
New Publication Operation	This initiates the process of creating a new publication operation.
Previous Publication Operation	If you want to replay a publication operation, when you select the radio button, the browse function is enabled so you can find the publication file you want to re-use.

Server Connection

Here you select the MySQL server and database(s) that you want to include in this particular publication instance.

Object Type Selection

Specify the server object(s) you want to include in the schema report. You can also select database objects organized by owner. To create a customized publication script based on individual objects, click the customize box.

Object Selection

Select the objects you had in mind when you opted to customize the object list.

Publication Option

Give the report a name, and decide if you want to create separate html files for each object encompassed by the report.

Summary

A summary of the publication operation as you created it in the preceding steps is displayed. If you want to make chances, you need to go **Back**. You can also designate the file where you want the output script to live and decide whether or not you want to save this operation for use again. Files are saved with the .sav extension.

Execute

Watch the progress of the publication as it is generated and see any errors if they occur. Then click Continue or Finish to see the results.

Server Configuration

DBArtisan places Configuration under the Server node of the Datasource Explorer. When you select Configuration, DBArtisan shows configuration information about your datasource in the right pane of the Explorer window. You can use the Datasource Explorer or shortcut menu to open the System Configuration information panel and make changes to the server's configuration.

You can:

- Open
- Extract
- Report
- or Refresh

each variable.

Configuring a Server Variable

To Open the System Configuration information panel:

1. On the **Explorer**, find the **server** where you want to make changes to the configuration.
2. Expand the **Server** node and double-click the **Configuration** branch.
3. Double-click any variable in the right-pane.

   OR

   Right-click in the variable pane and select **Open**.

The System Configuration information panel displays the name, value, and dynamic status of each system variable. Make changes to the value column and click **Alter** on the toolbar. When you do this, the Preview SQL box opens.

Database Backups

Implementing a plan to back up critical databases is perhaps the most important job for any database administrator. Without a sound disaster recovery plan, your company has no safeguards against hardware and software failures.

The Database Backup Wizard lets you backup your databases without knowing the underlying syntax. It provides support for backing up databases as well as transaction logs, including commands to stripe the backup across multiple devices.

The Database Backup Wizard lets you:

- Create a new dump operation.
- Re-use a previously defined operation.
- Specify the target server, database, and type of backup operation.
- Specify backup device type.
- Specify how to report the progress of the backup operation.

You can back up the server while the database is online and in use. However, the backup will fail if you are performing one of the following operations while attempting the back up:

- Creating indexes.
- Performing non-logged operations.
- Shrinking the database.
- Creating or deleting database files.

Creating or Editing a Backup Operation

DBArtisan lets you create different backup operations using the Database Backup Wizard. After you have successfully created a backup operation, you can reopen and edit it.

To Open the Database Backup Wizard

1. On the **Datasource Explorer,** select the **Databases** node.

   DBArtisan displays the databases in the right pane of the Explorer window.

2. Select the target database.
3. On the **Utilities** menu, click **Database Backup**.

In the following you'll find all the fields you'll encounter as you create or edit a backup operation:

Playback or New Operation

Backup Selection

Select Databases

Select Tables

Options Editor for the Wizard

Output File Type

Normal Output File Options

Datafile Options

Summary

>

Playback or New Operation

Required Information	Description
New Dump Operation	This initiates the process of creating a new backup operation.
Previous Dump Operation	If you want to replay a backup operation, when you select the radio button, the browse function is enabled so you can find the .DBDump file you want to re-use.

Backup Selection

Required Information	Description
Backup Databases	This is the default operation.
Backup tables within a single database	

Select Databases

From the checklist, choose one or more databases you want to include in the new dump operation.

Select Tables

From the drop-down list, choose the database you want to back-up tables from and then check one or more tables to include in the new dump operation.

Options Editor for the Wizard

NOTE: By default, the options on this wizard panel are sorted by category. You can also sort them alphabetically if that makes life easier.

Required Information	Description
TABLES: No Data	True/False: Backup DDL only, without saving any data rows.
Continue if Errors Occur	True/False: Force the dump to continue even in the event there is a SQL error when dumping a table.
Flush Server Logs	True/False: Flush MySQL logs before the dump begins.
Drop Table Statements	True/False: DROP TABLE statements will be added.
Lock/Unlock Statements	True/False: Add LOCK/UNLOCK table statements around each table dump.
Lock All Tables	True/False: Lock all tables prior to beginning the dump.
Allow Keyword Column Names	True/False: Permits creation of column names that are keywords.
Complete Insert Statements	True/False: Complete insert statements will include column names.
Use MySQL Table Options	True/False: Include all MySQL-specific table options in the CREATE TABLE statements.

Required Information	Description
Use Insert Delayed Statements	True/False: Insert rows using INSERT DELAYED statements.
Disable Keys	True/False **NOTE:** The disable keys command is only valid for MyISAM tables.
Use Extended Inserts	True/False: Enables extended INSERT syntax that includes several VALUES lists.
Use Create Table Statements	True/False: Script CREATE TABLE statements for each table.
GENERAL OPTIONS: Suppress Information	True/False: Suppress additional information like program version, host, and so forth.
Exclude Create Database Statements	True/False: Will not write CREATE DATABASE statements.
Retrieve Single Rows	True/False: Instead of retrieving the row set, retrieve one row at a time.
Use Quotes	True/False: Use quotation marks around database, table, and column names.
Use Begin Statement	True/False: Before dumping data from the server, issue a BEGIN statement.
REPLICATION: Use First-Slave Option	True/False: Lock all tables across all databases.
Use Master-Data Option	True/False: Lock all tables and generate CHANGE MASTER TO statements.
Delete Binary Logs	True/False: Delete binary logs after the dump operation on the master replication server.
CHARACTER SET: Suppress Set Names Statement	True/False: Suppress SET NAMES default_character_set statement.
Default Character Set	Identify the character set that you want to use as the default.
NETWORK: Protocol	Choose TCP, Socket, Pipe, or Memory as the connection protocol.
Max Allowed Packet	Enter the maximum size of the client/server communication buffer.
Net Buffer Length	Enter the maximum size of the client/server communication buffer.
Compress	True/False: Compresses all information sent between the client and the server if both support compression.
COMPATIBILITY: ANSI	True/False: Produces ANSI-compatible content.
MySQL 3.23	True/False: Produce MySQL 3.23-compatible output.
MySQL 4.0	True/False: Produce MySQL 4.0-compatible output.
PostgreSQL	True/False: Produce PostgreSQL-compatible output.
Oracle	True/False: Produce Oracle-compatible output.
MSSQL	True/False: Produce MSSQL-compatible output.

Required Information	Description
DB2	True/False: Produce DB2-compatible output.
MaxDB	True/False: Produce MaxDB-compatible output.
No key options	True/False: Produce output without key options.
No table options	True/False: Produce output without table options.
No field options	True/False: Produce output without field options.

Output File Type

Required Information	Description
Create a normal output file	Lets you create an output file and save it to the location of your choice in the next panel. It is important that you specify the full path the to the file, including the file name, for example C:\Documents\Dumps\Dump1.dbdump
Create a tab-separated data file. **NOTE:** This option should only be used when mysqldump is run on the same machine as the mysql server.	For each dumped (backed up) table, mysqldum creates a 'tbl_name.sql' file that contains the CREATE TABLE statement that creates the table and a 'tbl_name.txt' file that contains its data.

Normal Output File Options

Required Information	Description
Results File	Type the path and file name of the results file you are creating. It is important that you specify the full path the to the file, including the file name, for example C:\Documents\Dumps\Dump1.dbdump
Prevent newline "\n' characters from being converted to "\r\n' carriage return/newline sequences (recommended for Windows systems)	This option, intended for Windows, allow the output to be written to the named file. If not, linefeeds (newlines) will be converted to carriage return/linefeed pairs.
Write dump output as well-formed XML	Stores the file output as XML.

Datafile Options

Required Information	Description
Tab-separated data file: Enter the path and filename of the results file	Type the path and file name where the tab file will be generated.
Fill in a terminal value or leave a blank if you want to retain the default: Fields terminated by, Fields enclosed by, Fields optionally enclosed by, Fields escaped by, Lines terminated by	

Summary

Here you find a summary of all the choices you have made in creating this operation. If you want to make changes, you can make them in this grid or go **Back**. You also have the option of saving the operation. Saved files are identified with a .DBDump extension.

Database Restore/Import

To both validate your back ups and practice the disaster recovery process, you should regularly restore your database dumps. If possible, every day, you should restore your databases on a separate server to enhance system availability. In this manner, you can maintain a warm backup of your database guarding against downtime from system failure.

The Database Restore/Import Wizard lets you:

- Create a new load operation.
- Reuse a previously defined operation.
- Specify the target server, database, and type of restore operation.
- Specify the recovery completion state.
- Specify how to report the progress of the load operation.
- Specify which file to load from the device.
- Specify whether to perform safety checks before replacing the database.

Creating or Editing a Restore/Import Operation

DBArtisan lets you create different restore/import operations using the Database Import Wizard. After you have successfully created an import operation, you can reopen and edit it.

To Open the Database Import Wizard

1 On the **Datasource Explorer,** select the **Databases** node.

DBArtisan displays the databases in the right pane of the Explorer window.

2 Select the target database.

3 On the **Utilities** menu, click **Database Restore**.

In the following you'll find all the fields you'll encounter as you create or edit a restore/import operation:

Playback or New Operation

Required Information	Description
New Import Operation	This initiates the process of creating a new import operation.
Previous Import Operation	If you want to replay an import operation, when you select the radio button, the browse function is enabled so you can find the .DBDump file you want to reuse.

Target Database

Select the target database from the drop-down list.

Select Data Files

Type or browse to the file(s) that has the data you want to import.

Database Import Wizard Options

NOTE: By default, the options on this wizard panel are sorted by category. You can also sort them alphabetically if that makes life easier.

Required Information	Description
NETWORK OPTIONS: Protocol	Choose TCP, Socket, Pipe, or Memory as the connection protocol.
Compress	True/False: Compresses all information sent between the client and the server if both support compression.
PRE-LOAD OPTIONS: Delete Table Data	True/False: Delete all table data before the data file is imported.
Force Continue	True/False: If there's an error, force the program to continue processing the load file.
Lock Tables	True/False: Keep tables locked to writing before processing text files.
Input File Location	Local/Server: Where the input files are stored.
Ignore Header Rows	Supply a value or leave blank. The default is zero. You are asking to ignore x number of rows from the data file.
FORMATTING: Fields Terminated By	Include the value by which the fields are terminated.
Fields Enclosed By Display	Include the value by which the fields are enclosed.
Fields Optionally Enclosed By	Include the value by which the fields are optionally enclosed.
Fields Escaped By	Include the value by which the fields are escaped.
Lines Terminated By	Include the value by which the fields are terminated.

Duplicate Key Options

Choose among handling duplicate entries by enabling new rows to replace existing rows, ignore input rows that duplicate existing rows, or allowing an error to occur when a duplicate is found and ignore the remainder of the input file.

Summary

Here you find a summary of all the choices you have made in creating this operation. If you want to make changes, you can make them in this grid or go **Back**. You also have the option of saving the operation. Saved files are identified with a .DBImport extension.

Utilities

DBArtisan incorporates a number of powerful utilities that let you create, edit and manage your development environment. You can use utilities to:

- Conduct extensive database searches across multiple databases.
- Perform complex cross-platform database migrations.
- Extract schema for multiple database objects.
- Load and unload table data.
- Configure your server.
- Monitor and analyze all processes on the server.
- Schedule tasks.

Cross-Platform Utilities

The table below describes the utilities for IBM DB2 for Linux, Unix, and Windows, IBM DB2 for z/OS and OS/390, Microsoft SQL Server, MySQL, Oracle, and Sybase ASE:

Utility	Description
Using the schema extraction wizard	The Schema Extraction Wizard lets you reverse-engineer an entire database or any portion of it.
Using the schema migration wizard	The Schema Migration Wizard offers robust cross-DBMS-platform object migration capabilities that can migrate schema and table data between databases on the same or different servers.
Using the Schema Publication wizard	The Schema Publication Wizard lets you generate HTML reports documenting the details of your databases.
Database Monitor	The Database Monitor displays all processes connected to the server and lets you scan your databases to determine potential performance problems.
SQL Logfile	DBArtisan lets you log all of its SQL statements to a file so you have an audit trail of application activity.

IBM DB2 for Linux, Unix, and Windows Utilities

The table below describes the utilities for IBM DB2 for Linux, Unix, and Windows:

Utility	Description
Data Export Wizard for IBM DB2 for Linux, Unix, and Windows	The Data Export Wizard lets you export table data.
Data Import Wizard for IBM DB2 for Linux, Unix, and Windows	The Data Import Wizard lets you import table data.
Database Manager	The Database Manager lets you view and manage your databases.

>

Microsoft SQL Server Utilities

The table below describes the utilities for Microsoft SQL Server:

Utility	Description
Data Unload Wizard for Microsoft SQL Server and Sybase ASE	The Data Unload Wizard lets you export table data.
Data Load Wizard for Microsoft SQL Server and Sybase ASE	The Data Load Wizard lets you load data.
Database Manager	The Database Manager lets you view and manage your databases.
Database Backup Wizard	The Database Backup Wizard lets you backup your databases without having to know the underlying syntax.
Database Recovery Wizard	The Database Recovery Wizard lets you restore your databases on a separate server to enhance system availability.

Oracle Utilities

The table below describes the utilities for Oracle:

Utility	Description
Data Export Wizard for Oracle	The Data Export Wizard lets you extract object definitions and table data from an Oracle database and create a binary export file.
Data Import Wizard for Oracle	The Data Import Wizard lets you extract object definitions from an Oracle Export file and inputting them into the database.
Archive Activity	DBArtisan lets you operate an Oracle database in ARCHIVELOG or NOARCHIVELOG mode.
Backup Control Files	DBArtisan lets you use the graphical interface for constructing the ALTER DATABASE command needed to backup your control files.
Chained Row Analysis	DBArtisan lets you analyze chained rows.
Free Space Deficit	The Free Space Deficit Monitor identifies any objects not be able to acquire their next extent.
Instance Parameter Configuration	The Instance Parameter Configuration lets you view and edit the settings for each Oracle instance.
Configuring Resource Cost	DBArtisan lets you use the visual interface for constructing an ALTER RESOURCE COST command.
Session Information	DBArtisan lets you view session information.

Schema extraction and migration utilities

DBArtisan lets you migrate schema and table data between databases on the same or different servers. Similarly, you can build and store a schema extraction script. Both features employ a wizard that walks you through the process and make use of similar options.

For detailed information, see the following topics:

- Using the schema extraction wizard
- Using the schema migration wizard

Using the schema extraction wizard

The Schema Extraction wizard walks you through the process of building an extraction script, executing it for storage of the database or object schema in a file. The basic process includes choosing the server and database objects for extractions and specifying options that dictate how the extraction script is built. This utility allows you to:

- Extract and archive entire databases or sets of database objects.
- Generate the complete definition of any database object type.
- Reverse-engineer schema, and also place the SQL statements in their proper sequence, regardless of the objects you select.
- Create a new schema extraction operation.
- Reuse a previously defined schema extraction operation.
- Specify the target object types to extract.
- Specify the dependent object types automatically extracted for any object type.
- Set the general script type and options.
- Preview the definition of your schema extraction operation.

To build and execute a schema extraction script:

1. Connect to the datasource where you want to extract schema. For details, see Datasource Management.
2. On the Datasource Explorer, select the target database or database object.
3. On the **Utilities** menu, click **Schema Extraction**.

   The Extraction Wizard opens.
4. Use the following table as a guide to understanding and setting options in the wizard:

Panel	Tasks and Settings	Description
Playback	**Perform new extraction** and **Playback an existing extractio**n	Lets you start a new extraction or select an existing extraction script you want to replay from the drop-down list. **NOTE**: When the script opens, you can make changes to the operation you are replaying. The changes you make will overwrite the current script.
Connection	**Source Datasource**	Lets you select the datasource from which you want to extract schema.
Catalogs (only displayed for multiple database DBMS datasources)	**Source Database**	Lets you select one or more databases from which you want to extract schema.

Panel	Tasks and Settings	Description
Object Selection	**Schemas**	Opens a dialog that lets you narrow the choice of candidate objects by selecting only those associated with specified schema.
	Object Types	Lets you select the object types that will be extracted. You can use the expand/collapse icons to hide and show the dependent object types for each object. As you select object types, all objects of that type (that satisfy the **Schemas** criteria) are made available for selection in the **Objects** list. Selected dependent object types are also made available in the **Objects** list. For information on setting the dependent object type automatically selected to be extracted along with each object type, see DDL extract options.
	Objects	Lets you select the specific objects that will be extracted. For each object selected, you can also select the specific dependent objects that are to be extracted.
Options	This panel lets you choose migration/extraction options and view an **Example Preview** script that is updated as you choose options. It also lets you save your options choices as an template, with an option to use the saved template as the default. For a detailed descriptions of the options presened on this panel, see Common schema migration and extraction options.	
Summary	Presents the following options:	
	Output script to file	Lets you output the script performing the extraction to a file with the specified path and name.
	Save existing operation	Lets you save a copy the choices you made on the wizard panels. On subsequent schema extractions, you can then playback the existing extraction, optionally changing option settings, in creating an extraction script. You specify the path and name of the file in which to save the extraction operation.
	Click **Finish** when ready to proceed with the extraction.	
Execute	Displays the status of the extraction with details on the number of objects retrieved, the number of errors detected, and the elapsed time. On completion, the schema definition is stored in the file you specified on the previous panel. Filter options (Show All, Show Errors, and Show Warnings) let you display specific types of messages generated during the extraction process. When the extraction completes, you can also use the following options:	
	Continue	Opens the extraction script in the ISQL Editor. For more information, see SQL Scripting.
	Report	Opens a detailed report on the results of the schema extraction.

Using the schema migration wizard

DBArtisan offers robust cross-DBMS-platform object migration capabilities that can migrate schema and table data between databases on the same or different servers. DBArtisan's migration capabilities automatically resolve differences between DBMS platforms, alleviating the responsibility of knowing the syntax and object variations inherent in database platforms. The Schema Migration Wizard sequences the creation of objects in the proper order to eliminate dependency errors. It also has the intelligence to use the fastest method available for copying table data.

>

The Schema Migration Wizard lets you:

- Create a new schema migration operation.
- Re-use a previously defined schema migration operation.
- Specify a source and a target datasource.
- Specify server object types, database object types, and general migration options.
- Select specific objects and to set certain migration options for each object type.
- Specify the dependent object types automatically migrated for any object type.
- Specify object ownership.
- Specify if the migration operation is to be saved for reuse.

CAUTION: Schema migration jobs created in earlier versions of DBArtisan that were saved with a .migrate file extension must be re-created in DBArtisan 8.1 or later. Only migration jobs with a .sav extension can be reused.

The table below describes the object types that can be migrated to different supported platforms and the utility used to accomplish the migration operation:

Source	Target	Object Type	Data Out Method	Data In Method
DB2	DB2	All Object Types & Data	DB2 Export API Call	DB2 Import API Call
	SQL Server	Tables, data, views, indexes, users	SQL SELECT	BCP Utility
	Sybase ASE	Tables, data, views, indexes, users	SQL SELECT	BCP Utility
	Oracle	Tables, data, views, indexes, users	SQL SELECT	SQL * Loader Utility
SQL Server	SQL Server	All Object Types & Data	BCP Utility	BCP Utility
	Sybase ASE	All Object Types & Data	BCP Utility	BCP Utility
	Oracle	Tables, data, views, indexes, logins	SQL SELECT	SQL * Utility
	DB2	Tables, data, views, indexes, logins	SQL SELECT	DB2 Import API Calls
Oracle	SQL Server	All Object Types & Data	BCP Utility	BCP Utility
	Sybase ASE	Tables, data, views, indexes, logins	BCP Utility	BCP Utility
	Oracle	All Object Types and Data	SQL SELECT	SQL * Utility
	DB2	Tables, data, views, indexes, logins	SQL SELECT	SQL * Utility
Sybase ASE	Sybase ASE	All Object Types & Data	BCP Utility	BCP Utility
	SQL Server	All Object Types & Data	BCP Utility	BCP Utility
	Oracle	Tables, data, views, indexes, logins	SQL SELECT	SQL * Utility
	DB2	Tables, data, views, indexes, logins	SQL SELECT	DB2 Import API Calls

To copy objects between servers, the Schema Migration Wizard establishes separate connections to both the source and target servers. Based on the copy options you choose, it then migrates schema from the source to the target database. If the bulk copy command is used when copying table data, a temporary BCP file is built in the temp directory of your computer.

NOTE: Before migrating, ensure that any object requirements, restrictions, or restrictions are satisfied. For details, see Migration and object restrictions/requirements.

>

To Open the Schema Migration Wizard:

1. Connect to the datasource from which you want to migrate schema. For details, see Datasource Management.
2. On the Datasource Explorer, select the target database or database object.
3. On the **Utilities** menu, click **Schema Migration**.

   The Migration Wizard opens.
4. Use the following table as a guide to understanding and setting options in the wizard:

Panel	Tasks and Settings	Description
Playback	**Perform new migration** and **Playback an existing migratio**n	Lets you start a new migration or select an existing migration script you want to replay from the drop-down list. **NOTE**: When the wizard opens, you can make changes to the operation you are replaying..
Connection	**Source Datasource** and **Target Datasource**	Lets you select the datasource from which you want to migrate the schema and the datasource to which you want to migrate.
Catalogs (only displayed for multiple database DBMS datasources)	**Source Datasbase** and **Target Database**	Lets you select a source and target database.
Object Selection	**Schemas**	Opens a dialog that lets you narrow the choice of candidate objects by selecting only those associated with specified schema.
	Object Types	Lets you select the object types that will be migrated. You can use the expand/collapse icons to hide and show the dependent object types for each object. As you select object types, all objects of that type (that satisfy the **Schemas** criteria) are made available for selection in the **Objects** list. Selected dependent object types are also made available in the **Objects** list. For information on setting the dependent object type automatically selected to be extracted along with each object type, see DDL extract options.
	Objects	Lets you select the specific objects that will be extracted. For each object selected, you can also select the specific dependent objects that are to be extracted.
Options	This panel lets you choose migration options and view an **Example Preview** script that is updated as you choose options. It also lets you save your options choices as an template, with an option to use the saved template as the default. For a detailed descriptions of the options presened on this panel, see Common schema migration and extraction options.	
Summary	This panel presents the following options:	
	Output script to file	UNDER WHAT CIRCUMSTANCES IS THIS FEATURE, IN GENERAL< ENABLED
	Save existing operation	Lets you save a copy the choices you made on the wizard panels. On subsequent schema migrations, you can then playback the saved migration settings, optionally changing option settings. You specify the path and name of the file in which to save the migration operation.
	Click **Finish** when ready to proceed with the extraction.	

Panel	Tasks and Settings	Description
Execute	Displays the status of the migration with details on the number of objects retrieved and created, the number of errors detected, and the elapsed time. Filter options (Show All, Show Errors, and Show Warnings) let you display specific types of messages generated during the migration process. When the extraction completes, you can also use the following option:	
	Continue	Opens the extraction script in the ISQL Editor. For more information, see SQL Scripting.
	Report	Opens a detailed report on the results of the schema migration.

Executing statements before or after creation of functions or procedures

During migration, it can be necessary to have statements executed before and after creation of the procedure or function. This can be useful for example, if you need to create or drop temporary tables used by the function or procedure. DBArtsan supports the use of two tag pairs, ETStart and ETEnd, that let you embed statements in the first comment block of a stored procedure or function. The following shows the expected syntax:

```
create procedure dbo.procname(@a numeric) as
/*
<ETStart>SQL Statement</ETStart>
<ETEnd>SQL Statement</ETEnd>
*/
begin
...
```

Common schema migration and extraction options

When extracting or migrating schema, the wizards present a set of options that let you specify how the extraction or migration script is created. The following table describes these options.

NOTE: Depending on the DBMS you are working with, the types of objects you choose to migrate/extract, and the specific migration or extraction operation, not all options are available.

Tasks and Settings	Description
Group Dependencies	If set to **True**, objects that are dependent on each other are placed together in the script. For example, Table A would be followed by all objects dependent on Table A. If set to **False,** all objects of the same type are extracted simultaneously.
Script... options	Depending on the DBMS you working with and the types of objects you selected, these options let you specify whether you want to include Create statements, Drop statements, System Privileges, Object Privileges, and If Not Exists statements in the script.
Script Owner	True/False to include the owner in the scripts or not.
Owner	Retain Owner/Change Owner
Change Owner To	If you opted to change the owner, choose the new owner to use when scripting options from the drop-down list.
Copy Owner (Migration only)	Lets you retain the owner in the extraction/migration.

Tasks and Settings	Description
Overwrite Existing (Migration only)	Choose one of the following options: **Overwrite Existing**: objects are overwritten with no prompt. Objects on the target datasource are dropped before new objects are created **Prompt Whether To Overwrite**: The user will be prompted before the object is dropped and re-created. (**Note:** this will be done on an object-by-object basis). **Do Not Overwrite Existing**: If the object exists on the target datasource, no attempts is made to recreate it on the target. **Ignore Overwrite**: (default) Artisan will try to create the migrated objects, without checking the existence on the target datasource.
Generate Script (Migration only)	If set to FALSE, the generated DDL is executed on the target datasource. If set to TRUE, DBArtisan generates a script which could be used to migrate the schema at a later date.

For context information, see the following topics:

- Using the schema extraction wizard
- Using the schema migration wizard

Migration and object restrictions/requirements

Some objects have additional requirements, exceptions, or options. Before migrating a database object, check your database platform for additional information:

- Migrating IBM DB2 for Windows/Unix Objects
- Migrating IBM DB2 for OS/390 and z/OS Objects
- Migrating Microsoft SQL Server Objects
- Migrating Oracle Objects
- Migrating Sybase ASE Objects

Migrating IBM DB2 for Windows/Unix Objects

The table describes the condition(s) that need to be satisfied on the target database to migrate objects:

Object	Condition(s) on Target Database
Aliases	The login must already exist on the target server. The user must already exist in the target database.
Indexes	The base tables must already exist in the target database.
Nodegroups	The base tables must already exist in the target database.
Procedures	Any objects referenced by the stored procedure, including tables, views and other stored procedures, must already exist in the target database. The owner of the stored procedure must already exist in the target database.

Object	Condition(s) on Target Database
Tables	Any referenced user datatypes must already exist in the target database. Any defaults referenced by the table must already exist in the target database. Any rules referenced by the table must already exist in the target database. If the tables have foreign key references to tables not included in the migration, these tables must already exist in the target database.
Tablespaces	The container mappings and directories must already exist on the target server.
Triggers	The base tables and their owners must already exist in the target database. Any objects referenced by the trigger, including other tables, views and other stored procedures, must already exist in the target database.
User Datatypes	The owner of the user datatype must exist in the target database.
Users	The database objects must already exist in the target database.

For more information, see:

Migrating IBM DB2 for OS/390 and z/OS Objects

The table describes the condition(s) that need to be satisfied on the target database to migrate objects:

Object	Condition(s) on Target Database
Aliases	The login must already exist on the target server. The user must already exist in the target database.
Indexes	The base tables must already exist in the target database.
Procedures	Any objects referenced by the stored procedure, including tables, views and other stored procedures, must already exist in the target database. The owner of the stored procedure must already exist in the target database.
Tables	Any referenced user datatypes must already exist in the target database. Any defaults referenced by the table must already exist in the target database. Any rules referenced by the table must already exist in the target database. If the tables have foreign key references to tables not included in the migration, these tables must already exist in the target database.
Tablespaces	The container mappings and directories must already exist on the target server.
Triggers	The base tables and their owners must already exist in the target database. Any objects referenced by the trigger, including other tables, views and other stored procedures, must already exist in the target database.
User Datatypes	The owner of the user datatype must exist in the target database.
Users	The database objects must already exist in the target database.
Views	Any objects referenced by the view, including tables and other views, must already exist in the target database. The owner of the view must already exist in the target database.

Migrating Microsoft SQL Server Objects

The table describes the condition(s) that need to be satisfied on the target database to migrate objects:

Object	Condition(s) on Target Database
Aliases	The login must already exist on the target server. The user must already exist in the target database.
Defaults	The owner of the default must exist in the target database. Any bound user datatypes must exist in the target database. Any bound table columns must exist in the target database.
Groups	To migrate any object privileges granted to groups, the database objects must already exist in the target database.
Linked Servers	Links do not migrate. So, after you migrate a linked server, open the Linked Server Editor and create new linked server logins. For more information, see Linked Servers Editor (SQL Server).
Procedures	The user who owns the procedure(s) must already exist in the target database. Any objects referenced by the procedure(s) must already exist in the target database.
Remote Servers	DBArtisan lets you migrate the schema for one or more remote servers between cross-platform DBMSs. Specifically, you can migrate remote servers between different versions of Microsoft SQL Servers and Sybase ASE. To migrate remote servers, DBArtisan establishes separate connections to both the source and target servers and copies the DDL between the source and target databases.
Roles	To migrate any object privileges granted to roles, the database objects must already exist in the target database.
Rules	The owner of the rule must exist in the target database. Any bound user datatypes must exist in the target database. Any bound table columns must exist in the target database.
Segments	To migrate segments successfully, all referenced database fragments must already exist in the target database.
Tables	Any referenced user datatypes must already exist in the target database. If you are copying segment placement also, then the segments must already exist in the target database. Any defaults referenced by the table must already exist in the target database. Any rules referenced by the table must already exist in the target database. If the tables have foreign key references to tables not included in the migration, these tables must already exist in the target database.
Triggers	The base tables and their owners must already exist in the target database. Any objects referenced by the trigger, including other tables, views and other stored procedures, must already exist in the target database.
User Datatypes	The owner of the user datatype must exist in the target database. Any bound defaults and rules must exist in the target database.
Users	The logins must already exist in the target server. The database objects must already exist in the target database.
Views	Any objects referenced by the view, including tables and other views, must already exist in the target database. The owner of the view must already exist in the target database.

Migrating Oracle Objects

The table describes the condition(s) that need to be satisfied on the target database to migrate objects:

Object	Condition(s) on Target Database
Clusters	The user who owns the cluster(s) must already exist in the target database. The tablespace on which the cluster(s) are placed must already exist in the target database.
Directories	To migrate directories successfully, the object must exist in the target database.
Functions	The user who owns the function(s) must already exist in the target database. Any objects referenced by the function(s) must already exist in the target database.
Indexes	The tables on which the indexes are being built must already exist in the target database. The user who owns the indexes must already exist in the target database. The tablespace on which the indexes are being placed must already exist in the target database.
Libraries	The object type must exist on the target server.
Materialized Views	The user who owns the materialized view(s) must already exist in the target database. Any objects referenced by the materialized view(s) must already exist in the target database.
Materialized View Logs	The user who owns the materialized view log(s) must already exist in the target database. Any objects referenced by the materialized view log(s) must already exist in the target database.
Packages	The user who owns the package(s) must already exist in the target database. Any objects referenced by the package(s) must already exist in the target database.
Procedures	The user who owns the procedure(s) must already exist in the target database. Any objects referenced by the procedure(s) must already exist in the target database.
Rollback Segments	The user who owns the rollback segment(s) must already exist in the target database. The tablespaces where the rollback segment(s) are to be placed must already exist in the target database.
Sequences	The user who owns the sequence(s) must already exist in the target database. Any objects referenced by the sequence(s) must already exist in the target database.
Synonyms	The user who owns the synonym(s) must already exist in the target database. The base objects referenced by the synonym(s) already exist in the target database.
Tables	Any referenced user datatypes must already exist in the target database. If you are copying tablespace placement also, then the tablespaces must already exist in the target database. If the tables have foreign key references to tables not included in the migration, these tables must already exist in the target database. **NOTE:** Index-Organized tables are coded to automatically migrate as though Copy Table Constraints in Alter Form is not selected, and Include Table Constraints is selected.
Triggers	The user who owns the trigger(s) must already exist in the target database. The base triggers must already exist in the target database. Any objects referenced by the trigger(s) must already exist in the target database.
Types	The user who owns the object type(s) must already exist in the target database. The base object types must already exist in the target database. Any objects referenced by the object type(s) must already exist in the target database.

Object	Condition(s) on Target Database
Users	The default and temporary tablespaces assigned to the user(s) must already exist in the target database. The profiles assigned to the user(s) must already exist in the target database. Any tablespaces on which quotas are defined must already exist in the target database. Any roles granted to the user(s) must already exist in the target database.
Views	Any objects referenced by the view, including tables and other views, must already exist in the target database. The owner of the view must already exist in the target database.

Migrating Sybase ASE Objects

The table describes the condition(s) that need to be satisfied on the target database to migrate objects:

Object	Condition(s) on Target Database
Aliases	The login must already exist on the target server. The user must already exist in the target database.
Defaults	The owner of the default must exist in the target database. Any bound user datatypes must exist in the target database. Any bound table columns must exist in the target database.
Groups	To migrate any object privileges granted to groups, the database objects must already exist in the target database.
Procedures	The user who owns the procedure(s) must already exist in the target database. Any objects referenced by the procedure(s) must already exist in the target database.
Remote Servers	DBArtisan lets you migrate the schema for one or more remote servers between cross-platform DBMSs. Specifically, you can migrate remote servers between different versions of Sybase ASE and Microsoft SQL Servers.
Roles	To migrate any object privileges granted to roles, the database objects must already exist in the target database.
Rules	The owner of the rule must exist in the target database. Any bound user datatypes must exist in the target database. Any bound table columns must exist in the target database.
Segments	To migrate segments successfully, all referenced database fragments must already exist in the target database.
Tables	Any referenced user datatypes must already exist in the target database. If you are copying segment placement also, then the segments must already exist in the target database. Any defaults referenced by the table must already exist in the target database. Any rules referenced by the table must already exist in the target database. If the tables have foreign key references to tables not included in the migration, these tables must already exist in the target database. **NOTE:** DBArtisan no longer changes the table lock scheme from ALLPAGES back to DATAROWS on the target server when using BCP.exe version dated after December 1997.

>

Object	Condition(s) on Target Database
Triggers	The base tables and their owners must already exist in the target database. Any objects referenced by the trigger, including other tables, views and other stored procedures, must already exist in the target database.
User Datatypes	The owner of the user datatype must exist in the target database. Any bound defaults and rules must exist in the target database.
Users	The logins must already exist in the target server. The database objects must already exist in the target database.
Views	Any objects referenced by the view, including tables and other views, must already exist in the target database. The owner of the view must already exist in the target database.

Using the Schema Publication wizard

The Schema Publication Wizard lets you generate HTML reports documenting the details of your schema, databases, and objects, which you can share over an intranet or the Internet. DBArtisan uses HTML Templates to create the reports which you can manually edit to create customized reports.

Schema publication produces a two-pane report. The left-hand pane lets you navigate through all objects contained in the report while the right-hand pane shows details for an individual object. THe following example shows a report on three Microsoft SQL Server default objects, displaying details for a default named **UTDS.DefaultObject**.

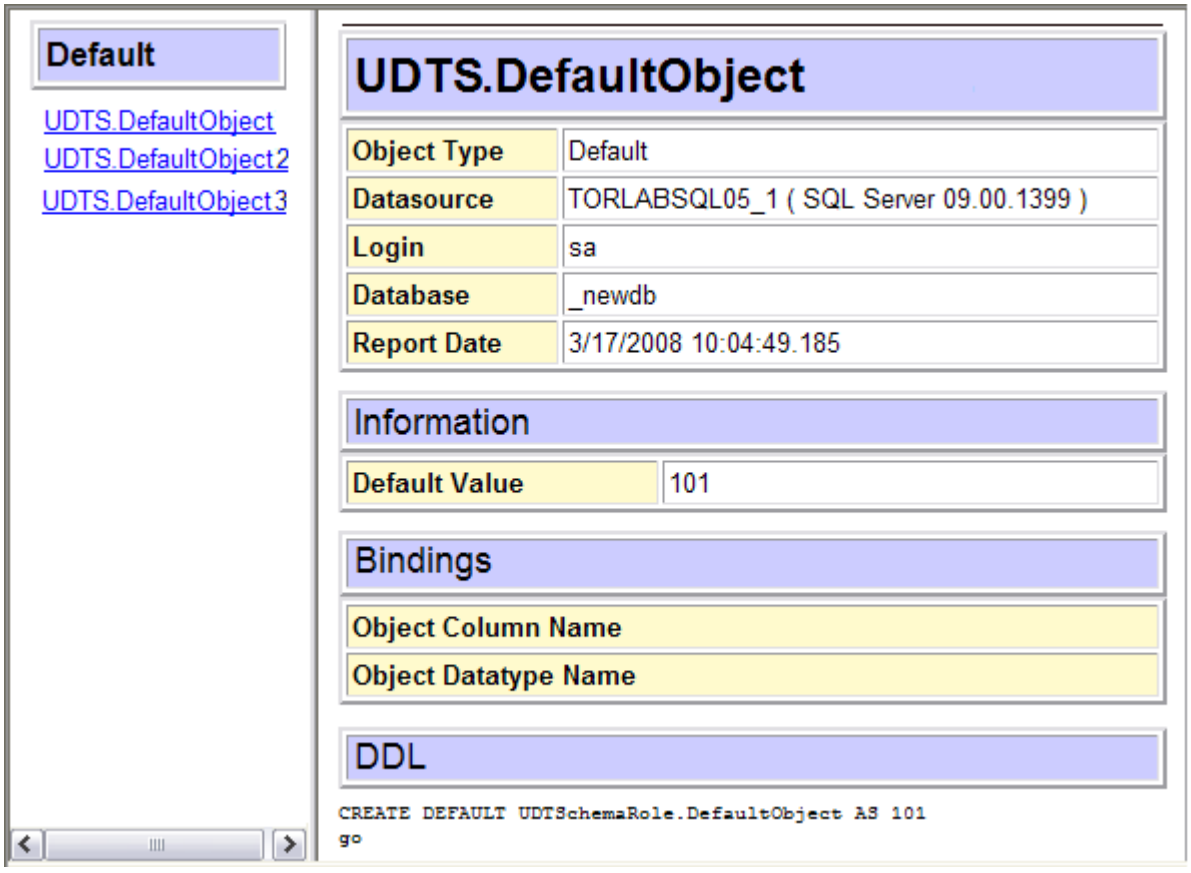

The Schema Publication Wizard lets you:

- Generate reports for all or any level of your database.
- View, find, and print the HTML reports directly in DBArtisan with the built-in HTML browser.
- Create a new schema publication operation.
- Re-use a previously defined schema publication operation.
- Select a datasource for which you want to generate an HTML report.
- Select the server object type and database object type for which you want to publish a schema report and lets you customize the report.
- Select specific objects and to set certain script options for each object type.
- Specify the report directory, HTML file name, and report title.
- Save the schema publication operation for re-use.
- Preview the definition of your schema publication operation.

To create a schema report:

1. Connect to the datasource for which you want to produce a schema report. For details, see Datasource Management.
2. On the **Utilities** menu, click **Schema Publication**.
3. Use the following table as a guide to understanding and setting options in the wizard:

Panel	Tasks and Settings	Description
Welcome	**New Publication Operation**	Select to define a new schema publication operation.
	Previous Publication Operation and **Publication Definition File Name**	Lets you enable use of a previously defined publication operation and provide the path and file name of the operation.
Select a Datasource	Lets you select the datasource to be published.	
Select Database	Only displayed for Microsoft SQL Server or Sybase datasources.	Lets you select the target database.
Select Object Types	**Server Object Types** and **Database Object Types**	Lets you select the server object types and database object types (optionally specifying database objects with a specified owner) that will be published. If you do not enable the **Customize Object List** option, all objects of the selected types will be published.
Select Objects	Only displayed if you selected the **Customize Object List** option on the previous panel.	For the types you chose on the previous panel, lets you use the **Object Type** dropdown to display available objects. You can then select the specific objects for publication.

Panel	Tasks and Settings	Description
Customization	**Report Directory** and **HTML File Name**	Lets you provide the location and name of the file that is to contain the published schema. **NOTE:** To grant access to others, you should specify a shared directory.
	Report Title	Lets you type the title of the report.
	View report file when the operation is complete	Select if you want to view the report immediately after it is generated.
	Create one HTM file for all objects	Select to create one HTM file for all objects.
Save this Operation?	If you want to reuse this schema publication operation, type the location and name of the file that is to store the schema publication operation. Use a **.report** file extension.	
Preview	Lets you review your selections. The settings shown on this panel are read-only, and you must move back to previous panels in order to make changes.	

4 When ready to publish, click **Execute**.

DBArtisan starts the Schema Publication operation and generates an HTML publication report displaying the details of the operation.

Database Monitor

The Database Monitor displays all processes connected to the server. The Database Monitor:

- Lets you scan databases for potential performance problems.
- Includes tabs that provide a wealth of information about your servers.
- Color-codes each process to reflect its state and let you identify problems quickly.

> **TIP:** You can view some Database Monitor Information by clicking the Performance node of the Datasource Explorer.

Important Notes

- Processes and Locks are results of DB2 API calls and are not located in the system catalog tables.
- The Database Monitor may display incomplete SQL due to limitations set by the DBMS. Platforms set limitations on the values returned by the commands that gather the information.

Opening the Database Monitor

1 On the **Utilities** menu, click **Database Monitor**.

DBArtisan opens the Database Monitor.

For more information, see:

- Color-Coding
- Setting Scope Level in the Database Monitor
- Configuring the Database Monitor
- Processes Tab
- Locks Tab

>

- Utilities Tab
- Procedures Tab
- Functions Tab
- Traces Tab
- Logs Tab
- Bufferpools Tab
- All Locks Tab
- Blocking Locks Tab
- Active Rollback Segment Tab
- Data Dictionary Cache Tab
- Datafile I/O Tab
- General Statistics Tab
- Hit Ratio By User Tab
- Latches Tab
- Library Cache Tab
- MTS Dispatcher Contention Tab
- MTS Server Contention Tab
- Open Cursors Tab
- Parallel Queries Tab
- SGA Detail Tab
- SGA Hit Ratios Tab
- SGA Summary Tab
- Tablespace I/O Tab
- System I/O Summary Tab
- Database I/O Summary Tab
- User I/O Summary Tab
- Waits Tab
- Top SQL Cache Tab
- Last 50 Backups Tab
- File I/O Summary Tab
- Top Twenty SQL Tab
- Wait Events Tab

Color-Coding

The Database Monitor color-codes each process to reflect its state, to let you quickly spot problem conditions.

The table below describes the Database Monitor color-codes:

Color	State
Yellow	Sleeping or Inactive
Green	Running
Red	Blocked by another process

For more information, see Database Monitor.

Setting Scope Level in the Database Monitor

NOTE: The functionality is available for IBM DB2 for z/OS and OS/390 only.

The Scope list lets you specify the scope of the element being monitored in a datasharing environment. Data Sharing lets one instance of the DB2 catalog to be shared equally among all members of the group.

The Scope list is available on the following tabs:

Traces Tab

Functions Tab

Procedures Tab

Utilities Tab

To set the scope level, do the following:

1. On the **Database Monitor**, click the target tab.
2. On the **Threads**, **Traces**, **Procedures** and **Functions** tabs:
   - In **Scope**, select **Group** for the monitoring results of all datasharing members in the set.

   OR

   - In **Scope**, select **Local** for the monitoring results for the member that your machine is are connected to.

   DBArtisan displays the results in the Database Monitor.
3. On the **Utility Tab**:
   - In **Scope**, select **Group** for the monitoring results of all datasharing members in the set.

   OR

   - In **Scope**, select **Member** for data by the member name.

   DBArtisan displays the results in the Database Monitor.

Refreshing the Database Monitor

1. On the **Database Monitor** tool bar, click the box to raise or lower the numbers, or type a new value in the box.

For more information, see Database Monitor.

Configuring the Database Monitor

The Database Monitor lets you:

- Toggle to split the screen and display the target process SQL For details, see SQL Toggle.
- Open a SQL window and display the SQL. For details, see ISQL.
- View a Query Plan for SQL running for selected processes. For details, see Query Plan.

SQL Toggle

DBArtisan lets you split the Database Monitor screen and display the SQL running for selected processes.

1 On the **Database Monitor** tool bar, click **Toggle**.

NOTE: To open a SQL window or view a Query Plan for SQL running on selected processes, you must have your Database Monitor Screen split to display SQL on the bottom.

For more information, see Configuring the Database Monitor.

ISQL

DBArtisan displays the SQL running for selected processes in a SQL Editor. You must turn on the SQL Toggle to display the target SQL in the split screen. To display the SQL in a separate window do the following:

1 On the **Database Monitor**, click the target process.

NOTE: To open a SQL Editor or view a Query Plan for processes, SQL Toggle must be on, displaying the target SQL in the split screen.

2 If **SQL Toggle** is not on, on the **Database Monitor** tool bar, click **Toggle**.

3 On the **Database Monitor** tool bar, click **ISQL**.

DBArtisan opens a SQL window with the target process SQL.

4 To execute the SQL, click the **SQL Window** tool bar, click **Execute**.

For more information, see Configuring the Database Monitor.

Query Plan

The Database Monitor lets you view the Query Plan for target processes. To view the Query Plan, do the following:

1 On the **Database Monitor**, select the target process**.**

NOTE: To open a SQL Editor or view a Query Plan for processes, SQL Toggle must be on, displaying the target SQL in the split screen.

2 If **SQL Toggle** is not on, on the **Database Monitor** tool bar, click **Toggle**.

3 On the **Database Monitor** tool bar, click **Plan**.

DBArtisan executes the SQL with query plan and opens a Results Editor with the results.

For more information, see Configuring the Database Monitor.

Session Details

NOTE: The functionality is available for Oracle only.

The Database Monitor lets you view user session details through the User Session Information dialog box. Once you open the User Session Information dialog box, you can view the SQL associated with the session.

The table below describes the tabs of the User Session Information dialog box:

Tab	Description
Statistical Details	Lets you can scroll up and down to see parameters and values.
Waits	Displays either historic or current wait statistics and lets you display or filter out idle wait events. **NOTE:** Data returned for historic session waits is different than the data returned for current session waits.
Current SQL	Lets you view the SQL for this session.

Completing the Session Details Dialog Box

1. On the **Database Monitor** tool bar, click **Session**.

   DBArtisan opens the User Session Information dialog box.

2. On the **Statistical Details Tab**, in the **Session Parameters** grid, you can scroll up and down to see parameters and values.

3. On the **Waits Tab**, view select **Historical Waits** or **Current Waits**.

4. Select **Include Idle Waits** to display idle wait events.

5. On the **Current SQL Tab**, view the SQL for this session.

6. When you finish viewing SQL, click **OK**.

   DBArtisan returns to the Database Monitor dialog box.

For more information, see Configuring the Database Monitor.

Processes Tab

The Database Monitor Processes Tab displays all processes connected to the server. The Processes Tab of the Database Monitor lets you:

- Scan databases for potential performance problems.
- View each process for potential performance problems.
- View login time and seconds idle since last call.

For more information, see Killing Processes.

Killing Processes

If you have system administrator privileges, DBArtisan lets you kill processes from the Database Monitor.

1. On the **Utilities** menu, click **Database Monitor**.

   DBArtisan opens the Database Monitor.

2 Click the target process to kill.

3 On the **Database Monitor** tool bar, click **Kill**.

4 Click **Yes**.

DBArtisan confirms you want to kill the process and closes the DBArtisan message.

For more information, see Configuring the Database Monitor.

Threads Tab

NOTE: This tab is available for IBM DB2 for z/OS and OS/390 only.

The Threads Tab displays the threads on the target datasource.

For more information, see Setting Scope Level in the Database Monitor.

Locks Tab

The Database Monitor Locks Tab displays all processes that are currently holding locks on an IBM DB2 database.

Locks include:

- Shared Lock

A Shared Lock lets other users view and read objects concurrently. The Shared Lock owner or other concurrent processes cannot alter data in tablespaces, tables, or partitions.

- Exclusive Lock

DBArtisan promotes locks from Shared to Exclusive when a user decides to update an object. Once DBArtisan promotes the lock, the user must complete the update before the lock can return to shared mode; another user cannot demote a shared lock. Exclusive locks exclude other users from viewing updates before they are committed. There can only be one exclusive lock for a set of objects. The first thread gets the exclusive lock and holds it until the update is committed. Once the update is committed, the exclusive lock returns to shared mode.

NOTE: You cannot update a lock in DBArtisan.

For more information, see Database Monitor

Utilities Tab

NOTE: This tab is available for IBM DB2 for z/OS and OS/390 only.

The Utilities Tab displays wizards that are typically run from the command line and do not apply to a particular type of schema object. Some of these wizards are also available via the command menus for objects. The wizards include:

- Check Data
- Check Index
- Copying Images
- Copy to Copy
- Load Java

- Merge Copy
- Modify Recovery
- Modify Statistics
- Rebuild Index (DB2 z/OS)
- Reorganize
- Reorganize Tablespace
- Runstats Index
- Runstats Tablespace
- Unload
- Maintain Template

  TIP: To run a wizard, select the command in the right pane of the application and click the Execute button on the Datasource Explorer toolbar.

For more information, see:

Setting Scope Level in the Database Monitor

Database Monitor

Procedures Tab

NOTE: This tab is available for IBM DB2 for z/OS and OS/390 only.

The Procedures Tab displays the procedures on the target datasource.

For more information, see:

Setting Scope Level in the Database Monitor

Database Monitor

Functions Tab

NOTE: This tab is available for IBM DB2 for z/OS and OS/390 only.

The Functions Tab displays the functions on the target datasource.

For more information, see:

Setting Scope Level in the Database Monitor

Database Monitor

Traces Tab

NOTE: This tab is available for IBM DB2 for z/OS and OS/390 only.

>

The Traces Tab displays the traces on the target datasource.

For more information, see:

Setting Scope Level in the Database Monitor

Database Monitor

Logs Tab

NOTE: This tab is available for IBM DB2 for z/OS and OS/390 only.

The Logs Tab displays the logs on the target datasource.

For more information, see Database Monitor.

Bufferpools Tab

NOTE: This tab is available for IBM DB2 for z/OS and OS/390 only.

The Threads Tab displays the bufferpools on the target datasource.

For more information, see Database Monitor.

All Locks Tab

The Database Monitor All Locks Tab displays all processes that are currently holding locks on a database.

For more information, see Database Monitor.

Blocking Locks Tab

The Blocking Locks Tab displays the locks that are blocked.

TIP: Detect blocking locks as they occur, as they can escalate into dead lock situations which can force you to shut down and restart your server to clear the traffic.

For more information, see Database Monitor.

Active Rollback Segment Tab

Rollback or undo segments are used to guarantee transactional integrity. When a transaction has not been committed or rolled back, a rollback segment will be in use with live transactional data. The Active Rollback Segment tab displays information concerning active rollback segments.

For more information, see Database Monitor.

Data Dictionary Cache Tab

The Data Dictionary Cache Tab displays key indicators in determining the performance of the shared pool. It shows how often object definitions are found in memory vs. having to read them in from disk. Because Oracle references the data dictionary many times when an SQL statement is processed, it is imperative that as much of this vital reference information be kept in RAM as possible.

For more information, see Database Monitor.

Datafile I/O Tab

The Datafile I/O Tab displays details concerning the physical I/O activity at the datafile level.

Physical I/O consists of Oracle going to disk to gather or write data. Logical I/O refers to data access performed in memory. The database writer (DBWR) and log writer (LGWR) processes typically perform all the I/O work in the database against the physical datafile used to hold information The checkpoint and archive processes (CKPT and ARCH), also perform I/O work in the database.

For more information, see Database Monitor.

General Statistics Tab

The General Statistics Tab displays general statistics data, including names and values.

For more information, see Database Monitor.

Hit Ratio By User Tab

The database server reads and updates all data from a bufferpool because memory access is much faster than disk access. Data is copied from disk to a bufferpool as needed by the applications using the database. When the server needs to read/write data and the data is already in the bufferpool, no disk access is required. However, if the data is not in the bufferpool, it needs to be read from the disk, which is significantly slower process.

The buffer pool hit ratio indicates the percentage of time that the database server did not need to load a page to service the request for a page.

For more information, see Database Monitor.

Latches Tab

Latches protect many memory structures in Oracle's SGA. They ensure that one and only one process at a time can run or modify any memory structure at the same instant. Much more restrictive than locks (which at least allow for some collective user interaction), latches have no queuing mechanism so either you get it or you do not and are forced to continually retry.

For more information, see Database Monitor.

Library Cache Tab

The library cache holds commonly used SQL statements - basically database code objects. A method for improving performance in Oracle is to encourage the reuse of SQL statements so expensive parse operations may be avoided. The library cache assists this tuning effort.

For more information, see Database Monitor.

MTS Dispatcher Contention Tab

The MTS Dispatcher Contention Tab displays Microsoft Transaction Server data. DBArtisan lets you develop and deploy COM-based applications using Microsoft Transaction Server.

For more information, see Database Monitor.

MTS Server Contention Tab

The MTS Server Contention Tab displays Microsoft Transaction Server data. DBArtisan lets you develop and deploy COM-based applications using Microsoft Transaction Server.

For more information, see Database Monitor.

Open Cursors Tab

Open cursors consume Oracle resources and have the potential to degrade performance, especially if the SQL code being used in the cursor is inefficient. The Open Cursors Tab lets you quickly spot user accounts that have many cursors opened as well as the actual performance statistics for each opened cursor.

For more information, see Database Monitor.

Parallel Queries Tab

The Parallel Queries Tab displays the total number of queries eligible to be run in parallel.

For more information, see Database Monitor.

SGA Detail Tab

The Oracle System Global Area (SGA) is Oracle's memory structural area devoted to facilitating the transfer of data and information between clients and the Oracle database.

For more information, see Database Monitor.

SGA Hit Ratios Tab

The Oracle System Global Area (SGA) is Oracle's memory structural area devoted to facilitating the transfer of data and information between clients and the Oracle database.

For more information, see Database Monitor.

SGA Summary Tab

The Oracle System Global Area (SGA) is Oracle's memory structural area devoted to facilitating the transfer of data and information between clients and the Oracle database.

For more information, see Database Monitor.

Tablespace I/O Tab

The Tablespace I/O tab displays details concerning the physical I/O activity at the tablespace level.

Physical I/O consists of Oracle going to disk to gather or write data. Logical I/O refers to data access performed in memory. The database writer (DBWR) and log writer (LGWR) processes typically perform all the I/O work in the database. Other processes like the checkpoint and archive processes (CKPT and ARCH) may also be used.

For more information, see Database Monitor.

System I/O Summary Tab

NOTE: This tab is available for Microsoft SQL Server only.

The System I/O Summary Tab displays a summary of I/O activities that occur at the system level.

For more information, see Database Monitor.

Database I/O Summary Tab

NOTE: This tab is available for Microsoft SQL Server only.

The Database I/O Summary Tab displays an overview of I/O operations at the database level.

For more information, see Database Monitor.

User I/O Summary Tab

NOTE: This tab is available for Microsoft SQL Server only.

The User I/O Summary Tab displays an overview of I/O operations at the user operation level.

For more information, see Database Monitor.

Waits Tab

NOTE: This tab is available for Microsoft SQL Server only.

The Waits Tab displays a detailed view of all wait activity occurring on the server.

For more information, see Database Monitor.

Top SQL Cache Tab

NOTE: This tab is available for Microsoft SQL Server only.

The Top SQL Cache Tab displays the top 100 objects in the SQL cache, ranked by use counts.

For more information, see Database Monitor.

Last 50 Backups Tab

NOTE: This tab is available for Microsoft SQL Server only.

The Last 50 Backups Tab displays the last fifty backups performed by the Microsoft SQL Server and lets you reference hot backup details at the server level within DBArtisan.

For more information, see Database Monitor.

File I/O Summary Tab

NOTE: This tab is available for Microsoft SQL Server 8.0 only.

The File I/O Summary Tab displays an overview of I/O operations at the database file level.

For more information, see Database Monitor.

Top Twenty SQL Tab

NOTE: This tab is available for Oracle only.

The Top Twenty SQL Tab lets you view the Top Twenty SQL, the top (worst) SQL running in the Oracle shared pool.

For more information, see:

Database Monitor

SQL Toggle

Query Plan

ISQL

Wait Events Tab

NOTE: This tab is available for Oracle only.

The Wait Events Tab lets you analyze bottlenecks and discover what a database is waiting on.

For more information, see Database Monitor.

Instance Configuration

To let you configure and tune your database(s), DBArtisan offers a utility that lets you view and edit the configuration parameters for each server. These configuration parameters are stored in the system catalog and used when restarting the server. For detailed information on the meaning and purpose of each server configuration parameter, see the System Administration Guide and Performance and Tuning Guide.

NOTE: Many changes to instance configuration parameters require you to shutdown and to restart the server before the changes take effect.

1 On the **Datasources** menu, click **Configure**.

2 To modify a parameter, double-click the parameter or click **Edit**.

DBArtisan opens the Edit Configuration dialog box.

NOTE: To modify a parameter, the parameter must have a Yes in the Session Modifiable or System Modifiable column of the Instance Parameter grid.

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

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

SQL Logfile

You can instruct DBArtisan to log all of its SQL statements to a file so that you have an audit trail of application activity. The SQL Logfile captures all the SQL sent by the DBArtisan to populate its editors or to execute your commands. You can use this logfile not only to audit application activity, but also to build SQL scripts for querying system catalogs.

Topics

Activating SQL Logging

You can activate SQL Logging directly from the DBArtisan main menu or toolbar.

Activating SQL Logging

1 On the **Logfile** menu, click **Log SQL**.

Opening the Logfile

You can open the SQL Logfile directly from the main menu.

Opening the Logfile

1 On the **Logfile** menu, click **Open**.

DBArtisan opens the SQL Logfile into an SQL window.

Logfile Options

You can open the Options Editor from the Logfile menu option and make any changes to your logging options.

Opening the Options Editor

1 On the **Logfile** menu, click **Options**.

DBArtisan opens the Options Editor to the Logging Tab.

2 To make any changes to your logging options, make the appropriate changes and then click the **OK** button.

DBArtisan closes the Options Editor.

Truncating the Logfile

If the SQL Logfile is growing too large, you can flush it directly from the main menu.

Truncating the Logfile

1 On the Logfile menu, click **Truncate**.

DBArtisan deletes the contents of the SQL Logfile.

Data Export Wizard for IBM DB2 for Linux, Unix, and Windows

DBArtisan offers a simple, five-panel wizard that lets you export table data. After defining the data export operation, you have the opportunity to save the job definition for future reuse.

The Export Wizard lets you:

- Create a new Export operation.
- Re-use a previously defined Export operation.
- Specify source server, data export file format, and data export specification.
- Select the tables and view to be exported.
- Specify the Output Directory.
- Specify the Message Directory.
- Save the Job Definition for future reuse.
- Preview the definition of your data export operation.

Starting the Export Wizard

1 On the **Datasource Explorer**, select the database object.

DBArtisan displays the database objects in the right pane of the Explorer Window.

>

2 On the **Utilities** menu, click **Data Unload**.

For more information, see Export Wizard - Panel 1.

Export Wizard - Panel 1

The table below describes the options and functionality on the first panel of the Export Wizard:

Option	Description
Do you wish to initiate a new export operation or 'play back' a previously saved operation?	New Export Operation - Select to define a new data unload operation. Previous Export Operation - Select to use a previously defined schema extraction operation. Then in the Export Definition File Name box type the location and name of the job file or click the browse button.

Export Wizard - Panel 2

The table below describes the options and functionality on the second panel of the Export Wizard:

Option	Description
What datasource has the data you wish to export?	Lets you click the list and select the datasource.
Specify the format for the data export file(s) to be generated.	Lets you click the list and select the file format.
Option Button	Available if you select a delimited option. Click to open the File Type Options Dialog Boxesx.
How would you like to specify the data to export?	Enter a SELECT statement. Select tables/views from a list (one export file per table/view).

Export Wizard - Panel 3

Option	Description
Specify the data to be extracted from the database by entering a SELECT statement	Enter the statement that will retrieve the row(s) or column(s) you want from one or more tables. Statements can be complicated. Go to http://msdn.microsoft.com/library and do a search on SELECT for help.
Would you like to store large objects (LOBs) in separate files?	When you choose to store LOBs in separate files, you will be given the chance to store them in files outside the database tablespaces. When you choose not to store your LOBs in separate files, they will exist in a table row in the database.

Export Wizard - LOB Paths

Option	Description
List paths where the LOB files are to be stored	Add, Edit, or Remove files where you want LOBs to be stored
List LOB file names	Add, Edit, or Remove LOB file names.

File Type Options Dialog Boxes

The table below describes the options and functionality of the Integrated Exchange Format File Type Options dialog box:

Option	Description
Accept data despite code page mismatches and suppress translation	
Drop existing indexes and create new ones	
Load each row without checking target length	
Index Schema	

The table below describes the options and functionality of the Integrated Exchange Format File Type Options dialog box:

Option	Description
Accept data despite code page mismatches and suppress translation	
Drop existing indexes and create new ones	
Load each row without checking target length	
Index Schema	

The table below describes the options and functionality of the Delimited ASCII File Type Options dialog box:

Option	Description
Delimiters	Lets you select character string and decimal point specifiers.
Prefix positive decimal values with a blank	
Use ISO date format	

The table below describes the options and functionality of the Non-Delimited ASCII File Type Options dialog box:

Option	Description
Do not recognize x'1A' as the end of the file	
Determine the location of an implied decimal point from the column definition	
Record Length	Maximum is 32,767

Option	Description
Truncate trailing blanks OR Truncate trailing NULLs (0x00) characters	Select one or the other

Export Wizard - Panel 3

The table below describes the options and functionality on the third panel of the Export Wizard:

Option	Description
Owner	Lets you select the tables and/or views you want to export or click the Check All button.
Count Rows	Lets you update the table row counts.
Check All	Lets you select all tables or views.
Uncheck All	Lets you deselect all tables or views.

Export Wizard - Panel 4

The table below describes the options and functionality on the fourth panel of the Export Wizard:

Option	Description
Select a directory for export files (including and LOB files)	Type the location and name of the export file or click the Browse button.
Specify a path and name for the message file	Type the location and name of the message file or click the Browse button.
Do you wish to re-use this data unload operation?	To re-use this data unload operation, in type the location and name of the data unload definition file or find the file by clicking the Browse button. DBArtisan writes a migrate definition file using the *.unload file extension that includes your job definition.

Export Wizard - Panel 5

Review the data export operation description to verify its accuracy.

1. If you need to make any changes, click **Back** to scroll back to the appropriate panels of the wizard to make your corrections.
2. If you are satisfied with the data export operation, click **Execute**.

   For IBM DB2 for z/OS and OS/390, DBArtisan opens the Preview dialog box.

Data Import Wizard for IBM DB2 for Linux, Unix, and Windows

DBArtisan offers a simple, six-panel wizard that lets you load data via DBArtisan's import utility. After defining the data import operation, you have the opportunity to save the job definition for future reuse.

The Import Wizard lets you:

- Create a new Data Load operation.

- Re-use a previously defined Data Load operation.
- Specify data import file.
- Specify file format target datasource.
- Specify import mode.
- Specify target database.
- Specify target schema.
- Specify target table or view.
- Specify the method for selecting columns in a table and select the columns into which the table or view data should be imported.
- Specify import LOBs from separate files, and if so, then specify the location for the message file.
- Define the number of rows to skip in the data file import if the operation restarts.
- Define the number of rows imported before committing to the database.
- Specify the use of Compound SQL statements.
- Save the operation for re-use at a later time.
- Preview the definition of your data unload operation.

Starting the Import Wizard

1 On the **Datasource Explorer**, select the database object.

DBArtisan displays the database objects in the right pane of the Explorer Window.

2 On the **Utilities** menu, click **Data Load**.

DBArtisan opens the Import Wizard - Panel 1.

Import Wizard - Panel 1

The table below describes the options and functionality on the first panel of the Import Wizard:

Option	Description
Do you wish to initiate a new import operation or 'play back' a previously saved operation?	New Data Load Operation - Select to define a new data load operation. Previous Data Load Operation - Select to open a previously defined schema extraction operation. In the Data Load Definition File Name box, type the location and name of the job file or click the browse button.

Import Wizard - Panel 2

The table below describes the options and functionality on the second panel of the Import Wizard:

Option	Description
What file has the data you wish to import?	Lets you type the location and name of the file or click the browse button.

Option	Description
Specify the format of the file to be imported	Lets you click the list and click the file format. Options Button - Click if you want to set the ASCII file type options for a delimited option. Opens the File Type Options Dialog Boxes.
What mode would you like to use for this import operation?	INSERT - The target table must exist and imported rows are inserted into the table without affecting existing rows in the table. INSERT_UPDATE - The target table must exist and have primary key constraint defined on it. Imported rows that match the existing row of data on which the primary key is defined are updated. New rows that do not match the primary key are inserted into the table. REPLACE - The target table must exist and all data is replaced with the contents of the import data file. CREATE - The table defined in the import data file is created and rows of data are inserted into the newly created table. You must use a *.IXF file to use this mode. REPLACE_CREATE - If the target table exists, the data will be replaced with the data contained in the import data file and the table definition remains unchanged. If the target table does not exist, the table defined in the import data file is created and rows of data are inserted into the newly created table. You must use a *.IXF file to use this mode.
What is the target object into which the data will be imported?	Database - Lets you select the database. Schema - Lets you select the schema. Object - Lets you select the table or view.

Import Wizard - Panel 3

The table below describes the options and functionality on the third panel of the Import Wizard:

Option	Description
What method should be used for selecting columns for the export file?	Lets you specify a method: **Default**, **Position**, **Location**, or **Names**.
Indicate the columns into which data should be imported, specifying the relevant column selection attributes	Lets you select the columns. Click Edit to open the Column Position Attributes dialog box.

Column Position Attributes

The table below describes the options and functionality on the Column Position Attributes dialog box:

Option	Description
Include	Select to include the column.
Position	Lets you type the position.

Import Wizard - Panel 4

The table below describes the options and functionality on the fourth panel of the Import Wizard:

Option	Description
Would you like to retrieve large objects (LOB's) from separate files?	Add Button - Click to open the Browse for Folder dialog box. Click the LOB path. Edit Button - Click to open the Browse for Folder dialog box. Remove - Click to remove the LOB path.
Specify a path and name for the message file	Type the directory location of the message file or click the Browse button to find the file.

Import Wizard - Panel 5

The table below describes the options and functionality on the fifth panel of the Import Wizard:

Option	Description
How many of the file's records should be skipped before the import begins?	Restart - Lets you type the value of the number of rows that should be skipped before the import operation begins.
How many records should be imported before committing them to the database?	Commit Records - Lets you type the value of the number of records that should be imported before they are committed to the database.
Use of Compound SQL may improve performance.	Statements - Lets you type the value of the number, which cannot exceed 100, of the number of SQL statements per compound block.

Import Wizard - Panel 6

The table below describes the options and functionality on the sixth panel of the Import Wizard:

Option	Description
Do you wish to re-use this data unload operation?	To re-use the operation, type the location and name of the data load definition file or click the browse button. DBArtisan gives the operation a default extension of *.LOAD.

Import Wizard - Panel 7

Review the data load operation description to verify its accuracy.

1. To make any changes, click **Back** to scroll back to the appropriate panels of the wizard to make your corrections.
2. If you are satisfied with the data load operation, click **Execute**.

   For IBM DB2 for z/OS and OS/390, DBArtisan opens the Preview.

Load Data Utility Wizard for DB2 OS390

The DBArtisan Load Data Utility for DB2 z/OS enables you to load records into one or more tables in a tablespace without creating an intermediate exported file. If the tablespace already contains data, you can opt to add the new data to the existing data or to replace it. The data that's loaded is subject to the same editing or validation routines currently associated with the table and any field procedure associated with particular table columns. As output, you get a loaded tablespace or partition, a file of records that have been discarded, and an error report if you specify the Enforce Constraints option or if unique indexes were part of the load operation.

To use this utility, you need a sufficient set of privileges.

One of the following conditions must be met:

- You must own the table
- Have load privileges for the database
- DBAdmin or DBCTRL authority for the database.
- SYSCTRL or SYSADM authority.

For more complete information than can be presented here, refer to the IBM Utility Guide and Reference. You can find it here: http://www-306.ibm.com/software/data/db2/zos/v8books.html

To Open the Load Data Utility Wizard

1 Connect to a DB2 OS390 database on the **Explorer** tree.

2 Expand the database and click the **Tablespace** subnode.

3 In the right window, select a tablespace and click **Load Data** from the **Command**  drop-down menu.

OR

On the **Explorer** node, click **Utility Execution**, and then click **Load** in the right pane.

Creating a Load Data Utility Operation

The DBArtisan Load Data Utility for DB2 z/OS enables you to load records into one or more tables in a tablespace without creating an intermediate exported file. If the tablespace already contains data, you can opt to add the new data to the existing data or to replace it.

The following sections describe fields you may encounter as you move through the Load Data Utility Wizard:

Initial Settings

Required Information	Description
Do you want the utility to be restartable?	The utility restart information will be included if a recover operation is required.
Tape Stacking	Select this if you are backing up this operation onto physical media. Click **Properties** to specify the particulars for the dataset as it's recorded by opening the Load Utility - Tape Stacking Dialog Box.
Do you want to specify a cursor?	If you want to specify a cursor you need to type the name of the cursor for the input data set in the field below. *No* is the default. Any cursor name you specify cannot exceed eight characters.

Required Information	Description
Specify the In DDN	The input template name. *SYSREC* is the default.
Would you like to preformat?	If you say yes, you are telling the operation to serialize at the tablespace level. This can inhibit concurrent processing of separate partitions. *No* is the default.

Resumption/ Replace/Statistics

Required Information	Description
Resume	If you say **No**, records load into an empty tablespace. This is the default. If you say **Yes**, records are loaded into a non-empty tablespace. Loading begins at the current end of data in the tablespace. If you specify the sharelevel as **change**, the before and after triggers for each row that's loaded are activated. The default sharelevel is **none--**this specifies that applications cannot concurrently access the tablespace during the load operation**.** When you select **Yes**, the remaining options on this panel are unavailable.
Replace	Here you indicate whether the tablespace and its indexes need to be resert to empty before records load. If you select **Yes**, you are telling the load operation to serialize at the tablespace level. **No** is the default.
Would you like to specify the statistics option?	Here you indicate whether you want to gather statistics for the tablespace. **No** is the default. If you select **Yes**, statistics are gathered on a base tablespace, but not an LOB tablespace. If you gather statistics on encrypted data, the results may not be useful.
Do you want to output message to SYSPRINT?	This is the output dataset for messages and printed output. **Default** is the default.
Do you want to force aggregation or rollup processing to be done even though some parts do not contain data?	Yes, No, or default. D**efault** is the default. Here you specify whether you want statistics to aggregate or rollup when RUNSTATS is executed even in the event that some parts are empty.

Select Tables (for running statistics)

Here you are presented with a dialog box where you identify the table you want to run statistics on. Click **ADD** to open the Table/Object Selector dialog box. Click **Set Properties** to open the Select Column for Table dialog box. Click **Delete** to remove a table you selected from the run statistics operation.

Select Tables to Load Data Into

Here you are presented with the dialog box where you identify the tables you want to load the data into. Click **ADD** to select one or more tables. Click **Set Properties** to open the Into Table Specification Dialog Box.

Select Indexes

Here you are presented with a dialog box where you identify the indexes you want to run statistics on. Click **ADD** to open the Index/Object Selector dialog box. Click **Set Properties** to open the Set Correlation option dialog box where you set the KeyCard and Frequent Item options. When you ADD Frequent Options, the Add Frequent Options dialog box opens which identifies the number of columns and the count.

Update/History Statistics

Required Information	Description
Update Statistics	**All**: Indicates that all collect statistics are to be updated in the catalog. *All* is the default. **Access Path**: Indicates that updates are to be made only to the catalog table columns that provide statistics used for selecting the access path. **Space**: Indicates that updates are to be made only to the catalog table columns that provide statistics that help assess the status of a particular tablespace or index. **None**: Indicates no catalog tables will be updated. Note: This is enabled only when the Report option is set to Yes.
History Statistics	This records all catalog table inserts or updates to the catalog history tables. Default: The value specified on **All**: Indicates all collected statistics are to be updated in the catalog history tables. **Access Path:** Indicates that updates are to be made only to the catalog history table columns that provide statistics used for selecting the access path. **Space**: Specifies that only space-related statistics are to be updated in catalog history tables. **None**: Specifies that no catalog history tables are to updated with the statistics that have been collected.

Dictionary/Log

Required Information	Description
Keep dictionary	If you specify **No**, the load operation is prevented from building a new compression dictionary, thereby saving you the costs related to such a build operation. *No* is the default.
Reuse	When used with Replace (on an earlier panel of the wizard), this option indicates that the load operation should reset and reuse the datasets without deleting or refining them. *No* is the default.
Log	The **Default** specifies normal logging during the load operation and is equivalent to **Yes**. All loaded records will be logged. **No** specifies no logging of data in the load operation and disables the copy-pending option below.
LOAD is not to set the tablespace in the COPY-pending status	This, the NOCOPYPEND option will not turn off any copy pending options set prior the load operation. **NOTE**: Specify this option only if the data in this table can be recreated by another load operation if, for some reason, the data is lost. You may not be able to recover the tablespace if you don't take an image copy after the load operation. *No* is the default.

Enforce/Specify

Required Information	Description
Enforce	Specifies whether the load operation should enforce check and referential constraints. **Constraints** indicates constraints will be enforced. If the operation detects a constraint violation, the errant row is deleted and identified by a message. If you elect this option, and there are referential constraints, you must define sort input and output datasets. **Default** enforces the constraints. **No** specifies constraints are not enforcused. It also puts the tablespace in the Check-pending status if at least one referential or check constraint exists.

Required Information	Description
Would you like to specify the Work DDN?	If you elected to enforce constraints, you must specify the work DDN. Here you establish DD names for the temporary sort input and output files. The default input is *SYSUT1* and output is *SORTOUT*.
Would you like to specify the SORTKEYS?	This specifies that the index keys are to be sorted in parallel during the sort build phase to improve performance. If you want, you can specify an integer value to provide an estimate of the number of index keys that are to be sorted. The default is 0 if the input is on tape or a cursor.
Format	This selection is for the format of the input data. **Unload** specifies that the inpurt record format is compatible with DB2 unload format. **SQL/DS** indicates the input record format is compatible with SQL/DS format.

Input/CCSID

Required Information	Description
Input data file type	Select between EBCDIC, ASCII, or Unicode.
CCSID(s)	These coded character set identifiers for the input file generally correspond to the input data file type. If that's not the case, the input data will be converted to the table CCSIDs before being loaded.
Do not accept substitution characters in a string	Specifies that the load operation should not accept substitution characters in the string. Use substitutions when a string is being converted from one data type to another. If you do not accept substitutions and the load operation determines that a substitution appears in the string because of a conversion, if discard processing is active, the record will be moved to the discard file. If discard processing is inactive, the load operation will terminate.
Float	S390 specifies that the load operation should expect floating point numbers are in the System/390 hexadecimal floating point format. IEEE specifies that the load operation should expect that floating point number are in IEEE binary floating point format. If you specify IEEE, the binary format floats will be converted to hexadecimal floating points. If there's a conversion error, the record goes in the discard pile.
Specify the error DDN	Here you specify the DD statement or template for a work dataset that's used during error processing. *SYSERR* is the default.

Map DDN/Discard DDN

Required Information	Description
Would you like to specify the Map DDN?	Specifies the DD statement or template for a work data set to be used in error processing. This is used to correlate the identifier of a table row with the input record that's causing the error. *SYSMAP* is the default.
Would you like to specify the Discard DDN?	Specifies the DD statement or template for a discard dataset that will hold copies for records that aren't loaded in the operation (e.g., a record with a conversion error). *SYSDISC* is the default.
Would you like to specify the maximum number of source records to be written on the discard data set?	Specify an integer with a value between 0 and 2146483647. If the maximum value is reached, the load operation abnormally terminates. When you enter 0, you are specifying that you do not want to set a maximum value and the entire input data set can be discarded.

Required Information	Description
Would you like to specify the sort device type?	Here you indicate, if available, the device type for temporary datasets that are to be dynamically allocated by DFSORT.

Specify Sort Number/CONTINUE-IF

Required Information	Description
Do you want to specify the sort number?	This is the number of temporary datasets that will be allocated by the sort application program.
Would you like to specify a CONTINUE-IF condition?	**Yes** indicates that you want to treat each input record as a portion of a larger record. If a record is true, the next record is concatenated with it before it is loaded.
Start/Finish	Specifies the column numbers of the first and last columns.
Hexidecimal	The byte-string value in the continuation field indicats that the next input record is a continuation of the current load record. Records with this byte-string value are concatenated until the value in the continuation field changes.
Condition	After the CONTINUE-IF completes, write a condition as X'byte-string' or 'character-string'.

Into Table Specification Dialog Box

At least one table INTO STATEMENT is required for each table that is to be loaded. Each Into Table Specification:

- Identifies the table that is to be loaded.
- Describes fields within the input record.
- Defines the format of the input dataset.

NOTE: All tables that are specified must belong to the same tablespace.

Required Information	Description
Page 1:	
Skip fields in the input data set that do not correspond to columns in the target table	Specifies whether or not the load operation should skip fields in the input dataset that do not correspond to tables in the target table.
Partition	An integer value that applies to the number of the partition where records are to be loaded. This is an option that is enabled only for partitioned tablespaces.
Would you like to preformat?	Specifies that any remaining pages should be preformatted up to the high-allocated RBA in the partition and its corresponding partitioning index space. Preformatting happens after the data is loaded and the indexes are built.
Load records into a non-empty partition	**No** loads records into an empty partition. *No* is the default. **Yes** loads records into a non-empty partition. If the partition is empty, you get a warning message, but the partition is loaded anyway.
Replace only the contents of the partition specified above, rather than the entire tablespace.	**NOTE**: You cannot use Load Replace with the parition integer replace option. You can either replace an entire tablespace using the load replace option or the single partition option here.

Required Information	Description
Reset and reuse DB2-managed data sets	When used with the replace option above, this specifies that the load operation should reset and reuse the DB2-managed datasets without deleting and redefining them.
Would you like to keep the dictionary?	This is an option only if a dictionary exists and the partition being loaded has the COMPRESS YES attribute. If the partition has the COMPRESS YES attribute but no dictionary exists, a dictionary is built and you get an error message.
Would you like to specify a cursor?	If desired, specify a cursor for the input dataset.
Select statement for cursor	
In DNN	Specifies the data definition statement or template for the input dataset for the partition. The default is *SYSREC*.
Discard DNN	Specifies the data definition statement or template for the partition. This dataset also holds copies of records that were loaded and removed. It must be a sequential data set.
Page 2:	
Do you want to set criteria for which records are loaded?	If you do not want to specify load criteria, all records in the input dataset are loaded, except for data that is beyond the range of the specified partition.
Do you want to use fleld selection criteria?	Yes or No.
SQL/DS	**Table Owner:** Select a table owner from the drop-down list. **Table:** This is the name of the table that the SQL/DS table is loaded into.
Field selection Criteria	The criteria describes a field and a character constant. A field must contain a character or graphic string. (When compared to a string constant, no datatype conversions are performed.) If the field and the constant are not the same length, the shorter of the two is padded with blanks before a comparison is made.
Fields	Click **Add** to open the Field Specification Dialog Box.

Field Specification Dialog Box

Required Information	Description
Field Name	The name of a field defined by a field specification. This can be any name you choose.
Would you like to specify a position?	Indicates where a field appears in the assembled load record. Start/End are the locations of the first and last field records. The first record is column 1.
Datatype	Pick the appropriate datatype from the drop-down list.
Length/Scale	**Length** is the length in bytes of the input field. **Scale** is the number of digits to the right of the decimal point and must be greater than or equal to 0. *0* is the default.
Would you like to add a condition that causes the db2 column to be loaded with NULL or it's default value?	**NULLIF/CONDITION:** Write a condition that that causes the DB2 column to be loaded with NULL. **DEFAULTIF/CONDITION:** Write a condition that causes the DB2 column to be loaded with its default value. No

>

Set Dataset Info Dialog Box

Required Information	Description
Dataset Name	The name of the input dataset from which the table data should be loaded.
Template Name	
SMS: Management class	The management class name may not exceed 8 characters. This enables the dataset to be catalogued. MGMTCLAS name.
Storage class	The storage class name must be valid and cannot exceed 8 characters. STORCLAS name.
Data class	The data class name cannot exceed 8 characters. DATACLAS name.
Unit: Device Type	Specifies the device type or group name for the data set.
Count	
Volume: Private	A communications connection that's specific to DB2.
Sequence Number	
Count	
Serial no(s)	The serial number of the volume on which a new temporary or permanent data set resides.
Space: Unit	
Primary Allocation	
Secondary Allocation	
Release	
Label: Sequence Number	
Retention period	
Expiration date	

Load Utility - Tape Stacking Dialog Box

Tape stacking allows you to make four copies of the same data set simultaneously. DB2 OS390's Copy to Copy utility allows you to make multiple image copies as well.

Required Information	Description
Local Site Primary Copy	Name the data set where you want the output image copy to go at the local primary site. This relates to the copyddn1 parameter in the stored procedure.
Local Site Backup Copy	Name the data set where you want the output image copy to go at the local backup site. This relates to the copyddn2 parameter in the stored procedure.
Recovery Site Primary Copy	Name the data set where you want the output image copy to go at the recovery primary site. This relates to the recoveryddn1 parameter in the stored procedure.
Recovery Site Backup Copy	Identify the image set where you want the image copy to go at the recovery backup site. This relates to the recoveryddn2 parameter in the stored procedure.
In DDN	Identify the input data set that contains the table to be loaded.
Discard DDN	Identify the work data set that will hold copies of records not included in the load operation.

Data Unload Wizard for Microsoft SQL Server and Sybase ASE

DBArtisan offers a simple, five-panel wizard that can export table data via Microsoft SQL Server's BCP (bulk copy program) utility or Sybase ASE's BCP (bulk copy program) utility. After defining the data export operation, you can save the job definition for future re-use.

The Data Unload Wizard lets you:

- Create a new data unload operation.
- Reuse a previously defined data unload operation.
- Specify source server, data unload file format, and data unload specification.
- Select the tables that you want to unload.
- Save the job definition for future re-use.

For more information, see Data Unload Wizard - Panel 1.

Data Unload Wizard - Panel 1

The table below describes the options and functionality on the first panel of the Data Unload Wizard:

Option	Description
Do you wish to initiate a new data unload operation or 'play back' a previously saved operation?	New Data Unload Operation - Select to define a new data unload operation. 'Playback' Previous Data Unload Operation - Select to use a previously defined schema extraction operation. Then in the Data Unload Definition File Name box type the location and name of the job file or click the browse button.

Data Unload Wizard - Panel 2

The table below describes the options and functionality on the second panel of the Data Unload Wizard.

Option	Description
In which directory should the unloaded files be placed?	Lets you specify the directory by typing the location and name of the directory or find the directory by clicking the browse button.
Which database has the data you wish to unload?	Lets you click the list, and then select the server. Then click the list, and then click the database.
What character should delimit columns in the unloaded data files?	Lets you click the list, and then select the character.
Use BCP Utility	Select if you want to use the Microsoft SQL Server's BCP (bulk copy program) utility.

Data Unload Wizard - Panel 3

The table below describes the options and functionality on the third panel of the Data Unload Wizard for Microsoft SQL Server.

Option	Description
Table Owner	Lets you select the target tables to unload.
Upd Stats button	Select the target tables and click button to update the table row counts.
Select All button	Click to select all tables.
Unselect All button	Click to unselect all tables.

The table below describes the options and functionality on the fourth panel of the Data Load Wizard for Sybase ASE:

Option	Description
Row Terminator	Lets you enter a value for this common BCP option.
Packet Size	Lets you enter a value for this common BCP option.
Text/Image Size	By default, Sybase only copies the first 32K of data in a text or image field. To override default size, type a new size.
Client Char Set	Lets you enter a value for this common BCP option.

Data Unload Wizard - Panel 4

The table below describes the options and functionality on the fourth panel of the Data Unload Wizard.

Option	Description
Do you wish to re-use this data unload operation?	To re-use this data unload operation, type the location and name of the data unload definition file or find the file by clicking the Browse button. DBArtisan writes a migrate definition file using the *.unload file extension that includes your job definition.

Data Unload Wizard - Panel 5

Review the data unload operation description to verify its accuracy.

1 If you need to make any changes, click **Back** to scroll back to the appropriate panels of the wizard to make your corrections.

2 If you are satisfied with the data unload operation, click **Execute**.

DBArtisan starts the Data Unload operation.

To schedule the data unload operation, click **Schedule**.

DBArtisan lets you schedule the task. For details, see Scheduling.

Data Load Wizard for Microsoft SQL Server and Sybase ASE

DBArtisan offers a simple, six-step wizard that can load data via Microsoft SQL Server's BCP (bulk copy program) utility or via Sybase ASE's BCP (bulk copy program) utility. After defining the data import operation, you have the opportunity to save the job definition for future re-use.

The Data Load Wizard lets you:

- Create a new data load operation.
- Re-use a previously defined data load operation.
- Specify data import file.
- Specify target database.
- Specify target schema.
- Specify target table or view.
- Specify the method for selecting columns in a table and select the columns into which the table or view data should be imported.
- Specify import LOBs from separate files, and if so, then specify the location for the message file.
- Define the number of rows to skip in the data file import if the operation restarts.
- Define the number of rows imported before committing to the database.
- Specify the use of Compound SQL statements.
- Save the operation for re-use at a later time.
- Preview the definition of your data unload operation.

For more information, see Data Load Wizard - Panel 1.

Data Load Wizard - Panel 1

The table below describes the options and functionality on the first panel of the Data Load Wizard.

Option	Description
Do you wish to initiate a new data load operation or 'play back' a previously saved operation?	New Data Load Operation - Select to define a new data load operation. Previous Data Load Operation - Select to open a previously defined schema extraction operation. In the Data Load Definition File Name box, type the location and name of the job file or click the browse button.

Data Load Wizard - Panel 2

The table below describes the options and functionality on the second panel of the Data Load Wizard.

Option	Description
Which database do you want to load the data into?	Lets you click the list, and then select the server. Then click the list and select the database.
Which table do you want to load the data into?	Lets you click the list, and then select the table.

Option	Description
Use BCP Utility	Select to use the Microsoft SQL Server's BCP (bulk copy program) utility.

Data Load Wizard - Panel 3

The table below describes the options and functionality on the third panel of the Data Load Wizard.

Option	Description
Which rows do you want to import?	All Rows Range - Lets you type the range in the First Row and Last Row boxes.
What restrictions do you want to place on the import?	**Max Errors** or **Batch Size**.

Data Load Wizard - Panel 4

The table below describes the options and functionality on the fourth panel of the Data Load Wizard for Microsoft SQL Server.

Option	Description
Data File	Lets you type the directory location of the data file or click the browse button to find the file.
Format	To use a format file, type the directory location of the format file or click the browse button to find the file.
Error File	To output errors to an error file, type the directory location of the format file or click the browse button to find the file.
What character delimited the columns in the datafile?	Lets you click the list, and then click the character.

The table below describes the options and functionality on the fourth panel of the Data Load Wizard for Sybase ASE:

Option	Description
Row Terminator	Lets you enter a value for this common BCP option.
Packet Size	Lets you enter a value for this common BCP option.
Text/Image Size	By default, Sybase only copies the first 32K of data in a text or image field. To override default size, type a new size.
Client Char Set	Lets you enter a value for this common BCP option.
Max Connections	Lets you specify the maximum connection of connections in the connection pool.
Show-fi	Provides the option to add columns for each functional index.
Hide-vcc	Provides the option to hide virtual computed columns

Data Load Wizard - Panel 5

The table below describes the options and functionality on the fifth panel of the Data Load Wizard for Microsoft SQL Server.

Option	Description
Do you wish to re-use this data load operation?	To re-use the operation, type the location and name of the data load definition file or click the browse button to find the file. DBArtisan gives the operation a default extension of *.LOAD.
Do you wish to DROP the index(es) on the target table BEFORE copying data?	Drop Index(es) Before Loading of Data - Select to drop the indexes on the target table before copying data. Recreate Index(es) After Loading of Data - Select to recreate these indexes after the data loads.

The table below describes the options and functionality on the fifth panel of the Data Load Wizard for Sybase ASE.

Option	Description
Data File	Lets you type the directory location of the data file or click the browse button to find the file.
Format	To use a format file, type the directory location of the format file or click the browse button to find the file.
Error File	To output errors to an error file, type the directory location of the format file or click the browse button to find the file.
What character delimited the columns in the datafile?	Lets you click the list, and then click the character.

Data Load Wizard - Panel 6

For Microsoft SQL Server, the sixth panel of the Data Load Wizard lets you preview the definition of your data load operation.

1. If you need to make any changes, click **Back** to scroll back to the appropriate panels of the wizard to make your corrections.
2. If you are satisfied with the data load operation, click **Execute**.

   DBArtisan starts the Data Load operation.
3. To schedule the data load operation, click **Schedule**.

   DBArtisan lets you schedule the task. For details, see Scheduling.

The table below describes the options and functionality on the sixth panel of the Data Load Wizard for Sybase.

Option	Description
Do you wish to re-use this data load operation?	To re-use the operation, type the location and name of the data load definition file or click the browse button to find the file. DBArtisan gives the operation a default extension of *.LOAD.
Do you wish to DROP the index(es) on the target table BEFORE copying data?	Drop Index(es) Before Loading of Data - Select to drop the indexes on the target table before copying data. Recreate Index(es) After Loading of Data - Select to recreate these indexes after the data loads.

Data Load Wizard - Panel 7

Review the data load operation description to verify its accuracy.

1 If you need to make any changes, click **Back** to scroll back to the appropriate panels of the wizard to make your corrections.

2 If you are satisfied with the data load operation, click **Execute**.

DBArtisan starts the Data Load operation.

3 To schedule the data load operation, click **Schedule**.

DBArtisan lets you schedule the task. For details, see Scheduling.

Database Manager

DBArtisan lets you view and manage your databases through Database Manager. The Database Manager lets you:

- Add, modify, and delete database files.
- Add, modify, and delete transaction logfiles.
- Add or delete database options.
- View data space and transaction log usage.
- View DDL.

Opening the Database Manager

1 On the **Datasource Explorer,** select the **Databases** node.

DBArtisan displays the databases in the right pane of the Explorer window.

2 In the right pane of the Explorer window**,** click the database link(s).

3 On the **Utilities** menu, click **Database Manager**.

Database Backup Wizard

Implementing a plan to back up critical databases is perhaps the most important job for any database administrator. Without a sound disaster recovery plan, your company has no safeguards against hardware and software failures.

The Database Backup Wizard lets you backup your databases without knowing the underlying syntax. It provides support for backing up databases as well as transaction logs, including commands to stripe the backup across multiple devices.

The Database Backup Wizard lets you:

- Create a new dump operation.
- Reuse a previously defined operation.
- Specify the target server, database, and type of backup operation.
- Specify backup device type.
- Specify how to report the progress of the backup operation.

You can back up the Microsoft SQL Server while the database is online and in use. However, the backup will fail if you are performing one of the following operations while attempting the back up:

- Creating indexes.
- Performing non-logged operations.
- Shrinking the database.
- Creating or deleting database files.

NOTE: Backups created with Microsoft SQL Server version 7 cannot be restored to an earlier version of SQL Server.

Creating or Editing a Backup Operation

DBArtisan lets you create different backup operations using the Database Backup Wizard. After you have successfully created a backup operation, you can reopen and edit it using the same wizard and make the changes you want.

To Open the Database Backup Wizard

1. On the **Datasource Explorer,** select the **Databases** node.

   DBArtisan displays the databases in the right pane of the Explorer window.

2. Select the target database.
3. On the **Utilities** menu, click **Database Backup**.

In the following sections, you'll find all the fields you'll encounter as you create or edit a backup operation:

Job Category Wizard

Database Backup Wizard

Job Category Wizard

The Job Category Wizard panels appear ONLY when you are performing a backup operation for a MS SQL datasource. The job categories are relevant to the Backup Analyst and are described below:

Option	Description
Standard Backup	Any backup job which is a full database, differential database, file and file group, transaction log or table backup which does not include the compression or encryption options.
Advanced Backup	Any backup job which is a full database, differential database, file and file group, transaction log or table backup which includes at least one of the compression or encryption options.
Recovery	Any recovery job created by Backup Analyst.
Would you like to create the 16 popular SQL Server Alerts for common Backup and Recovery Errors?	Yes - Opens the Alerts wizard. No - Opens the Database Backup Wizard.

Database Backup Wizard

The first panel of the Database Backup Wizard welcomes you to the Backup Wizard.

>

Subsequent panels of the wizard are as follows:

Playback or New Operation

Database Selection for SQL Server

Backup Type for SQL Server

Backup Type for SQL Server

Backup Style

Backup Location for MS SQL

Sybase Server Selection

Backup Device and Location for Sybase

Backup and Media Information

Specify Device-handling Options for Sybase

Backup Options

Backup Operation Summary

Backup Operation Execution Selection

Playback or New Operation

Required Information	Description
New Dump Operation	This initiates the process of creating a new backup operation.
Previous Dump Operation	If you want to replay a backup operation, when you select the radio button, the browse function is enabled so you can find the .DBDump file you want to reuse.

Database Selection for SQL Server

Here you select the database you want to backup from a drop-down list and name the backup operation. You can keep the default name for the operation or replace it with one you type. The job description is optional and you can accept the default or write another.

Sybase Server Selection

Required Information	Description
What are the source Sybase server and database?	Select a server from the drop-down list and then the appropriate datasource.

>

Required Information	Description
What do you want to dump?	Database - Backs up the entire database. The size of the database is listed in parentheses. Transaction Log - Backs up the inactive portion of the transaction log. The size of the transaction log is listed in parentheses. Transaction Log With No Log - Removes the inactive portion of the transaction log without making a backup of it. You should only use this command when you run out of room in the database and cannot issue DUMP TRAN WITH TRUNCATE_ONLY. Transaction Log With Truncate Only - Removes the inactive portion of the transaction log without making a backup of it. You should use this command prior to backing up a database. Transaction Log With No Truncate - Makes it possible to dump a transaction log even if the database is inaccessible. You can use this option when the data portion of your database is damaged and the transaction log resides on a separate device.

>

Backup Type for SQL Server

Required Information	Description
Full Database	Backs up the entire database. The size of the database is listed in parentheses.
Differential Database	Records all modifications made to a database after the last database backup, not just the changes made after the last differential database backup.
File and File Group	Select to specify a file or file group.
Transaction Log	Full Truncate- Backs up the inactive portion of the transaction log. The size of the transaction log is listed in parentheses. With No Log - Removes the inactive portion of the transaction log without making a backup of it. You should only use this command when you run out of room in the database and cannot issue DUMP TRAN WITH TRUNCATE_ONLY. Full No Truncate - Makes it possible to dump a transaction log even if the database is inaccessible. You can use this option when the data portion of your database is damaged and the transaction log resides on a separate device. With Truncate Only - Removes the inactive portion of the transaction log without making a backup of it. You should use this command prior to backing up a database.

Backup Style

Required Information	Description
Backup Analyst Backup	Select to use the Backup Analyst for your backup, assuming the Backup Analyst is installed on the database you are backing up. Also lets you compress the backup and apply an encrypted password.
Standard SQL Server Backup	Select to continue the Database Backup Wizard.

Backup Location for MS SQL

Required Information	Description
Backup Device	Depending on what is available to you, you can choose to backup to a disk , floppy, pipe,or tape device.
Disk Directory and Name	Lets you specify a backup location on your machine or click the browse button and select a location.
View Contents	Click to open the Backup Media Contents Dialog Box.

Backup Device and Location for Sybase

Required Information	Description
What type of dump device do you want to use?	Choose among tape device, disk device, or file. When you select file, you need to type the file's name.
Where should messages about the dump operation be sent?	Your options are the Operator Console, Client, or Default.
Dump file	It's up to you to specify the name of the dump file.

Backup and Media Information

Here, for MS SQL, you specify the backup set name, backup description, backup set password, media name, media description, and media name password.

>

For Sybase, select the dump device you want to use, if applicable. Clicking **Attribute** opens the Backup Device editor. For details, see Backup Devices Editor (SQL Server).

Specify Device-handling Options for Sybase

Required Information	Description
Overwrite Contents	Select to overwrite the contents of a tape, and in the Retain Days box, type the number of days to retain the contents of the tape.
Compress	Lets you compress dumps and logs using different levels of compression. The valid levels are 1-9, with level 9 compressing the file to its smallest size.
Dismount After Completion	Select to dismount the tape after backup.
Rewind and Unload After Completion	Select rewind the tape and unload after the backup.

Backup Options

Required Information	Description
Expire Date	The date when the backup set expires and can be overwritten. If supplied as a variable (@takeover), this date is specified as either a string constant (@takeover = date), as a variable of character string data type (except for the text or text data types), a small datetime, or datetime variable, and must follow the configured system datetime format.
Retain Days	The number of days that must elapse before this backup media set can be overwritten. If supplied as a variable (@days_var), it must be specified as an integer.
Skip backup set expiration and name checking before overwriting	Microsoft SQL Server will skip the set expiration and name checking before overwriting.
Overwrite all backup sets on the device	Microsoft SQL Server will overwrite the contents of the device.
Format the entire media before backup operation	Will format the entire media before the backup operation starts.
Read and verify the integrity after the backup after finish	After the operation completes, the integrity of the backup is verified.
Backup Threads	0 is the default. Zero enables the server default value.

Backup Operation Summary

Here you indicate whether or not you want to save the operation, and if so, where the file should be saved. Also displayed are the elements of the function you are creating. You can go back to earlier panels if you want to make changes before the job runs.

Backup Operation Execution Selection

Required Information	Description
Open the backup script in an ISQL editor	The backup script will be opened into a new ISQL session. **NOTE**: This option is not available for Sybase backups.
Perform the backup now	The backup will be executed immediately.
Schedule the backup for later execution	The backup will be scheduled to run at a specified date/time.

For more information, see:

Executing a Backup Script in the ISQL Window

Executing a Backup Script in the ISQL Window

To execute the actual command DBArtisan uses to backup your database in the ISQL Window, do the following:

1 On Backup Operation Execution Selection, click **Open**.

DBArtisan opens the ISQL Window, which is used to execute SQL and Transact-SQL commands. The ISQL Window displays the procedure call that Backup Analyst makes to back up SQL Server databases, complete with the parameters needed to perform the backup.

2 To execute your backup, on the **ISQL Editor** toolbar, click **Execute**.

DBArtisan completes your backup and displays a message "Procedure Succeeded" on the Results tab of the ISQL Editor.

Backup Media Contents Dialog Box

Here you find a read-only display of the backup media available for the backup/restore operation you are creating.

Database Recovery Wizard

To both validate your back ups and practice the disaster recovery process, you should regularly restore your database dumps. If possible, every day, you should restore your databases on a separate server to enhance system availability. In this manner, you can maintain a warm backup of your database guarding against downtime from system failure.

The Database Recovery Wizard lets you:

- Create a new load operation.
- Reuse a previously defined operation.
- Specify the target server, database, and type of restore operation.
- Specify the recovery completion state.
- Specify how to report the progress of the load operation.
- Specify which file to load from the device.
- Specify whether to perform safety checks before replacing the database.
- Save the definition of the dump operation for future reuse.

NOTE: Backups created with Microsoft SQL Server version 7 cannot be restored to an earlier version of SQL Server.

NOTE: If you perform a Database Backup, and then Perform a Data Unload on the same database, close all open windows or disconnect from the database before starting the Database Recovery Wizard.

Creating or Editing a Recovery Operation

DBArtisan lets you create different recover operations using the Recovery Wizard. After you have successfully created a recovery operation, you can reopen and edit it using the same wizard and make the changes you want.

>

To Open the Database Recovery Wizard

1. On the **Datasource Explorer,** select the **Databases** node(s).

   DBArtisan displays the databases in the right pane of the Explorer window.

2. Select the target database.

3. On the **Utilities** menu, click **Database Restore**.

The first panel of the Database Backup Wizard welcomes you to the Backup Wizard.

Subsequent panels of the wizard are as follows:

Playback or New Operation

Database Selection for MS SQL

Database Selection for Sybase ASE

Recovery Completion for MS SQL

Recovery Completion for Sybase ASE

Backup and Media Set Information for MS SQL

Dump Device Selection for Sybase ASE

Tape Handling Options for Sybase ASE

Recovery Options

Recover Operation Summary

Action Selection

Playback or New Operation

Required Information	Description
New Recovery/Load Operation	This initiates the process of creating a new recovery operation.
Previous Recovery/Load Operation	If you want to replay a recovery operation, when you select the radio button, the browse function is enabled so you can find the .sqrc file you want to reuse.

Database Selection for MS SQL

Required Information	Description
Target Database	From the drop-down list, specify the target database for the recovery.
Restore Type	Elect to restore the database, filegroups or files, or a transaction log.
Backup Source	From Backup History - Find backup sets between the dates you specify and make your selection. From Device - Specify the device. Click **Add** or **Edit** to open the Recovery Wizard Dialog Box. Or, click Remove or Remove All to eliminate the device(s) from this operation.

Database Selection for Sybase ASE

Here you select the Sybase server and database you want to recover from the drop-down lists.

Recovery Completion for MS SQL

Required Information	Description
Leave database operational	No more transaction logs can be recovered.
Leave database non-operational	You can restore more transaction logs.
Leave database read-only	You can restore more transaction logs.
Force restore over existing database	
Eject tapes (if any) after restoring each backup	Eject tapes after the backup is restored.
Point in time restore	Select the date and time from which you want the restore operation to take place.

Recovery Completion for Sybase ASE

Required Information	Description
Which operation do you want to perform?	Load Database- Specifies that the complete database and transaction log should be restored. Load Transaction Log - Specifies that only the transaction log should be applied to the database. List Database - Validates a load database operation without actually performing the load. List Transaction Log - Validates a load transaction operation without actually performing the load.
Information type you want to see	**Header Only**, **Brief Description**, or **Detailed Description**.
If you are loading or listing a transaction log...	If you are loading the transaction log, in the If you are loading or listing a transaction log... box, type the specific log time you want to load.

Backup and Media Set Information for MS SQL

Here, for MS SQL, you specify the backup set name, backup set password, media name, and media name password. Any existing backup files are listed.

If you are creating at Leave database read-only operation, the Undo File option is enabled. Browse to find the file you want to undo.

Backup and Media Information for Sybase ASE

Required Information	Description
What type of dump device do you want to use?	Choose among: Tape or Disk Device: Existing options for either of these choices become apparent on the next panel. If none are available, you won't see any devices. File: You must type the name of the file you want use
Where should messages about the load operation be sent?	Select among: **Default**, **Client**, or **Console**.
Dump File	You must type the name of the file where you want to send the load operation messages.

Dump Device Selection for Sybase ASE

If you elected to use a tape or disk device, here you select which specific tape/disk device you want to use. You can choose one or more of the devices that appear in the panel. If none appear, you have to go **Back** to pick another device type. IYou can specify tape or disk properties by clicking each device, and then clicking the **Attributes** button. If you click **Attribute**, the Device Attribute Editor dialog box opens.

Tape Handling Options for Sybase ASE

You can opt to dismount after the operation completes or rewind and unload after the operation completes.

Recovery Options

For **MS SQL**, the backup file is identified here, and if you are creating a new database from another database's backup, you can rename files in the Recover As box.

For **Sybase**, you see the default backup file, which you can easily change. You also have the option, for tape device operations, to load the data with decompression.

Recover Operation Summary

You can elect to save the recovery operation. You can see the default location where the recovery operation will be saved or you can browse to another location. The summary panel displays the choices you have made to this point. Go **back** to make changes.

Action Selection

Option	Description
Open the recovery script in an ISQL editor	The recovery script will be opened into a new ISQL session. See Executing a Recovery Script in the ISQL Window. **NOTE**: This option is not available for Sybase recoveries.
Perform the recovery now	The recovery will be executed immediately.
Schedule the recovery for later execution	The recovery will be scheduled to run at a specified date/time.

Device Attribute Editor Dialog Box

The Device Attribute Editor dialog box lets you configure tape/disk device attributes. The choices you make here apply to the device, not only to the current restore job you are configuring.

The table that follows describes the required information:

Required Information	Description
Backup Server Name	Type the name of the backup server where the dump device is located.
Density	Use the drop-down list to select the appropriate value. This is the density of the tape, if a reel to reel tape is the storage device.
Block Size	Enter the appropriate number of bytes. This is amount of data to read/write in a single operation.
Capacity	Enter the appropriate value in megabytes. This is the backup storage capacity.
Volume Name	Enter name of the volume where the backup is stored.

>

Recovery Wizard Dialog Box

The Recovery Wizard dialog box records your selection for the backup device and file you want to include in the backup operation.

NOTE: The Backup Device and Disk Directory options are mutually exclusive.

The table below explains the information fields you encounter:

Required Information	Description
Restore Type	From the drop-down list, choose among Database, Filegroups or files, and Transaction Log recovery operations.
Backup Device	Your choices from the drop-down list are Disk Device, Floppy Device, Pipe Device, and Tape Device. Only when a device is selected, and available, does the right drop-down list populate.
Disk Directory and Name	The default directory populates the directory and name field. You can browse or type to make changes to the default.
Backup Number	Indicate the backup number. When you select View Contents, the Backup Media Contents Dialog Box opens.
Password/Password Confirmation	Create a password if you want to restrict access to the recovery file.

Executing a Recovery Script in the ISQL Window

To execute the actual command DBArtisan uses to recover your database in the ISQL Window, do the following:

1. On Database Recovery Wizard - Last Panel, click **Open**.

   DBArtisan opens the ISQL Window, which is used to execute SQL and Transact-SQL commands. The ISQL Window displays the procedure call that Backup Analyst makes to recover SQL Server databases, complete with the parameters needed to perform the recovery.

2. To execute your recovery, on the **ISQL Editor** toolbar, click **Execute**.

   DBArtisan completes your recovery and displays a message "Procedure Succeeded" on the Results tab of the ISQL Editor.

Database Recovery Wizard - Last Panel

Review the database restore description to verify its accuracy.

1. If you need to make any changes, click **Back** to scroll back to the appropriate panels of the wizard to make your corrections.
2. If you are satisfied with the database restore operation, click **Execute**.

   For IBM DB2 for z/OS and OS/390, DBArtisan opens the Preview

Data Export Wizard for Oracle

DBArtisan provides comprehensive graphical support for the Oracle Export Utilities. DBArtisan offers a simple wizard that guides you through the process of extracting object definitions and table data from an Oracle database and creating binary export file. DBArtisan's Export Wizard ensures that the proper parameters and privileges are used, provides access to desired database information and resolves version compatibility differences.

The Export Wizard supports all four export modes:

- Table
- Users
- Full Database
- Point-in-time Recovery (available in Oracle8)

The benefits of using the Oracle Export Utilities range from moving data between different versions of Oracle quickly and efficiently, saving space on the database, to selectively backing up an entire or portion of a database using cumulative or incremental export functions.

To use the Export Wizard, you must specify the path and name for the Oracle Export utility in the Oracle Utility Tab of the Options Editor.

NOTE: If you do not have the EXP_FULL_DATABASE role assigned to you, you may not be able to access all the functionality of the export wizard.

Starting the Data Export Wizard

1 On the **Datasource Explorer**, select the database object.

DBArtisan displays the database objects in the right pane of the Datasource Explorer.

2 On the **Utilities** menu, click **Database Export**.

DBArtisan opens the Data Export Wizard - Panel 1.

Data Export Wizard - Panel 1

The first panel of the Export Wizard lets you specify the source datasource, name an existing parameter file or use the existing Oracle default parameter file.

The table below describes the options and functionality on the first panel of the Data Export Wizard:

Option	Description
Which database has the data you wish to export?	Click the list, and then click the database.
Edit an existing parameter file	Type the location and name of the parameter file or click the browse button to find the file.
Default values	To use the Oracle default values for the parameter file, leave this box blank.

Data Export Wizard - Panel 2

The second panel of the Export Wizard lets you select the export mode that determines the database objects you will export.

- The full database mode makes available incremental, cumulative or complete export operations.
- The point-in-time export operation can export one or more tablespaces, which you can use in an import operation to recover a tablespace at a prior point in time without affecting the rest of the database.

The table below describes the objects, related objects and structures that are exported for each mode:

Mode	Object Types and Structures
Table	Table definitions, object types defined on the table, table data by partition, nested table data, indexes defined on the table, table constraints, table permissions, column and table comments, table analysis information, triggers that refer to the table, and auditing information.
User	Libraries, object types, database links, sequence numbers, cluster definitions, table definitions, object types defined on the table, table data by partition, nested table data, indexes defined on the table, table constraints, table permissions, table analysis information, column and table comments, private synonyms, stored procedures, functions, packages, auditing information, views, cluster analysis information, referential constraints, triggers, snapshots, and snapshot logs.
Full Database	Tablespace definitions, tablespace quotas, profiles, user definitions, roles, role privileges, system privileges, resource costs, rollback segment definitions, database links, sequence numbers, directories, libraries, object types, cluster definitions, password information, default and system auditing, object type tables, table definitions, table data by partition, nested table data, table indexes, table constraints, table privileges, table analysis information, column and table comments, auditing information, synonyms, views, stored procedures, functions, packages, triggers, cluster analysis information, snapshots, snapshot logs, refresh groups and children.
Point-In-Time (available in Oracle8)	Tablespaces.

The table below describes the options and functionality on the second panel of the Data Export Wizard:

Option	Description
Which mode would you like to use for this export?	Click the option button that corresponds to the export mode you want to implement: **Table**, **User**, **Full Database**, or **Point-in-Time Recovery**.

Data Export Wizard - Panel 3

The third panel of the Export Wizard lets you specify the objects you want to export based on the mode you selected in the second panel of the Export Wizard.

The table below describes the options and functionality on the third panel of the Data Export Wizard for Table mode:

Option	Description
Select the tables and/or partitions you would like to export	To display list of available tables, click the Owner list, and then click the owner.
Tables	In the box, select the check boxes that correspond to the tables you want to export. If a table has partitions, select the check boxes that correspond the partitions you want to export. To select all tables, click the Select All button.

The table below describes the options and functionality on the third panel of the Data Export Wizard for User mode:

Option	Description
Select the users whose objects you would like to export	To display list of available users, select the check boxes that correspond to the objects you want to export. To select all objects, click the Select All button.

The table below describes the options and functionality on the third panel of the Data Export Wizard for Full Database mode:

Option	Description
Would you like to perform an incremental export	An incremental operation exports all database objects that have changed since the last incremental, cumulative, or complete export operation. NOTE: To perform an incremental export operation, you must have the EXP_FULL_DATABASE role. A cumulative operation exports all database objects that have changed since the last cumulative or complete export. A complete operation exports all objects.
Do you want to records this export in the system tables that track exports?	To record the export operation in the SYS.INEXP, SYS.INCFIL, and SYS.INVID system tables, click the Yes button.

The table below describes the options and functionality on the third panel of the Data Export Wizard for Point-in-Time Recovery mode:

Option	Description
Select the tablespaces to be recovered	To display list of available tablespaces, select the check boxes that correspond to the tablespaces you want to recover. To recover all tablespaces, click the Select All button.

Data Export Wizard - Panel 4

The fourth panel of the Export Wizard lets specify:

- The objects you want to export.
- To view the progress meter.

The table below describes the options and functionality on the fourth panel of the Data Export Wizard:

Option	Description
Check the objects you would like to export	**NOTE:** This option is not available if you are performing a full database incremental export operation. The Oracle Export Utility automatically determines if a row has been changed since the last export operation. Select the check box that corresponds to the object(s) you want to export: To export indexes associated with the tables, select the Indexes check box. To export check constraints or referential integrity constraints, select the Constraints check box. To export object privileges, select the Grants check box. To export table data, select the Rows check box.
Rows/Dot	To view the progress meter which displays the status of rows being exported, in the box, type the value of the number of rows per dot you want displayed in the meter.

Data Export Wizard - Panel 5

The fifth panel of the Export Wizard lets you set the following options:

- The Export Path

- The Length of the File Record
- The Buffer Size

The table below describes the options and functionality on the fifth panel of the Data Export Wizard:

Option	Description
Would you like to perform a direct path export?	A direct path export extracts data much faster as it bypasses the SQL Command processing layer and saves data copies whenever possible.
Would you like to accept your platform's value for BUFSIZ as the length of the file record?	Lets you select the Yes or No option button. If you clicked the No option button, in the Specify length box, type the KB value of the length. The highest value you can specify is 64 KB.
Would you like to accept your platform's default value for the buffer size used to fetch rows?	**NOTE:** This option is only valid if you are performing a conventional path export. Lets you select the Yes or No option button. If you clicked the No option button, in the Specify size box, type the KB value of the length; the highest value you can specify is 64 KB.

Data Export Wizard - Panel 6

The sixth panel of the Export Wizard lets you specify:

- Handling storage of table data.
- Type of optimizer statistics to generate when the file is imported.
- Make the data remain unchanged during the export operation.

The table below describes the options and functionality on the sixth panel of the Data Export Wizard:

Option	Description
Would you like table data to be consolidated into one initial extent upon import?	Lets you select the No or Yes option button.
What type of database optimizer statistics should be generated upon import?	Lets you click the appropriate option button: **Estimate**, **Compute**, or **None**.
Would you like to ensure that the data seen by this export operation does not change during execution?	**NOTE:** This option is not available for an incremental export operation. Lets you select the No or Yes option button. Click Yes to ensure that the data remains intact throughout the export operation.

Data Export Wizard - Panel 7

The seventh panel of the Export Wizard lets you specify names and locations for the export, parameter and logfiles.

The table below describes the options and functionality on the seventh panel of the Data Export Wizard:

Option	Description
Supply the export file name	Lets you type the full location and name of the export file. Export file names use the.dmp file extension. (for example, C:\Temp\EXPBerlin.dmp). If you expect to perform the same export operation again, type the location and name of the parameter file in the corresponding box. Parameter files are stored as.txt files. (for example, C:\Temp\EXPPARBerlin.txt).
Logfile	To capture informational and error messages, click the option button and type the location and name of the logfile in the corresponding box. Export logfile names use the.logfile extension. (for example, C:\Temp\EXPLOGBerlin.log).
None	If you do not want to capture informational and error messages, click the option button.

Data Export Wizard - Panel 8

Review the data export operation description to verify its accuracy.

Completing Panel 8

Review the export operation description to verify its accuracy.

1. If you need to make any changes, click **Back** to scroll back to the appropriate panels of the wizard to make your corrections.
2. If you are satisfied with the export operation, click **Execute**.

   DBArtisan starts the Export operation.

Data Import Wizard for Oracle

DBArtisan provides comprehensive graphical support for the Oracle Import Utilities. DBArtisan offers a simple wizard that guides you through the process of extracting object definitions from an Oracle Export file and inputting them into the database. The Oracle Import utility only reads export files created with the Oracle Export utility. DBArtisan's Import Wizard ensures that the proper parameters and privileges are used, provides access to desired database information and resolves version compatibility differences.

The Import Wizard supports all four import modes:

- Table
- Users
- Full Database
- Point-in-time Recovery (available in Oracle8)

The Oracle Import utilities build database objects in the proper sequence to avoid dependency problems. Tables are created first, then table data is input into the proper tables, next indexes are built, then triggers are imported and finally integrity constraints are built.

To use the Import Wizard, you must specify the path and name for the Oracle Import utility in the Oracle Utility Tab of the Options Editor.

NOTE: If you do not have the EXP_FULL_DATABASE role assigned to you, you may not be able to access all the functionality of the import wizard.

>

Starting the Import Wizard

1 On the **Datasource Explorer**, select the database object.

DBArtisan displays the database objects in the right pane of the Datasource Explorer.

2 On the **Utilities** menu, click **Data Import**.

DBArtisan opens the Data Import Wizard - Panel 1.

Data Import Wizard - Panel 1

The first panel of the Import Wizard asks you to specify the target datasource and to name an existing import parameter file or use the existing Oracle default parameter file.

The table below describes the options and functionality on the first panel of the Data Import Wizard:

Option	Description
Into which database would you like to import data?	Click the list, and then click the database.
Edit an existing parameter file	Type the location and name of the parameter file or click the browse button to find the file.
Default values	To use the Oracle default values for the parameter file, leave this box blank.

Data Import Wizard - Panel 2

The second panel of the Import Wizard lets you specify the export file from which the data should be imported. You must also select the import mode you want to use which determines the database objects that DBArtisan imports. The table below describes the objects, related objects and structures that are imported for each mode:

Mode	Object Types and Structures
Table	Table definitions, object types defined on the table, table data by partition, nested table data, indexes defined on the table, table constraints, table permissions, column and table comments, table analysis information, triggers that refer to the table, and auditing information.
User	Libraries, object types, database links, sequence numbers, cluster definitions, table definitions, object types defined on the table, table data by partition, nested table data, indexes defined on the table, table constraints, table permissions, table analysis information, column and table comments, private synonyms, stored procedures, functions, packages, auditing information, views, cluster analysis information, referential constraints, triggers, snapshots, and snapshot logs.
Full Database	Tablespace definitions, tablespace quotas, profiles, user definitions, roles, role privileges, system privileges, resource costs, rollback segment definitions, database links, sequence numbers, directories, libraries, object types, cluster definitions, password information, default and system auditing, object type tables, table definitions, table data by partition, nested table data, table indexes, table constraints, table privileges, table analysis information, column and table comments, auditing information, synonyms, views, stored procedures, functions, packages, triggers, cluster analysis information, snapshots, snapshot logs, refresh groups and children.
Point-In-Time (available in Oracle8)	Tablespaces.

The table below describes the options and functionality on the second panel of the Data Import Wizard:

Option	Description
Specify the export file from which you want to import data?	Lets you type the location and name export file to be used in the import operation, or find the file by clicking the Browse button.
Which database was used to create this export file?	Lets you choose the database used for export file creation from the list.
Which mode would you like to use for this export?	Lets you click the option button that corresponds to the import mode you want to implement. Table User Full Database ORACLE 8 OR LATER ONLY: Point-In-Time

Data Import Wizard - Panel 3

The third lets you specify which objects you want to import based on the mode you selected in the second panel of the Import Wizard.

The table below describes the options and functionality on the third panel of the Data Import Wizard for Table mode:

Option	Description
Select the tables and/or partitions you would like to import	Displays a list of available tables. Lets you click the Owner list, and then click the owner.
Tables	To In the box, select the check boxes that correspond to the tables you want to import. If a table has partitions, click the check boxes that correspond the partitions you want to import. To select all tables, click the Select All button.

The table below describes the options and functionality on the third panel of the Data Import Wizard for User mode:

Option	Description
Select the users whose objects you would like to import	To display list of available users, select the check boxes that correspond to the objects you want to import. To select all objects, click the Select All button.

The table below describes the options and functionality on the third panel of the Data Import Wizard for Full Database mode:

Option	Description
Would you like to perform an incremental import	An incremental operation imports all database objects that have changed since the last incremental, cumulative, or complete export operation.

Data Import Wizard - Panel 4

The fourth panel lets you specify options to import dependent objects, privileges and table data, and to create a file for index-creation commands.

The table below describes the options and functionality on the fourth panel of the Data Import Wizard:

Option	Description
Check the objects you would like to import	Select the check box that corresponds to the object(s) you want to export. To import object privileges, click the Grants check box. To import table data, select the Rows check box. To import indexes associated with the tables, select the Indexes check box.
Would you like to skip unusable indexes?	Lets you click the Yes or No option button.
Would you like specify a file to receive index-creation commands?	When this parameter is specified, index-creation commands for the requested mode are extracted and written to the specified file, rather than used to create indexes in the database. Tables and other database objects are not imported. If you clicked the Yes option button, type the location and name of the file or find the file by clicking the Browse button.

Data Import Wizard - Panel 5

The fifth panel of the Import Wizard lets you:

- Set the data rows transfer buffer size.
- Indicate if the import operation should commit after each array insert.
- Specify the option to re-use existing data files.
- Specify the option to ignore creation errors.

The table below describes the options and functionality on the fifth panel of the Data Import Wizard:

Option	Description
Would you like to accept your platform's default value for buffer size used to fetch rows?	If you select the No option button, in the Specify length box, type the KB value of the length. The highest value you can specify is 64 KB.
Should Import commit after each array insert?	Lets you select the No or Yes option button.
Would you like to reuse the existing datafiles making up the database?	Select Yes to set the reuse option in the datafile clause of the CREATE TABLESPACE command.
Should Import ignore object creation errors when attempting to create database objects?	Select Yes if you selected Table as your mode of import in panel 2, and the tables exist in the database in the target database.

Data Import Wizard - Panel 6

The sixth panel of the Import Wizard lets you specify:

- The length of the file record.
- How you want to implement SQL Analyze statements contained in the export file.

- To view the contents of the export file rather than import the database information.
- To view the progress meter.

The table below describes the options and functionality on the sixth panel of the Data Import Wizard:

Option	Description
Would you like to show the contents of the export file rather than import them?	Lets you select the No or Yes option button.
Would you like to view the progress meter as rows are imported	Lets you enter the number of rows per dot displayed in the meter box, type the value of the number of rows per dot you want displayed in the meter.
Would you like to accept your platform's value for BUFSIZ as the length of the file record?	If you select the No option button, in the Specify length box, type the KB value of the length; the highest value you can specify is 64 KB.
Would you like to accept your platform's default value for the buffer size used to fetch rows?	**NOTE:** This option is only valid if you are performing a conventional path export. If you click the No option button, in the Specify Size box, type the KB value of the length; the highest value you can specify is 64 KB.
Would you like Import to execute SQL Analyze statements found in the export file?	Lets you select the No or Yes option button.

Data Import Wizard - Panel 7

The seventh panel of the Import Wizard lets you specify names and locations for the parameter and logfiles. If you importing from an Oracle 6 export file, you can define the character set of the original file.

The table below describes the options and functionality on the seventh panel of the Data Import Wizard:

Option	Description
Character Set	For Oracle 6 export files, type the character set of the original file.
Logfile Button	Click to capture informational and error messages, and then type the location and name of the logfile in the corresponding box. Export logfile names use the *.logfile extension. (for example, C:\Temp\EXPLOGBerlin.log).
None Button	Click if you do not want to capture informational and error messages.
Parameter File	If you expect to perform the same export operation again, type the location and name of the parameter file in the corresponding box. Parameter files are stored as .txt files. (for example, C:\Temp\EXPPARBerlin.txt).

Data Import Wizard - Panel 8

Review the data import operation description to verify its accuracy.

Completing Panel 8

Review the import operation description to verify its accuracy.

1. If you need to make any changes, click **Back** to scroll back to the appropriate panels of the wizard to make your corrections.
2. If you are satisfied with the import operation, click **Execute**.

   DBArtisan starts the Import operation.

Archive Activity

Oracle uses Redo logs to record changes made to a database during operation. These files, which are active or online during normal operation, are referred to as the online redo logfiles. Oracle uses these files to reapply changes made during the recovery process that were not written to datafiles before shutdown. You can operate an Oracle database in ARCHIVELOG or NOARCHIVELOG mode. If you choose ARCHIVELOG mode, then Oracle uses a background process to copy the contents of the online redo logfiles as archive redo logfiles in a separate area. ARCHIVELOG mode lets you perform online backups and database media recovery of your Oracle database.

The Archive Activity Monitor lets you to analyze the status, activity levels, proper sizing and performance of your archive redo logs. At a glance, you can determine the ARCHIVELOG mode, configuration and vital statistics about redo log performance. In particular, the Archive Activity Monitor lets you to assess whether or not the redo logfiles are sized properly. It is very important to choose the right size for online redo logfiles. If the redo logs are too small, the Log Writer must switch logfiles too often, which hurts performance. If the redo logs are too big, recovery times are longer.

The table below describes the Archive Activity Monitor options:

Category	Option	Description
Log Archiving	Log Mode	Whether or not ARCHIVELOG mode is active for the Oracle instance.
	Auto Archival	Whether or not the database has been configured to archive redo logs automatically.
	Destination	Directory where the archive redo logs are contained.
	Oldest Log Sequence	Oldest system change number. As a precaution, your backups should include all archived logfiles starting with the oldest log sequence.
	Current Log Sequence	Current system change number. It marks the start of the backup.
Online Redo Archive Activity	Redo Entries	Number entries made to the online redo logs by the Log Writer (LGWR) process.
	Avg Redo Size	Average size, in kilobytes, of a redo log entry.
Redo Log Buffer Contention	Entry Requests	Number of entries made to the redo log buffer by the Log Writer (LGWR) process.
	Entry Wait Times	Number of times that the Oracle Log Writer has had to wait before making an entry to the redo log buffer. This number is accumulated since the last time that the database was started. A high number of waits indicates that the log buffer size is too small. To increase the size, you need to modify the LOG_BUFFER parameter in the Oracle configuration file (e.g., INITORCL.ORA).
Redo Log Switches	Time	When the redo log switch occurred. If redo log switches occur frequently, the redo logs are probably too small. As a rule-of-thumb, redo log switches should not occur more frequently than once every 30 minutes. If they are occurring more frequently, then you should increase the size of the redo logs.
	Record ID in Controlfile	Displays ID number in the control file.

Opening the Archive Activity Dialog Box

1 On the **Utilities** menu, click **Archive Activity.**

DBArtisan opens the Archive Activity dialog box.

Backup Control Files

You should keep many copies of your control files to safeguard against disk failure or loss of the primary control file. DBArtisan provides a graphical interface for constructing the ALTER DATABASE command needed to backup your control files.

Using the Backup Control Files Utility

1 On the **Utilities** menu, click **Backup Control File**

DBArtisan opens the Backup Controlfile Dialog Box.

Backup Controlfile Dialog Box

The table below describes the options and functionality available on the Backup Controlfile dialog box:

Option	Description
Backup Controlfile Name	Type the name of the backup control file in the box. To reuse an existing Back Up Controlfile, type the name of the existing file, and then select Reuse Existing Backup Controlfile. **NOTE:** When you type in the Backup Controlfile Name box, the check box is activated.
Reuse Existing Backup Controlfile check box	Select this check box if you want to reuse an existing file. The name of the file and the name entered in the Backup Controlfile Name box must be the same.
Trace SQL Statements	Click the option button that corresponds to how you want Trace SQL statements handled: None - No statements are traced. Reset Logs - Traces reset logs. No Reset Log - Does not trace reset logs.

Chained Row Analysis

Updates can cause data rows to migrate and to chain over multiple data blocks. Chaining degrades 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.

To avoid chained rows, you can tune the Percent Free storage settings for that table or cluster. By increasing the Percent Free on each data block, you decrease the likelihood that an update causes a data row to expand over multiple data blocks.

If the row length is greater than the size of a data block, you cannot avoid chained rows.

DBArtisan provides a utility that lets you identify chained rows quickly by displaying chained rows for all or any schema in your Oracle database.

Analyzing Chained Rows

1 On the **Utilities** menu, click **Chained Row Analysis**.

DBArtisan opens the Chained Row Analysis dialog box.

The table below describes the options and functionality on the Chained Row Analysis dialog box:

Option	Description
Owner	Lets you select the owner for whom you want to scan for chained rows.
Refresh Button	Click to refresh the list.

Free Space Deficit

One of the central tasks of every Oracle administrator is to monitor the availability of free space in tablespaces. When a data segment (e.g., table, index or cluster) is created on a tablespace, it acquires an initial extent of a determined size. As it grows, the data segment will eventually need to acquire another extent to store additional data. If a data segment cannot acquire its next extent, all further transactions on it will fail.

The Free Space Deficit Monitor identifies any objects that will not be able to acquire their next extent. Whenever it identifies a problem data segment, you should take corrective action as soon as possible. To address the problem, you can:

1 Reduce the object's Next Extent Size from the appropriate Object Editor.

2 Increase the availability of free space in the tablespace by:

- Adding datafiles.
- Activating autoextend for the tablespace's datafiles. Autoextend allows the datafiles to grow to fill available room on the physical drive. This option is only available with Oracle.

The table below describes the objects:

Category	Information	Description
Tablespace	Name	Name of the tablespace.
	Largest Free Section	Largest contiguous section of free space in a tablespace.
Objects	Object Name	Name of any object that will not be able to acquire its next extent because it will be larger than the Largest Free Section.
	Object Type	Type of the object, such as a table, index or cluster, that will not be able to acquire its next extent.
	Next Extent	Specified size of the object's next extent. One of the options you have to correct the problem is to reduce the next extent size so that it is smaller than the Largest Free Section.

Opening the Free Space Deficit Dialog Box

1 On the **Utilities** menu, click **Free Space Deficit.**

DBArtisan opens the Free Space Deficit dialog box.

The table below describes the options and functionality on the Free Space Deficit dialog box:

Option	Description
Owner	Lets you click the owner for whom you want to analyze free space.

Pinned Code

For Oracle, DBArtisan now offers a Pinned Code Utility that lets you pin or unpin code objects in memory. This ensures that frequently accessed database stored procedures, triggers, functions, etc., are available in memory, which improves their overall execution performance.

> **NOTE:** To create the DBMS_SHARED_POOL package log on as SYS and run the DBMSPOOL.SQL script that is found in the ORACLE_HOME\RDBMS\ADMIN directory.

Opening the Pin Code Dialog Box

1 On the **Utilities** menu, click **Pinned Code.**

DBArtisan opens the Pin Code dialog box.

The table below describes the options and functionality on Pin Code dialog box.

Option	Description
User	Lets you select the object user.
Pin Button	To pin an object, select the target object, and then click the Pin button.
UnPin Button	To unpin an object, select the target object, and then click the UnPin button.
Flush SGA Button	Click to flush the Oracle shared pool to remove fragmentation and any unwanted code objects from memory.

You must install the SYS.DBMS_SHARED_POOL package on your Oracle server to use this utility.

Confirm Shared Pool Flush

The Confirmed Shared Pool Flush dialog box lets you confirm that you want to flush the Shared Pool. Flushing the Shared Pool causes all non-pinned code objects to be removed. After you flush the Shared Pool, all non-pinned objects are parsed and reloaded back into the pool on the next execution. This can extend first-time execution response times.

Maximum Extent Limits

When an object reaches its maximum extent limit, it cannot acquire any new free space. This can cause critical problems in a production database. DBArtisan's Maximum Extent Limits for Oracle tool lets you quickly find and correct space problems.

Opening the Maximum Extent Limits Dialog Box

1 On the **Utilities** menu, click **Maximum Extent Limits.**

DBArtisan opens the Maximum Extent Limits dialog box.

The table below describes the options and functionality on Maximum Extent Limits dialog box.

Option	Description
Object Owner	Lets you select owner of the objects to display the Objects box.
Objects	Objects for the user account(s) that have reached their maximum extent limit.

Option	Description
Edit Button	To increase an object's maximum extent limit, edit the object and adjust its maximum extent property to a higher value or to unlimited.

Instance Parameter Configuration

DBArtisan lets you configure and tune your Oracle databases. DBArtisan offers a utility that lets you view and edit the settings for each Oracle instance. DBArtisan derives its information from the system catalog. To reconfigure your Oracle instance, you must edit the parameter settings in the applicable configuration file (for example, INITORCL.ORA).

NOTE: After changing the settings, shut down and restart the database server for the changes to take effect.

Changing an Instance Parameter

1 On the **Datasources** menu, click **Configure**.

DBArtisan opens the Instance Parameters dialog box.

2 To modify an instance parameter, double-click the target instance parameter or click the **Edit** button to open the **Edit Configuration** dialog box.

NOTE: To modify an instance parameter, the parameter must have a Yes in the Session Modifiable or System Modifiable column of the Instance Parameter grid.

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

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

Extracting Instance Configuration Parameters

DBArtisan lets you extract the statements for setting and displaying your instance configuration parameters.

From the Datasource Explorer

1 On the **Datasource Explorer**, select the **Configuration** node.

DBArtisan displays the parameters in the right pane of the Datasource Explorer.

2 In the right pane of the **Datasource Explorer**, click the target parameter(s).

3 On the **Datasource Explorer** toolbar, click **Extract**.

OR

In the right-pane of the **Datasource Explorer**, right-click the parameters, and then click **Extract**.

DBArtisan extracts the definition to a text window.

Generating a Report on Configuration Parameters

DBArtisan lets you generate a detailed report containing information about your instance configuration parameters. The reports derive their contents from the information displayed in the right hand side of the Datasource Explorer. To facilitate the online publication of these reports, DBArtisan generates the report in Hypertext Mark-Up Language (HTML). You can view the report directly in DBArtisan's built-in HTML-browser.

NOTE: For more information on generating reports, see Using the Schema Publication wizard.

Generating a Instance Configuration Parameters Report

1 On the **Datasource Explorer**, select the **Configuration** node.

DBArtisan displays the parameters in the right pane of the Datasource Explorer.

2 In the right pane of the **Datasource Explorer**, click the target parameter(s).

3 On the **Datasource Explorer** toolbar, click **Report**.

OR

In the right-pane of the **Datasource Explorer**, right-click the parameters, and then click **Report**.

DBArtisan opens the Report dialog box.

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

Option	Functionality
Report Home Page File Name	To change the report home page file name, in the box, type the location and name of the file or find the file by clicking the Browse icon.
Report Title	To change the report title, in the box, type the location and name of the file or click the Browse Button to find the file.

Configuring Resource Cost

When you create a profile, you can assign a composite resource limit to it. The composite limit is a weighted average of four individual resource limits:

- Connect time
- Private SGA
- CPU per session
- Logical reads per session

The default setting is to weight each of these measures equally. You can alter the weights by issuing an ALTER RESOURCE COST command. DBArtisan provides a visual interface for constructing such a command.

Altering Resource Cost

1 On the **Utilities** menu, click **Alter Resource Cost**.

DBArtisan opens the Alter Resource Cost dialog box.

The table below describes the options and functionality on the Alter Resource Cost dialog box:

Option	Description
Weightings	In the box, specify the weightings for each resource limit by typing the weight value in the corresponding box: **Connect time**, **Private SGA**, **CPU per session**, or **Logical reads per sessio**n.
Enable Resource Limits	Select to To enable the resource limit changes immediately.

Session Information

Oracle tracks detailed information about the resources used by each session on a database. This information is useful for analyzing resource utilization patterns between user sessions to monitor performance.

The information displayed is derived from the Oracle system tables, SYS.V_$STATNAME and SYS.V_$SESSTAT. Consult the Oracle documentation or detailed descriptions in the meanings of each session parameter.

> **TIP:** Waits lets you analyze bottlenecks and discover what a database is waiting on.

Viewing Session Information

1. On the **Utilities** menu, click **Session Information**.

   DBArtisan opens the User Session Information dialog box.

2. Complete the User Session Information Dialog Box.
3. When you finish viewing SQL, click **Close** to return to the **User Session Information** dialog box.
4. When you finish viewing the session information, click **Close**.

Resource Limits

Resource limits are sets of parameters that let you maximize your datasource assets. By setting resource limits, you can prevent transactions and queries from congesting server resources. This ensures that you do not:

- Exceed estimated or actual input/output demands determined by the optimizer.
- Return a number of rows beyond a set limit.
- Exceed a defined elapsed time.

DBArtisan includes the Resource Limits dialog box that lets you perform the following tasks:

- Showing Resource Limits
- Adding Resource Limits
- Editing Resource Limits
- Dropping Resource Limits

Opening the Resource Limits Dialog Box

You must connect to a datasource running Sybase before you can open the Resource Limits dialog box. To open the Resource Limits dialog box, do the following:

1 On the **Datasource Explorer**, select the target datasource node.

2 On the **Utilities** menu, click **Resource Limits**.

DBArtisan opens the Resource Limits dialog box.

The table below describes the columns in the Resource Limits table of the Resource Limits dialog box:

Column	Description
name	The login name for the resource limit.
appname	The application for the resource limit.
rangename	The time range in which the resource limit is enforced. The time range is a contiguous block of time.
rangeid	The value of the rangename.
limitid	The value and type of the resource limit.
limitvalue	The value of the resource limit.
enforced	The value and enforcement time for the resource limit.
action	The code number of the action for the resource limit. The codes numbers are: 1=Issue a warning 2=Abort the query batch 3=Abort the transaction 4=Kill the session
scope	The code number of the scope for the resource limit. The codes numbers are: 1=Query 2=Query batch 4=Transaction 6=Batch and transaction

Showing Resource Limits

DBArtisan displays all resource limits for the current datasource. In addition, DBArtisan lets you use the Resource Limits dialog box to search for and display specific resource limits. In the Resource Limits dialog box, you can specify search criteria for resource limits. DBArtisan displays the resource limits that meet your search criteria in the Resource Limits table of the Resource Limits dialog box.

Displaying Resource Limits

1 On the **Datasource Explorer**, select the target datasource node.

2 On the **Utilities** menu, click **Resource Limits**.

DBArtisan opens the Resource Limits dialog box.

3 Specify the search parameters. The table below describes the search parameters in the Resource Limits dialog box:

Parameter	Description
Login list	Lets you specify the target login or logins.
Application list	Lets you specify an application or applications for the limit criteria.
Limit Time check box	Lets you specify a time range in which to enforce the resource limit.
Limit Time list	Lets you specify a start time for the resource limit. The start time must be within the proceeding 24 hours. To specify the start time, you must select the Limit Time check box.
Limit Day list	Lets you specify the day or day(s) for the resource limit.
Scope list	Lets you specify the scope of the target action(s). On the Scope list, the numbers next to each scope display in the scope column of the Resource Limits table.
Action list	Lets you specify the action to take place when a resource limit is exceeded. On the Action list, the numbers next to each action display in the action column of the Resource Limits table.

4 When you finish specifying search parameters, click **List Limits**.

DBArtisan lists the resource limits meeting the target criteria in the Resource Limits table.

TIP: To sort the display of the Resource Limits list, click a column heading.

Adding Resource Limits

DBArtisan offers a simple graphical wizard that guides you through the process of creating a resource limit without requiring your to know any of the underlying commands.

Starting the Resource Limit Wizard

1 On the **Datasource Explorer**, select the target datasource node.

2 On the **Utilities** menu, click **Resource Limits**.

DBArtisan opens the Resource Limits dialog box.

3 In the **Resource Limits** dialog box, click **Add**.

DBArtisan opens the Resource Limit Wizard - Panel 1

Resource Limit Wizard - Panel 1

The first panel of the Resource Limit Wizard lets you identify the objects to which the resource limit will apply.

The table below describes the options and functionality on the first panel of the Resource Limit Wizard:

Option	Description
Add Button	Click to open the Add/Edit Resource Limit Object Dialog Box.
Edit Button	Click to open the Add/Edit Resource Limit Object Dialog Box.

Add/Edit Resource Limit Object Dialog Box

The table below describes the options and functionality on the Add/Edit Resource Limit Object dialog box:

Option	Description
Login	Lets you click the list, and then click the target user login.
Application	Click the list, and then click the target application.

Resource Limit Wizard - Panel 2

The second panel of the Resource Limit Wizard lets you specify:

- The type of resource limit.
- The value of the resource limit type.
- The time during which the limit will be enforced.

The table below describes the options and functionality on the second panel of the Resource Limit Wizard:

Option	Description
Select the type of resource limit	Lets you specify the type of resource limit by clicking one of the following option buttons: **Input/output cost**, **Elapsed time**, or **Row count**.
Limit Value box	**NOTE:** You must type the value for the limit to continue.
Time Range	Lets you click the list, and then click the target named range.
Modify, Add, and Edit buttons	Click to add or edit a named time range. DBArtisan opens the Named Time Ranges Dialog Box.
Pre-execution	To specify the execution times during which the limit will be enforced, select the target check box(es):
Execution	To specify the execution times during which the limit will be enforced, select the target check box(es):

Named Time Ranges Dialog Box

A time range is a contiguous set of days. For example if you designate a Friday start day and a Monday end day, the time range includes the weekend.

The table below describes the options and functionality on the Named Time Ranges dialog box:

Option	Description
Name	In the box, type a name for the new time range.
Start time	In the box, specify the start time.
End time	In the box, specify the end time.
Start day	In the box, specify the beginning of the time range.
End day	In the box, specify the end of the time range.

Resource Limit Wizard - Panel 3

The third panel of the Resource Limit Wizard lets you specify the scope of the resource limit and the action to perform when the limit is exceeded. The available scopes are dependent on the type of resource limit you selected on the previous panel of the wizard.

The table below describes the types and available scopes:

Resource Limit Type	Available Scope
Input/output cost	Query
Elapsed time	Query batch, Transaction, and Query batch and transaction
Row count	Query

The table below describes the options and functionality on the third panel of the Resource Limit Wizard:

Option	Description
Scope	Click one of the option buttons to specify the scope of the resource limit: Query Query batch Transaction Query batch and transaction
Action	Click one the following option buttons to specify the action to take when the resource limit is exceeded: **Issue a warning**, **Abort the query batch**, **Abort the transaction**, or K**ill the session**.

Resource Limit Wizard - Last Panel

The last panel of the Resource Limit Wizard lets you:

- Preview the SQL for function that adds the resource limit.
- Execute the SQL for function that adds the resource limit.

Completing the Last Panel

1 Review the SQL.

2 Click **Execute**.

Editing Resource Limits

DBArtisan lets you edit resource limits. To edit resource limits, do the following:

1 On the **Datasource Explorer**, select the target datasource node.

2 On the **Utilities** menu, click **Resource Limits**.

DBArtisan opens the Resource Limits dialog box.

3 In the **Resource Limits** grid, select the target resource limit.

4 Click **Edit**.

DBArtisan opens the Resource Limit Editor.

5 To modify the limit value, in the **Limit Value** box edit the value.

6 To change the type of action, click a different option button.

7 Click **OK**.

Dropping Resource Limits

DBArtisan lets you drop resource limits. To drop resource limits, do the following:

1 On the **Datasource Explorer**, select the target datasource node.

2 On the **Utilities** menu, click **Resource Limits**.

DBArtisan opens the Resource Limits dialog box.

3 In the **Resource Limits** grid, select the target resource limit.

4 Click **Drop**.

DBArtisan opens a confirm drop message.

5 Click **Yes**.

Top Sessions

Being able to identify the top resource users in an Oracle system is a valuable tuning asset. The Top Sessions Options Dialog Box lets you quickly pinpoint the heaviest users of I/O and memory usage. The Top Sessions utility quickly finds and ranks user sessions based on their resource usage.

Opening the Top Sessions Dialog Box

1 On the **Utilities** menu, click **Top Sessions.**

DBArtisan opens the Top Sessions dialog box.

The table below describes the options and functionality on the Top Sessions dialog box.

TIP: Click the Options button to open the Top Sessions Options Dialog Box.

Top Sessions Options Dialog Box

The table below describes the options and functionality on the Top Sessions Options dialog box:

Option	Description
Exclude Sys and System	Select to force the query used to display top session data to be changed.
Exclude Background Process	Select to force the query used to display top session data to be changed.
Limit Top Session Display to	Lets you specify the number of sessions to display. These options will force the query used to display top session data to be changed.
No Limit	
Sort Top Sessions By	Lets you custom soft the display

User Session Information Dialog Box

The table below describes the options and functionality on the User Session Information dialog box:

Option	Description
User	Lets you select a user.
OS User Name	Displays the operating system user name.
Session ID	Lets you select the application.
Machine Name	Displays the name of the datasource.
Status	Displays the session status.
Terminal Name	Displays the terminal name.
Command	Displays the command.
Program	Displays the program used.
Client Platform	**IBM DB2 FOR WINDOW/UNIX ONLY:** Lets you select the client platform.
Client Protocol	**IBM DB2 FOR WINDOW/UNIX ONLY:** Lets you select the client protocol.
Session Parameters	**IBM DB2 FOR WINDOW/UNIX ONLY:** Lets you select the session parameters.
User Session Information Dialog Box - Statistical Details Tab	**ORACLE ONLY:** Displays statistical details.
User Session Information Dialog Box - Waits Tab	**ORACLE ONLY:** Displays waits.
User Session Information Dialog Box - Current SQL Tab	**ORACLE ONLY:** Displays the current SQL.

User Session Information Dialog Box - Statistical Details Tab

The table below describes the options and functionality on the Statistical Details Tab of the User Session Information Dialog Box:

Option	Description
Statistical Details grid	Displays the session parameters and value (if any).

User Session Information Dialog Box - Waits Tab

The table below describes the options and functionality on the Statistical Details Tab of the User Session Information Dialog Box:

Option	Description
Historical Waits	If selected, displays historical waits in the waits grid.
Current Waits	If selected, displays current waits in the waits grid.
Include Idle Waits	If selected, displays idle waits in the waits grid.
Waits grid	Displays the wait information, including Session ID, Event, Seconds in Wait, Wait Time, State, and other information.

User Session Information Dialog Box - Current SQL Tab

The table below describes the options and functionality on the Statistical Details Tab of the User Session Information Dialog Box:

Option	Description
SQL box	Displays current SQL.

Tools

DBArtisan 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.
- Administer your ODBC datasources.

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

Section	Description
Find in Files	This section describes the Find in Files dialog box that lets you find a phrase or character in your files.
Database Search	This section describes the powerful database search utility that helps you to find instances of a string across multiple databases.
Script Execution Facility	This section describes the 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 for Linux, Unix, and Windows datasources.
File Execution Facility	This section describes the 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.
Scheduling	The DBArtisan scheduling programs and utilities let you schedule and execute jobs on local datasources anytime.
Visual Difference	DBArtisan 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.
Query Builder	This section describes Query Builder, a tool that lets you construct, structure, and manipulate up to five different types of queries simultaneously.
Import Data	How to bring data from an external source into a database table is described. You can import data from a text file or a Microsoft Excel document.
Data Editor	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.
Embarcadero Products	The Tools menu lists all installed Embarcadero Technologies products. This lets you toggle to or start another Embarcadero product.
Code Workbench	This section describes the Code Workbench that lets you 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.

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**.

DBArtisan 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**.

DBArtisan opens the Database Search Wizard - Panel 1.

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 datasource(s) you want to search	Search however many datasources you want. You can search more than one platform at a time and specific owners within each datasource.

1 Click **Next**.

DBArtisan 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 for OS/390, 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 for Linux, Unix, and Windows only.

1 Click **Execute** to start the operation.

DBArtisan displays a progress dialog box while the search runs. When the search completes, DBArtisan opens the Database Search Results window.

Database Search Results

DBArtisan 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 Database Search Wizard.
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

DBArtisan'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 for Linux, Unix, and Windows datasources. After completing a scheduled job, DBArtisan 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 DBArtisan's scheduling facilities, letting you schedule script execution jobs. When used in conjunction with a scheduler, DBArtisan 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 a scheduling program to schedule the script execution.
- Specify execution notification e-mail and Net Send addresses.

For more information, see File Execution Facility.

File Execution Facility

DBArtisan'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 DBArtisan's scheduling facilities, letting you schedule script execution jobs. After completing a scheduled job, DBArtisan generates a results report that lists errors, verifies execution, and details the output of the job. When used in conjunction with a scheduler, DBArtisan 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 DBArtisan scheduling programs to schedule the file execution.
- Specify execution notification e-mail and Net Send addresses.

Completing the Script/File Execution Facility

DBArtisan 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**.

   DBArtisan opens the Script or File Execution Facility dialog box.
2. Complete the Script Execution Facility - Script Tab
3. Complete the File Execution Facility - Files Tab
4. Complete the File/Script Execution Facility - Target Tab
5. Complete the File/Script Execution Facility - Output Tab
6. Complete the File/Script Execution Facility - Notify Tab

For more information, see 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

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

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

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.

>

Option	Description
File Output	If selected, specifies a file output. Directory - Type or browse to enter the full path and directory name in which you want to place the output file. File Type - Specifies a file type. Include column titles when saving - If selected, lets you save column titles. Open files with registered applications - If selected, opens files with registered applications.

For more information, see:

File Execution Facility

Script 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

Scheduling

The DBArtisan scheduling programs and utilities let database administrators schedule and execute jobs on local datasources 24-hours-a-day, 7-days-a-week. After completing a scheduled job, DBArtisan generates a results report that lists errors, verifies script execution, and details the output of the job. DBArtisan can automatically send the results report to any e-mail or network recipients.

DBArtisan offers the following programs and utilities to let you schedule routine tasks and jobs:

- Microsoft Task Scheduler
- ETSQLX Command Line Utility

DBArtisan's scheduling facilities let you:

- Execute scheduled jobs without running DBArtisan.
- Run batch files to automate tasks.
- Run a script in parallel against multiple datasources.
- Automatically send result reports to e-mail or network recipients when jobs complete.

Microsoft Task Scheduler

DBArtisan 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, DBArtisan provides you with a link to the Microsoft Web site where you can download the Microsoft Task Scheduler at no cost.

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

Option	Description
Job Name	Lets you type the name of the job.
Job Description	Lets you type a job description that will appear in the subject line in your e-mail and Net Send messages.
E-mail Address(es)	Lets you type the e-mail address(es) to which you want to send notifications.
Net Send	Lets you type the network user(s) to whom you want to send notifications.
Action Output Directory	Lets you type the full path for the directory in which to place the output file or click browse to open the **Select Directory** dialog box. For more information, see Select Directory.

The table below describes the tabs on the Job Scheduler dialog box:

Tab	Description
Schedule	Lets you set your scheduling parameters.
Settings	Lets you set your settings parameters.

NOTE: DBArtisan's ETSQLX command line utility runs a scheduled job even if DBArtisan is not running. For more information, see ETSQLX Command Line Utility.

This section includes the following:

- Opening the Microsoft Scheduled Task Directory Through DBArtisan

Select Directory

The select directory dialog box lets you locate the directory in which to place the output file.

For more information, see Microsoft Task Scheduler

Opening the Microsoft Scheduled Task Directory Through DBArtisan

In the Microsoft Scheduled Tasks directory, Microsoft lets you create, view, edit, and delete jobs.

1 On the **Tools** menu, click **Scheduler**.

DBArtisan opens the Microsoft Scheduled Tasks directory.

ETSQLX Command Line Utility

DBArtisan'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

DBArtisan 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**.

DBArtisan 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**.

DBArtisan 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.**

DBArtisan 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**.

DBArtisan 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

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**.

DBArtisan 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.**

DBArtisan 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.**

DBArtisan 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

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

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.

DBArtisan 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

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.

DBArtisan 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

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.

DBArtisan 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**.

DBArtisan accepts the options.

For more information, see Visual Difference

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
Building a SELECT Statement	Create, manipulate and execute SELECT Statements for tables and views.
Building an INSERT Statement	Create and manipulate INSERT Statements for tables.
Building an UPDATE Statement	Create and manipulate UPDATE Statements for tables.
Building a DELETE Statement	Create and manipulate DELETE Statements for tables.
Building a CREATE VIEW Statement	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 DBArtisan. 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:

- Statement Properties
- Workspace Windows
- Query Builder Explorer Window
- Query Builder Tool Bar
- SQL Diagram Pane
- SQL Statement Pane

For more information, see Query Builder

Workspace Windows

The Workspace Windows provide a comprehensive view of your data structure and query. The table below describes the Workspace Windows:

Pane	Description
Query Builder Explorer Window	Includes two tabs that display selected object details: Tables/Views DML
SQL Diagram Pane	Displays tables or views included in the current query.
SQL Statement Pane	Displays the SQL code, and when appropriate, a Results Tab.

For more information, see Query Builder

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 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. For details, see Saving and Reopening Queries.

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

Query Builder

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:

- Working with Tables and Views in the SQL Diagram Pane
- Joins
- Working with Columns in the SQL Diagram Pane

All changes in the SQL diagram reflect in correct SQL code in the SQL Statement Pane.

For more information, see Query Builder

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

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. DBArtisan 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 details, see Using Data Editor with Query Builder.

>

For more information, see Query Builder

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 Query Builder.

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	For more information, see Table Properties.

NOTE: Your selection applies to all selected tables and views.

For more information, see Query Builder

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 Query Builder

Query Builder Dialog Boxes

Query Builder includes a number of dialog boxes to assist you in building and customizing your query.

Dialog Box	Description
Statement Properties	Specifies general properties in an individual Query Builder session.
Table Properties	Specifies column selection and alias names for a table or view.
Column Properties	Specifies column functionality within SELECT and CREATE VIEW statements.

For more information, see Query Builder

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. These properties override the global options set in the Options editor for the current session. For details, see Query Builder options.

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 in the Options editor while Query Builder is running, DBArtisan displays a warning indicating that you are about to change options or properties. For details, see Query Builder options.

>

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**.
2. Set options.
3. Click **OK**.

   Query Builder saves the options.

For more information, see Query Builder

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.

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 Query Builder

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.
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 Query Builder

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 for Linux, Unix, and Windows and IBM DB2 for z/OS and OS/390 servers, there is an additional join object in the SQL Statement Tree. The Join On node displays join relations between columns in IBM DB2 for Linux, Unix, and Windows and IBM DB2 for z/OS and OS/390 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 Query Builder

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 DBArtisan. Query Builder lets you build five separate types of queries simultaneously:

- Building a SELECT Statement
- Building an INSERT Statement
- Building an UPDATE Statement
- Building a DELETE Statement
- Building a CREATE VIEW Statement

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. Saving and Reopening Queries.

DBArtisan 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 DBArtisan. You can use different tables and views for each query. You can also toggle back and forth among each of the queries. For more information, see Working with Tables and Views in the SQL Diagram Pane.

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. For details, see Saving and Reopening Queries.

Using Query Builder

To use Query Builder, do the following:

- Selecting a Database
- Selecting a Statement
- Selecting Tables and Views
- Selecting 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**.

DBArtisan 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, IBM DB2 for z/OS and OS/390 or IBM DB2 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.

DBArtisan clears the current query and displays a warning prompt.

>

5 To save the current query, click **Yes**.

DBArtisan opens the Save As dialog box.

6 To continue without saving, click **No**.

DBArtisan clears the SQL Diagram Pane and SQL Statement Pane.

For more information on saving queries in Query Builder, seeSaving and Reopening Queriess.

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. For details, see 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. For details, see:

- SQL Diagram Pane.
- Selecting Columns in the 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 details, see Column Properties.

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**.

For more information, see Using Query Builder.

Building a Query

Query Builder lets you build SELECT, INSERT, UPDATE, SELECT, and CREATE VIEW statements, which you can run separately or simultaneously, depending on your needs.

To build a Query, do the following:

- Selecting a Database
- Selecting a Statement
- Selecting Tables and Views
- Selecting 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

Selecting 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.

For related information, see Joins and Syntax Checker.

To build a SELECT statement, do the following:

1. On the **Tools** menu, click **Query Builder**.

   DBArtisan 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 Query Builder

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. For details, see Saving and Reopening Queries.

Building an INSERT Statement

To build an INSERT Statement, do the following:

1. On the **Tools** menu, click **Query Builder**.

   DBArtisan 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 Query Builder

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. For details, see Saving and Reopening Queries.

Building an UPDATE Statement

To build an UPDATE statement, do the following:

1. On the **Tools** menu, click **Query Builder**.

   DBArtisan 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 Query Builder

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. For details, see Statement Properties.

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. For details, see Saving and Reopening Queries.

Building a DELETE Statement

To build a DELETE statement, do the following:

1 On the **Tools** menu, click **Query Builder**.

DBArtisan 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 Query Builder

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. For details, see Saving and Reopening Queries.

To build a CREATE VIEW statement, do the following:

1 On the **Tools** menu, click **Query Builder**.

```
DBArtisan 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.

DBArtisan 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**.

DBArtisan executes the CREATE VIEW query.

11 To close the Editor, click **Close**.

DBArtisan opens the ISQL Editor save message.

12 Click **No**.

DBArtisan returns to Query Builder.

13 To add the view to the table tree, on the **Query Builder** menu, click **Refresh**.

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 Query Builder

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 a Table or View

Query Builder can automatically dictate a layout in the SQL Diagram Pane using the **Auto Layout** button. For details, see Auto Layout.

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 in the SQL Diagram Pane.

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 in the SQL Diagram Pane.

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 in the SQL Diagram Pane.

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 in the SQL Diagram Pane.

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. For details, see Column Properties.

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. For more information, see Selecting 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**.

> **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**.

   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
Inner Joins	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.
Left Outer Joins	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.
Right Outer Joins	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.
Self Joins	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. For details, see Editing Joins.

You can set Query Builder options, in the Options editor, to automatically create joins. For details, see Query Builder options.

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.

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 Left Outer Joins. 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.
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 for Linux, Unix, and Windows and IBM DB2 for z/OS and OS/390 servers, there is an additional join object in the SQL Statement Tree. The Join On node displays join relations between columns in IBM DB2 for Linux, Unix, and Windows and IBM DB2 for z/OS and OS/390 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. For details, see Auto Joins.

Using Auto Layout

To use Auto Layout, do the following:

1 On the **Query Builder** menu, 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 DBArtisan 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.For details, see Statement Properties.

Using Auto Join

To use Auto Join, do the following:

1 On the **Query Builder** menu, 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.

For more information, see the following topics:

- Creating a WHERE Clause
- Deleting a WHERE Clause
- Creating an AND Clause in a WHERE Clause
- Deleting an AND Clause
- Inserting an AND or OR Clause
- Deleting an OR Clause
- Creating an ORDER BY Clause
- Changing the Sort Order in an ORDER BY Clause
- Deleting an ORDER BY Clause
- Creating a GROUP BY Clause
- Deleting a GROUP BY Clause
- Creating a HAVING Clause
- Deleting a HAVING Clause

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 Selecting Columns in 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 Selecting Columns in 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 Selecting Columns in 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 Selecting Columns in 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 Selecting Columns in 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 Selecting Columns in 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 Selecting Columns in 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 Selecting Columns in 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 Selecting Columns in 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 Selecting Columns in 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 Selecting Columns in 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 Selecting Columns in 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 Selecting Columns in 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 Selecting Columns in 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 Selecting Columns in 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. The options for the Syntax Checker tool are set in the Options Editor. For details,see Query Builder options.

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 HAVING clause, is there a GROUP BY clause?	Query Builder returns an error message when it detects a HAVING clause without a GROUP BY clause.
If there are aggregates, or a GROUP BY 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 joins 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 DBArtisan.

Query Builder runs Schema Change detection any time you set a query to execute, refresh the data, or open a saved query.

DBArtisan 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**.

DBArtisan 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: DBArtisan lets you save data in text (*.txt) and XML (*.xml) file formats.

3 Click **OK**.

DBArtisan 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. DBArtisan 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**.

DBArtisan opens Query Builder.

2 Select a database or instance. For details, see Selecting a Database.

3 Select a table. For details, see Selecting Tables and Views.

4 Select a column, or columns. For details, see Selecting Columns.

5 On the **Query Builder** menu, click **Edit Data**.

DBArtisan opens Data Editor.

For more information on using the Data Editor, seeData Editor

Import Data

It is often necessary to import data into database tables from external sources. You may need to bring sample data into a database to establish test case scenarios or import data for use in the real world. Business analysts, for example, often use spreadsheets from outside sources. Putting such data into tables can help them perform detailed queries and make informed decisions based on meaningful data.

NOTE: When you import a text file, the wizard may give the impression of stalling. Press **ESC** and you should be able to continue without further incident.

To Open the Import Data Wizard:

1 On the **Tools** menu, click **Import Data**.

The fields you may encounter as you move through the wizard are described in the table that follows:

Data Properties

Required Information	Description
Specify the file to be used in this data load operation.	Type or browse to the file you will be importing the data from
Which table do you want to load data into? Catalog Schema Table	Catalog: The collection of tables for Sybase or MS SQL. This is where you want to import the data to. Schema: The user. Table: The particular table where you want to import the data.

Delimeter

Required Information	Description
What character delimited the columns in the data file?	Choose between Tab, Semicolon, Comma, Space, or Other. If you choose Other, you must indicate the specific delimiting character.
First Row Contains Field Names	By choosing this, if the first row contains column names, these will be displayed as the first row of the file. You can still change column mappings if you want. If you do not select this, generic Field 1, Field 2, etc. will indicate column mappings.

Column Mapping

Required Information	Description
NOTE: You must map all mandatory (NOT NULL) column fields to a column heading before you can enable the Finish button and import the data. To map columns, click the grey bar at the top of the grid and a drop-down list appears. The names in this drop-down list are columns in the table that's going to receive the data. These need to sync up to the data that's being imported. Red, or mandatory columns, turn blue in the column list when they have been assigned. A maximum of 10 rows are displayed here to make mapping columns easier. All rows will be imported when the job completes. If there are more columns in the file you are importing than in the receiving table, the extra columns will be ignored and you can create the insert statements. You cannot complete the import operation if columns are mapped more than once. The import operation will not let you generate invalid columns.	

Excel Column Mapping

Required Information	Description
Sheet	If the file you selected has more than one sheet, identify the worksheet that has the data you want to import.
Start/End Cell	Identify the starting and ending cells of data you want to import and click **Refresh**. Or, accept the default cells.
First Row Contains Field Names	By choosing this, if the first row contains column names, these will be displayed as the first row of the file. You can still change column mappings if you want. If you do not select this, generic Field 1, Field 2, etc. will indicate column mappings.
NOTE: You must map all mandatory (NOT NULL) column fields to a column heading before you can enable the Finish button and import the data. To map columns, click the grey bar at the top of the grid and a drop-down list appears. The names in this drop-down list are columns in the table that's going to receive the data. These need to sync up to the data that's being imported. Red, or mandatory columns, turn blue in the column list when they have been assigned. If there are more columns in the file you are importing than in the receiving table, the extra columns will be ignored and you can create the insert statements. You cannot complete the import operation if columns are mapped more than once. The import operation will not let you generate invalid columns.	

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. For details, see Query Builder.

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 Filter.

TIP: You can customize Data Editor options in the Options editor. For detals, see Data Editor options.

For more information, see:

- Data Editor Design
- Using Data Editor

Data Editor Design

The Data Editor includes the following components:

- Data Editor Edit Window
- Data Editor ISQL Window
- Data Editor Tool Bar

>

- Notes on XML types and Unicode display in the Data editor
- Data Editor Filter
- Editing Date and Time Functions
- 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:

Data Editor ISQL Window

Data Editor Tool Bar

Data Editor Filter

Editing Date and Time Functions

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:

Data Editor Edit Window

Data Editor Tool Bar

Data Editor Filter

Editing Date and Time Functions

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.

>

Description	Function
Insert Record button	Inserts new record for the target table. New records display at the end of the table. For related information, see Default value handling.
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 Default value handling.
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
Calendar button	Sets correct format for target date/time cell. Enables the Calendar window.
Date/Time Format Builder button	For details, see Editing Date and Time Functions.
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:

Data Editor Edit Window

Data Editor ISQL Window

Data Editor Filter

Editing Date and Time Functions

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:

Data Editor Edit Window

Data Editor ISQL Window

Data Editor Filter

Editing Date and Time Functions

Using Data Editor

Extract Data as XML

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 Uincode 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. For details, see Editing Date and Time Functions.

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:

Data Editor Edit Window

Data Editor ISQL Window

Data Editor Filter

Editing Date and Time Functions

Using Data Editor

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:

Data Editor Tool Bar

Editing Date and Time Functions

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 DBArtisan, 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, DBArtisan 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
Global	Lets you make global changes to the Data Editor date display from the Data Editor Tab of the Options Editor. For details, see Editing Date/Time Globally.	Options Editor
Grid	Lets you make changes to the date display of the entire Data Editor grid for that session only. For details, see Editing Grid Date/Time.	Data Editor grid
Column	Lets you make changes to the date display of a single column in the Data Editor for that session only. For details, see Editing Column Date/Time.	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 or Editing Column Date/Time.

To edit the date and time globally, do the following:

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

DBArtisan opens the Options Editor.

2 On the **Options Editor**, click the **Data Editor Tab**.

3 On the **Data Editor Tab**, click **...**

DBArtisan 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**.

DBArtisan 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**.

DBArtisan 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, DBArtisan uses the **Calendar Control** window. For details, see Editing Date and Time Functions.
Two-digit year system setting warning	If selected, DBArtisan sends a warning when you use a two-digit year system setting.

9 Click **OK**.

DBArtisan 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. For details, see Editing Date/Time Globally.

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**.

DBArtisan opens the Data Editor.

3 On the **Data Editor** tool bar, click **Date/Time Format Builder**.

DBArtisan 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**.

DBArtisan 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**.

DBArtisan 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. For details, see Editing Grid Date/Time.

To make this display permanent, Editing Global Date/Time Format. For details, see Editing Date/Time Globally.

To edit the column date and time, do the following:

1 On the **Datasource Explorer**, select the target table.

2 Right-click the table, and click **Edit Data**.

DBArtisan opens the Data Editor.

3 On the **Data Editor**, click the column header to select the column.

4 Right-click the column and click **Format**.

DBArtisan opens the Date/Time Format Builder.

5 On the **Date/Time Format Builder** dialog box, click the **Date/Time Format** list, and then click the target predefined date/time format.

6 To customize the date/time format to your specifications, click **Customize**.

DBArtisan 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**.

DBArtisan 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.

Embarcadero Products

The Tools menu lists all installed Embarcadero Technologies products. This lets you toggle to or start another Embarcadero product.

Performance Center

If you have Performance Center installed on a server, DBArtisan lets you view performance metrics across all of your Performance Center monitored datasources.

NOTE: This feature requires only a Performance Center server definition. No client upgrade is required. For details, see Perf Center options.

Performance Center offers an automated strategy for identifying problems that threaten the availability and performance of critical Oracle®, Sybase and Microsoft SQL Server databases. Embarcadero Performance Center ensures that relational database management systems enjoy maximum up-time and exhibit performance that exceeds user expectations. It quickly identifies performance problems, and helps prioritize performance-tuning workloads to keep your databases performing at peak efficiency.

TIP: On the **Perf Center** Tab of the Options Editor, you can specify your Performance Center server definition and whether the DBArtisan opens the standard client (PerfCntr.exe) or the Web Client URL.

For more information, Using the Performance Center Web Client Within DBArtisan.

Using the Performance Center Web Client Within DBArtisan

To open the Performance Center web client, do the following:

1 On the **Tools** toolbar, select **Performance Center**.

OR

On the **Tools** menu, select **Performance Center**.

The first time you launch the Performance Center Web Client from DBArtisan, DBArtisan opens the Select Performance Center Server dialog box.

2 In **Computer Name or IP Address** specify where the Performance Center Server is installed.

3 In **Port Number** specify the port for the Apache or IIS server.

4 To verify the settings, click **Test**.

DBArtisan displays the Performance Center Server Test Results dialog box.

5 Click **Close**.

DBArtisan returns to the Select Performance Center Server dialog box.

6 To hold these setting for future use, click **OK**.

DBArtisan opens Performance Center.

Code Workbench

The Code Workbench tool reduces the time needed for the most common day-to-day coding tasks. The Code Workbench lets you:

- Define auto replacement expressions that can be used to quickly insert commonly used blocks of SQL syntax or commands.
- Import and Export Code Workbench settings for client sharing purposes.

NOTE: No privileges are required to use the Code Workbench options.

Important Notes

When installing DBArtisan on a machine that already has Rapid SQL 7.3 or DBArtisan 7.3 or later, the Code Workbench settings are not installed by default. This is to protect any modifications that have been made to the current settings. To import the settings that are shipped with the latest release, select Tools, Code Workbench. Select the Import Settings option, navigate to the main DBArtisan installation folder, and select codeworkbench.xml.

CAUTION: This will overwrite the current settings.

The table below describes the options and functionality on the Code Workbench dialog box.

Tab	Option	Description
Settings	Enable Auto Replace	When enabled, you can define simple keystrokes that will be automatically replaced with predefined expressions when some activation keyboard event is performed (space, new line, special character, etc.). This also lets you define correction actions for common misspellings or SQL syntax coding errors.
	Replace substrings	When selected, this lets you enable automatic replacement for substrings: simple keystrokes will be automatically replaced with predefined expressions when some activation keyboard event is performed (space, new line, special character, etc.). This also lets you define correction actions for common misspellings or SQL syntax coding errors.
Auto Replace	Add/Edit/Clone	Click to open the Edit Auto Replace Expression dialog box. For details, see Using Code Workbench Auto Replace Shortcuts in the ISQL Window.
	Delete	Click to delete the selected auto replace expression.
	Import Settings	Click to import settings from a local or network directory.
	Export Settings	Click to export your settings to a local or network directory.
	Restore Settings	Click to restore the Code Workbench settings to the original settings.

Common Tasks

Creating and Modifying Code Workbench Auto Replace Shortcuts

Using Code Workbench Auto Replace Shortcuts in the ISQL Window

Importing and Exporting Settings in Code Workbench

Common Tasks

Creating and Modifying Code Workbench Auto Replace Shortcuts

Using Code Workbench Auto Replace Shortcuts in the ISQL Window

Importing and Exporting Settings in Code Workbench

Creating and Modifying Code Workbench Auto Replace Shortcuts

The Edit Auto Replace Expression dialog box lets you define shortcut expressions for commonly used SQL syntax or commands. You can define what keys or key combination activate the auto replacement of the defined shortcut.

1. Select **Tools**, **Code Workbench**.
2. On the **Auto Replace Tab**, select **Add**.
3. In **Expression** type a new expression.
4. In **Activation** type activation information. For example, CTRL+SPACE.
5. In **Replace With** type replace with information.
6. Click **OK**.

   The replace expression is now ready for use in the ISQL Window.

>

For more information, see:

Using Code Workbench Auto Replace Shortcuts in the ISQL Window

Code Workbench

Using Code Workbench Auto Replace Shortcuts in the ISQL Window

When Enable Auto Replacement is selected on the Setting Tab of the Code Workbench, you can type a replace expression in the ISQL Window instead of typing a complete command.

1 Select **File**, **New**, and then **SQL**.

2 In the **ISQL Window**, type a replace expression. For example, type **beg**.

The application displays begin.

For more information, see Code Workbench.

Importing and Exporting Settings in Code Workbench

The feature helps standardize your application settings. Exporting settings is useful when you want to give your auto replace expressions to another ISQL Window user. Importing settings is useful when another user wants to give you their auto replace expressions.

Importing Settings

1 Select **Tools**, **Code Workbench**.

2 Click **Import Settings**.

Code Workbench opens the Open dialog box.

3 Locate the .xml file, and then click **Open**.

Code Workbench imports the settings.

Exporting Settings

1 Select **Tools**, **Code Workbench**.

2 Click **Export Settings**.

Code Workbench opens the Save As dialog box.

3 Specify a location, and file name, and then click **Save**.

Code Workbench saves your settings as an .xml file. You can send the .xml file to the user you want to import your settings.

For more information, see Code Workbench.

Permissions Management

DBArtisan 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. DBArtisan 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. DBArtisan 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. DBArtisan 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. DBArtisan lets you revoke permissions with the Roles Editor.

DBArtisan lets you identify existing users and their database permissions, and to detect and manage changes to user permissions by rolling change back or forward. DBArtisan 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, DBArtisan gives you:

- Enhanced permissions management across the enterprise (Oracle, Sybase ASE, Microsoft SQL Server, and IBM DB2 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. DBArtisan 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. DBArtisan also shows if the privilege is grantable (granted with the GRANT OPTION.) The table below describes the icons:

Icon	Label	Description
Black check mark	User	The privilege(s) was granted explicitly by a user.
Blue check mark	Group	The privilege(s) was inherited from a group.
Green check mark	Role	The privilege(s) was inherited from a role.
Red check mark	Denied	**MICROSOFT SQL SERVER ONLY:** The privilege(s) was denied from a security account in the current database.
Two black check marks	User/Grant	The privilege(s) was granted by a user with the GRANTABLE option.
Two blue check marks	Group/Grant	The privilege(s) was granted from a group with the GRANTABLE option.
Two green check marks	Role/Grant	The privilege(s) was granted from a role with the GRANTABLE option.

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.

>

also shows if the privilege is grantable (granted with the GRANT OPTION.) The table below describes the icons:

Icon	Label	Description
Black check mark	User	The privilege(s) was granted explicitly by a user.
Blue check mark	Group	The privilege(s) was inherited from a group.
Green check mark	Role	The privilege(s) was inherited from a role.
Red check mark	Denied	**MICROSOFT SQL SERVER ONLY:** The privilege(s) was denied from a security account in the current database.
Two black check marks	User/Grant	The privilege(s) was granted by a user with the GRANTABLE option.
Two blue check marks	Group/Grant	The privilege(s) was granted from a group with the GRANTABLE option.
Two green check marks	Role/Grant	The privilege(s) was granted from a role with the GRANTABLE option.

NOTE: The available privileges depend on the target DBMS and object.

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.

Add-On Tools

Embarcadero Technologies offers add-on tools that enhance the usefulness of DBArtisan. The table below describes these add-on tools:

Add-on Tool	Description
Embarcadero SQL Debugger	Provides detailed information about using the Embarcadero SQL Debugger. It Includes a step-by-step walk through to familiarize you with the features of the Embarcadero SQL Debugger.
DBArtisan PL/SQL Profiler	Provides detailed information about using the DBArtisan PL/SQL Profiler. Includes a step-by-step walk through to familiarize you with the features of the PL/SQL Profiler.

Embarcadero SQL Debugger

The Embarcadero SQL Debugger lets you locate and fix bugs in the following elements:

- Procedures (IBM DB2 for Linux, Unix, and Widows, Microsoft SQL Server, Oracle, Sybase)
- Functions (Oracle)
- Triggers (IBM DB2 for Linux, Unix, and Widows, Microsoft SQL Server, Oracle, Sybase)

NOTE: The SQL Debugger lets you debug triggers by debugging the procedures that call them.

NOTE: For Oracle, you cannot debug packages, but you can debug the functions and procedures within packages.

NOTE: For Oracle, you cannot debug any objects contained in the Exclusion List. For more information, see Editing the Exclusion List.

The table below describes the sections of this chapter:

Section	Description
SQL Debugger modules and DBMS support	Describes the specific SQL Debugger modules and DBMS versions supported/
Embarcadero SQL Debugger features	Provides an overview of SQL Debugger functionality.
Setting up the Embarcadero SQL Debugger	Provides details on requirements and setup tasks for each supported DBMS.
Embarcadero SQL Debugger interface	Describes the Embarcadero SQL Debugger graphical interface that includes an editor window and four debug view windows.
Embarcadero SQL Debugger functionality	This section describes the functionality on the SQL Debugger.
Using the Embarcadero SQL Debugger	This section describes how to run a debug session.
Tutorial sessions	Provides walkthrough sessions for new users.

SQL Debugger modules and DBMS support

There are four Embarcadero SQL Debugger modules, each corresponding to a supported DBMS. Modules and DBMS support are as follows:

- **Embarcadero SQL Debugger for IBM DB2 for Linux, Unix, and Windows** - supports IBM DB2 for Linux, Unix, and Windows version 8 or later
- **Embarcadero SQL Debugger for Microsoft** - supports Microsoft SQL Server version 7.0 or later
- **Embarcadero SQL Debugger for Oracle** - supports all Oracle versions supported by DBArtisan
- **Embarcadero SQL Debugger for Sybase** - supports Sybase 12 , 12.0 (special version), and 12.5

Each SQL Debugger version is an optional and separate add-on module**.**

Embarcadero SQL Debugger features

The Embarcadero SQL Debugger lets you identify problems within your code. The Embarcadero SQL Debugger 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 DBArtisan main menu, the Procedures window, the DDL Editor and ISQL windows.

The Embarcadero SQL Debugger offers fundamental debugging features and options to fine tune debugging. The table below describes these features:

Debugging Feature	Description
Step Into	Lets you execute each instruction step-by-step and step inside a stored object.
Step Out	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.
Step Over	Lets you execute the current instruction without stepping into any child dependents.
Insert or Remove Breakpoint	Lets you specify positions in a program where the debugger stops execution.

To set specific Debugger values on DBArtisan's Options Editor, see Embarcadero SQL Debugger options.

Setting up the Embarcadero SQL Debugger

The following topics describe requirements and tasks for those DBMS that have setup requirements:

- Embarcadero SQL Debugger for Linux, Unix, and Windows setup
- The Embarcadero SQL Debugger for Microsoft setup
- Embarcadero SQL Debugger for Oracle setup

In addition, you should set your debugger preferences at this time. For details, see Embarcadero SQL Debugger options.

Embarcadero SQL Debugger for Linux, Unix, and Windows setup

Embarcadero SQL Debugger for IBM DB2 for Linux, Unix, and Windows requires proper configuration of the server and client.

For more information, see:

- Prerequisites for Adding and Compiling Stored Procedures
- Configuring the IBM DB2 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 for Linux, Unix, and Windows requires the following products and components.

Client

- IBM DB2 for Linux, Unix, and Windows 8 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 for Windows/Unix Support
- Administration and Configuration Tools

Server

- IBM DB2 for Linux, Unix, and Windows 8 or later
- DB2 Enterprise Edition
- Communications Protocols
- Stored Procedure Builder
- Applications Development Interfaces
- System Bind Files
- DB2 Connect Server Support
- Documentation
- Base DB2 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 8 and includes the C compiler options needed to compile the procedure on the server. If it is not found, install the IBM DB2 8 Administration and Configuration Tools option on the server.

Configuring the IBM DB2 for Linux, Unix, and Windows Server for Procedures

DBArtisan lets you create procedures on the targeted server using DBArtisan.

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 SQL Debugger.

The Embarcadero SQL Debugger for Microsoft setup

Embarcadero SQL Debugger for Microsoft requires proper configuration of Server Requirements and Client Requirements.

Important Note about Microsoft SQL Server 2000 Service Pack 4 (SP4)

When running Windows XP SP2, SQL Server 2000 SP4 is required on both the client and the server. To verify this is the case start MS Query Analyzer and press Help > About. The Version SQL should be 8.00.2039.

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 DBArtisan's T-SQL Debugger.

You may receive the following error message if you try to use the T-SQL Debugger from DBArtisan:

> "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 has disabled T-SQL Debugger for Application using earlier T-SQL Debugger clients for security reasons. To enable T-SQL Debugger for these Applications, a member of the sysadmins server role, such as sa must explicitly enable debugging by running the following code:

```
Exec sp_sdidebug 'legacy_on'
```

NOTE: You must repeat this procedure whenever you restart the server.

Server Requirements

Embarcadero SQL Debugger for Microsoft requires:

- Windows 2005
- Windows 2000
- Microsoft SQL Server version 7.0 or later

>

Setting Up the Server

There are three parts to setting up the server:

1. Installing the Microsoft SQL Debugger Interface Subcomponent
2. Configuring the Service
3. 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. 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, install them before running the Embarcadero SQL Debugger for Microsoft.

To install the Debugger Interface subcomponent on the server after the initial installation:

1. Start **Microsoft Visual Studio, Enterprise Edition Setup**.

   OR

   Start **Microsoft SQL Server 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.

For more information, see Server Requirements.

Configuring the Service

To configure the service, see the instructions for your server operating system:

- Windows 2000

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 who will be using the Embarcadero SQL Debugger for Microsoft.

NOTE: This person needs 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 Note about Microsoft SQL Server 2000 Service Pack 4 (SP4)

When running Windows XP SP2, SQL Server 2000 SP4 is required on both the client and the server. When running Windows XP SP2, SQL Server 2000 SP4 is required on both the client and the server. To verify this is the case start MS Query Analyzer and press Help > About. The Version SQL should be 8.00.2039.

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 DBArtisan's T-SQL Debugger.

You may receive the following error message if you try to use the T-SQL Debugger from DBArtisan:

> "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 has disabled T-SQL Debugger for Application using earlier T-SQL Debugger clients for security reasons. To enable T-SQL Debugger for these Applications, a member of the sysadmins server role, such as sa must explicitly enable debugging by running the following code:

```
Exec sp_sdidebug 'legacy_on'
```

NOTE: You must repeat this procedure whenever you restart the server.

For more information, see Configuring the Service.

Configuring DCOM on the Server

To configure DCOM on the server, 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, grant them remote access on the server. To grant remote access, 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.

For more information, see Server Requirements.

Client Requirements

There are three categories of client requirements for the Embarcadero SQL Debugger for Microsoft:

- Operating System
- Client Connectivity
- Installing the Microsoft SQL Debugger Interface Subcomponent

Operating System

The client must be running one of the following operating systems:

- Microsoft Windows 95
- Microsoft Windows 98
- Microsoft Windows XP

For more information, see Client Requirements.

Client Connectivity

When running the Debugger on SQL Server 2005 Clients

Before you proceed, verify that the file ssdebugps.dll is registered on the client machine. This file is REQUIRED for debugging and Microsoft only installs it on the server machine and not as a part of a client-only install.

Next you must configure DCOM by following these steps:

1 At a command prompt, type **dcomcnfg**, and then press **ENTER**.

Component Services opens.

2 In Component Services, expand **Component Services**, expand **Computers**, and then expand **My Computer**.

3 On the toolbar, click the **Configure My Computer** button.

The **My Computer** dialog box appears.

4 In the **My Computer** dialog box, click the **COM Security** tab.

5 Under **Access Permission**, click **Edit Limits**.

The **Access Permission** dialog box appears.

6 Under **Group or user names**, click **ANONYMOUS LOGON**.

7 Under **Permissions** for ANONYMOUS LOGON, select the **Remote Access** check box, and then click **OK**.

Finally, you need to go to the **Default Properties** tab of the MyComputer dialog.

1 Select **None** in the **Default Authentication Level** dropdown menu.

2 Select **Impersonate** in the **Default Impersonation Level** dropdown menu.

For more detailed information, you can refer to Microsoft's Help:
:http://support.microsoft.com/default.aspx?kbid=833977#XSLTH413412112412012112012O

When running Windows XP SP2, SQL Server 2000 SP4 is required on both the client and the server.

For Microsoft SQL the client must have the Client Connectivity component of Microsoft SQL Server.

For Microsoft SQL Server

The client must have the Development Tools, Debugger Interface subcomponent of Microsoft SQL Server. 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, 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, do the following:

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.

For more information, see Client Requirements.

Embarcadero SQL Debugger for Oracle setup

The only task required in setting up the Embarcadero SQL Debugger for Oracle is Editing the Exclusion List.

Editing the Exclusion List

Upon installation, DBArtisan sets up an Exclusion List on your computer which includes packages that the application cannot debug. The Exclusion List is located in the DBArtisan directory. You can add or remove packages from this file by 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 Microsoft 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: Press ENTER 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 not debug a package function or package procedure listed on the Exclusion List.

4. Save the changes to **deborcex.etd**.

Embarcadero SQL Debugger options

You can specify debugger options from the Debug Tab of DBArtisan'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**.

   DBArtisan 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 DBArtisan 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 dbms_output.put_line 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**.

DBArtisan closes the Options Editor.

For more information, see Embarcadero SQL Debugger interface and Embarcadero SQL Debugger functionality.

Embarcadero SQL Debugger interface

The Embarcadero SQL Debugger includes an editor window and four debug view windows. When you open a debug session, DBArtisan 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 includes five windows:

1 DDL Editor Window

2 Watch Window

3 Variables Window

4 Call Stack Window

5 Dependency Tree Window

Working with SQL Debugger Windows

DBArtisan 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.

DBArtisan resizes the window.

2 To move and dock the target window, click its grab bar and drag it.

DBArtisan 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.

DBArtisan frames the window in its own floating frame and moves the window to its new location.

DDL Editor Window

The DDL Editor displays your code in read-only format. When you start debugging, the Embarcadero SQL Debugger extracts your code into a DDL Editor. The DDL Editor uses the default DBArtisan syntax coloring.

For more information, see Embarcadero SQL Debugger interface.

> **NOTE:** For Oracle, DBArtisan displays LOB datatypes, and REF CURSOR variables, in the Results Tab.

Watch Window

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 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: 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 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 @. Drag the @ to the Watch Window.

2 In the **Watch** window**,** change the value of the variable.

3 On the **DDL Editor**, click **Debug** or **Go**.

The Embarcadero SQL Debugger 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 interface.

Variables Window

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 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.

DBArtisan opens a ScreenTip displaying the current value of that variable.

For more information, see Embarcadero SQL Debugger interface.

Call Stack Window

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, DBArtisan displays a green arrow on the line of the referenced call.

For more information, see Embarcadero SQL Debugger interface.

Dependency Tree Window

The Dependency Tree window displays any external database objects the script accesses. DBArtisan 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.

DBArtisan displays the SQL of the target object in the DDL Editor window.

For more information, see Embarcadero SQL Debugger interface.

Embarcadero SQL Debugger functionality

The Embarcadero SQL Debugger offers 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

You must be running a debugging session to use these functions. For details, see Opening 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, DBArtisan opens a dialog box to prompt 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 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 or Continue button	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 functionality.

Step Into

Step Into lets you execute the current instruction. If the current instruction makes a call to a stored SQL object, the Embarcadero SQL Debugger steps inside the nested child object.

To use the Step Into facility, do the following:

1 On the **Debug** menu, click **Step Into**.

OR

Press **F11**.

The Embarcadero SQL Debugger moves the arrow to execute the current instruction.

For more information, see Embarcadero SQL Debugger functionality.

Step Out

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

Press **SHIFT+F11**.

The Embarcadero SQL Debugger stops stepping through the current object and executes the remainder of the script.

For more information, see Embarcadero SQL Debugger functionality.

Step Over

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

   Press **F10**.

   The Embarcadero SQL Debugger executes the current instruction.

For more information, see Embarcadero SQL Debugger functionality.

Run to Cursor

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 places the cursor on the target line.
3. On the **Debug** menu, click **Run to Cursor**.

   OR

   Press **CTRL+F10**.

   The Embarcadero SQL Debugger executes all instructions between the pointer and the cursor.

For more information, see Embarcadero SQL Debugger functionality.

Insert or Remove Breakpoint

A breakpoint is a position in a program where a debugger stops execution. When you start debugging, Embarcadero SQL Debugger opens the script in a DDL Editor. A yellow arrow pointer indicates which line the Embarcadero SQL Debugger executes next.

The Embarcadero SQL Debugger executes all lines of code between the yellow arrow and the first breakpoint. If no breakpoints are present, Embarcadero SQL Debugger 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 Breakpoints, temporarily disable or enable breakpoints without having to add or remove breakpoints.

DBArtisan's Embarcadero SQL Debugger 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.

DBArtisan stores all breakpoints you set so that when you debug the same script on separate occasions, you can reuse the same breakpoints. 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

   Press **F9**.

   The Embarcadero SQL Debugger inserts a new breakpoint or removes an existing breakpoint on the target line of code.

For more information, see Embarcadero SQL Debugger functionality.

Toggle Breakpoint

After inserting a breakpoint, Toggle Breakpoint lets you enable or disable that breakpoint. Embarcadero SQL Debugger 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 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

   Press **CTRL+F9**.

   The Embarcadero SQL Debugger toggles the breakpoint indicated by the pointer.

For more information, see Insert or Remove Breakpoint.

Go

Go lets you execute all instructions stopping only 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

   Press **F5**.

   The Embarcadero SQL Debugger executes all instructions.

For more information, see Embarcadero SQL Debugger functionality.

Stop

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

Press **SHIFT+F5**.

The Embarcadero SQL Debugger stops the script execution and terminates the session.

For more information, see Embarcadero SQL Debugger functionality.

Restart

Restart lets you terminate the current debug session and open a new one. When the new session opens, DBArtisan prompts you for new input parameters.

To use the Restart facility, do the following:

1 On the **Debug** menu, click **Restart**.

OR

Press **CTRL+SHIFT+F5**.

The Embarcadero SQL Debugger restarts the debug session.

For more information, see Embarcadero SQL Debugger functionality.

Break

Break lets you pause the debug session.

To use the Break facility, do the following:

1 On the **Debug** menu, click **Break**.

The Embarcadero SQL Debugger suspends the debug session.

For more information, see Embarcadero SQL Debugger functionality.

Close

Close lets you close the DDL Editor and the Embarcadero SQL Debugger.

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 closes the debug session.

>

For more information, see Embarcadero SQL Debugger functionality

Using the Embarcadero SQL Debugger

This section offers a general overview of how to use Embarcadero SQL Debugger's full range of debugging functionality.

- Opening a Debugging session
- Debugging an SQL Script with Embarcadero SQL Debugger

Opening a Debugging session

When you open a debugging session, DBArtisan opens the Embarcadero SQL Debugger interface. For details, see Embarcadero SQL Debugger interface.

If the target script requires input parameters, DBArtisan opens a dialog box prompting you for the necessary input parameters before displaying the target code in the SQL Editor window. When DBArtisan displays the target script in the SQL Editor window, you can begin debugging. For details, see Debugging an SQL Script with Embarcadero SQL Debugger

NOTE: Embarcadero SQL Debugger for IBM DB2 for Linux, Unix, and Windows, mbarcadero SQL Debugger for Sybase, and Embarcadero SQL Debugger for Microsoft only let you debug the SQL script of procedures or triggers.

To debug an object, do the following:

1 In the **Datasource Explorer**, click the object node.

DBArtisan displays the objects in the right pane of the application.

- Right-click the target object, and select **Open**.

DBArtisan opens the Object Editor.

OR

- Right-click the target object and select **Extract**.

DBArtisan opens the DDL Editor displaying the code of the target object.

2 On the toolbar, click **Debug**.

OR

In the **DDL Editor** window, right-click, and then click **Debug**.

If the script requests input parameters, DBArtisan opens a dialog box prompting for parameter values. If the script does not require input parameters, DBArtisan displays the script in the DDL Editor window for you to begin debugging. For details, see Debugging an SQL Script with Embarcadero SQL Debugger.

NOTE: You cannot use the Embarcadero SQL Debugger until it has fully initialized.

3 In the dialog box, specify the appropriate parameter valuess, and then click **OK** or **Continue**.

DBArtisan displays the script in the DDL Editor window.

NOTE: If the script requires Sybase or Oracle types (tables, records, or Booleans) as input parameters, the Embarcadero SQL Debugger generates an anonymous block.

Debugging an SQL Script with Embarcadero SQL Debugger

After you open a debugging session and enter any required input parameters, you can begin working with your script in the Embarcadero SQL Debugger.

Debugging an SQL Script

To debug a SQL Script, do the following:

1. On the **Debug** menu, click one of the SQL Debugger options (Step Into, Step Over, and so forth) 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

   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).

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.

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.

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.
3. When you finish debugging the script, click **Close** or **Exit**.

   DBArtisan 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.

Tutorial sessions

The following topics provide walkthrough sessions geared to new users of the Embarcadero SQL Debugger:

- Debugging a Sample Script with Embarcadero SQL Debugger for Microsoft
- Debugging a Sample Script with Embarcadero SQL Debugger for Oracle

Debugging a Sample Script with Embarcadero SQL Debugger for Microsoft

This Getting Started section demonstrates basic debugging functionality. You will debug two procedures using the Embarcadero SQL Debugger for Microsoft.

The section topics are designed to familiarize you with basic debugging features and functionality:

Getting Started with Embarcadero SQL Debugger for Microsoft

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 will use the Embarcadero SQL Debugger for Microsoft to identify this bug.

The Getting Started section includes:

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, open DBArtisan, 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”.

1 Start **DBArtisan**.

2 Connect to a Microsoft SQL Server 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

Press **CTRL+N**.

DBArtisan opens an SQL Editor in the current workspace.

5 In the **SQL Editor**, type the DDL for procedure **check_modulo**.

NOTE: Substitute your user name once in the DDL for this procedure.

6 On the **SQL Editor** toolbar, click **Execute**.

DBArtisan 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 will use the Embarcadero SQL Debugger for Microsoft to identify this bug.

CAUTION: When you input parameters, enter the smaller number in the @p_num1_in int box.

To create this procedure, open DBArtisan, 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

1. Start **DBArtisan**.
2. Connect to a Microsoft SQL Server 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

   Press **CTRL+N**.

   DBArtisan opens an SQL Editor in the current workspace.

5. In the **SQL Editor**, type the DDL for procedure **calculate_sum_with_overflow_bug**.

   NOTE: Substitute your user name twice in the DDL for this procedure.

6 On the **SQL Editor** toolbar, click **Execute**.

DBArtisan 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 DBArtisan's Datasource Explorer.

Confirming the Creation of the Procedures

To confirm the 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.

DBArtisan 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**.

Testing a Procedure

After you confirm the creation of the procedures, 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, 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**.

DBArtisan 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**.

DBArtisan 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**.

DBArtisan 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**.

DBArtisan returns an error stating “Arithmetic overflow occurred”.

Proceed to Starting the Debugging Session.

Starting the Debugging Session

After you test the procedure, open the procedure in Embarcadero SQL Debugger for Microsoft and enter input parameters before debugging.

To start the debugging 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.

DBArtisan 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**.

DBArtisan closes the dialog box and opens the Embarcadero SQL Debugger interface.

The Embarcadero SQL Debugger includes the following five windows:

- DDL Editor Window
- Watch Window
- Variables Window
- Call Stack Window
- Dependency Tree Window

Proceed to Breakpoints.

Breakpoints

After you start the debugging session, insert a breakpoint into the code of the procedure calculate_sum_with_overflow_bug. Then 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

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

Press **F9**.

DBArtisan inserts a breakpoint (indicated by a dot) next to the number of the target line.

3 On the **Debug** menu, click **Go**.

OR

Press **F5**.

DBArtisan Embarcadero SQL Debugger for Microsoft displays the value of the variables before the breakpoint in the Variables Window.

Proceed to Step Into.

Step Into

After setting the breakpoint, step into the dependent procedure, check_modulo.

To use the Step Into facility, do the following:

1 On the **Debug** menu, click **Step Into**.

OR

Press **F11**.

DBArtisan extracts the DDL for the dependent, nested procedure into the DDL Editor.

2 Step Into again.

DBArtisan 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.

Proceed to Step Out.

Step Out

After you Step Into the modulo_check (nested procedure) code, step back out and return to the calculate_sum_with_overflow_bug (outside procedure) code.

To use the Step Out facility, do the following:

1 On the **Debug** menu, click **Step Out**.

OR

Press **SHIFT+F11**.

DBArtisan opens the DDL Editor containing the code for calculate_sum_with_overflow_bug.

2 On the **Debug** menu, 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, DBArtisan closes the Debugger and opens a DDL Editor to the Results Tab.

Proceed to Correcting the Script.

Correcting the Script

When you finished Stepping Out of the nested code and encounter the error, to fully fix the bug, do the following:

1 Locate the source of the error

2 Scroll to the line in the script displaying the error

3 Analyze the code

4 Correct the error

5 Compile the corrected script

When you first executed the procedure, DBArtisan 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, DBArtisan returns the error message and the procedure fails.

To correct the script, do the following:

1. On the **Explorer Tab**, right-click **calculate_sum_with_overflow_bug**, and then click **Extract**.

   DBArtisan extracts the DDL for the procedure into a DDL Editor.

2. On the **Edit** toolbar, click **Find**.

   DBArtisan opens the Find dialog box.

3. In the **Find What** box, type **TINYINT**.

4. Click **Find Next**.

   DBArtisan 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.

   DBArtisan executes the script and opens the Results Tab.

7. On the **Explorer Tab**, right-click **calculate_sum_with_overflow_bug**, and then click **Execute**.

   DBArtisan 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**.

    DBArtisan 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.

For more information, see Getting Started with Embarcadero SQL Debugger for Microsoft.

Debugging a Sample Script with Embarcadero SQL Debugger for Oracle

The DBArtisan 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 familiarize you with basic debugging features and functionality:

- Getting Started with Embarcadero SQL Debugger for Oracle guides you through creating the package you use in Debugging Sample Script 1 and Debugging Sample Script 2.
- Debugging Sample Script 1 guides you through debugging functionality and demonstrates the Embarcadero SQL Debugger for Oracle interface features.
- Debugging Sample Script 2 guides you through debugging functionality and error correction.

NOTE: For the purposes of this walk-though we have created this package under the user name DEMO_SPENCE.

Getting Started with Embarcadero SQL Debugger for Oracle

The DBArtisan 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 DBArtisan installation places the script in the \UsrScrpt subfolder of the main DBArtisan installation directory.

If you create the package included with the DBArtisan installation, you can delete it and its objects from your system when you finish working with them. The objects to delete:

- 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

Proceed to Embarcadero SQL Debugger for Oracle Overview.

Embarcadero SQL Debugger for Oracle Overview

The Getting Started section guides you through:

- Opening the sample debug script.
- Executing sample debug script.
- Changing the Explorer Tab display.
- Confirming the creation of the package, including its functions and procedures.

Getting Started

1. Start **DBArtisan**.
2. On the **File** Menu, click **Open**.

   DBArtisan opens the Open File(s) dialog box.
3. In the **Open File(s)** dialog box, go to **DBA710\UsrScrpt\DEBUGGER_DEMO.sql**, and then click **Open**.

   > **NOTE:** During the installation DBArtisan places DEBUGGER_DEMO.sql in the \UsrScrpt subfolder of the main DBArtisan installation folder.

   DBArtisan 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**.

   DBArtisan opens the target script in an SQL Editor.
5. On the **SQL Editor** toolbar, click **Execute** to execute the script and create the package.

   DBArtisan 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**.

DBArtisan displays a list of owners in the Datasource Explorer.

7 On the **Explorer**, double-click your owner name.

DBArtisan displays a list of your schema objects.

8 Under your owner node, double-click the **Packages** node.

DBArtisan displays COUNT_TIME_INTERVAL, confirming the package's creation.

Proceed to Debugging Sample Script 1.

Debugging Sample Script 1

Sample Script 1 demonstrates Embarcadero SQL Debugger's basic features and functionality with the function WORKING_DAYS( ), which counts the number of business days between two dates.

Debugging Sample Script 1 is divided into five parts:

- Sample Script 1 - Starting the Debug Session
- Sample Script 1 - Entering Input Parameters
- Sample Script 1 - 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( ), 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.

DBArtisan opens the COUNT_TIME_INTERVAL node and displays the following items:

2 Under the **COUNT_TIME_INTERVAL** node, double-click **Functions**.

DBArtisan 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.

DBArtisan opens the Function Execution dialog box with the current date in the boxes.

Proceed to Sample Script 1 - Entering 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 enter input parameters, do the following:

1 Click the **P_START_DATE DATE** box, and then click the drop-down arrow.

DBArtisan opens a calendar.

2 On the calendar, click the left arrow to set the month to **November 1999**.

3 Click **1**.

DBArtisan displays 11/01/1999 in the Value column of P_START_DATE.

4 Click the **P_END_DATE DATE** box, and then click the drop-down arrow.

DBArtisan opens a new calendar.

5 On the calendar, click the left arrow to set the month to **November 1999**.

6 Click **8**.

DBArtisan displays 11/08/1999 in the Value column of P_END_DATE.

7 Click **OK**.

DBArtisan closes the Function Execution dialog box, and then opens the following five Embarcadero SQL Debugger for Oracle interface windows:

- DDL Editor Window, which displays the SQL code for the function.
- Watch Window
- Variables Window.
- Call Stack Window
- Dependency Tree Window, which displays the dependent objects.

Proceed to Sample Script 1- 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, you insert the breakpoints in the extracted dependent object code. After you extract this code, 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.

DBArtisan displays the SQL code for the package body in the SQL Editor window.

2 On the **Edit** toolbar, click **Find**.

DBArtisan 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, DBArtisan highlights the first occurrence of DBMS_OUTPUT, on line 22.

5 On the **SQL Editor** toolbar, click **Breakpoint**.

DBArtisan inserts a breakpoint next to the target line number.

6 On the **Find** dialog box, click **Find Next**.

DBArtisan highlights the next occurrence of DBMS_OUTPUT.

7 Click **Find Next** a third time.

DBArtisan highlights the next occurrence of DBMS_OUTPUT, on line 35.

8 On the **Find** dialog box, click **Cancel**.

DBArtisan closes the Find dialog box.

9 On the **Edit** toolbar, click **Breakpoint** to insert a second breakpoint.

You should now have breakpoints set at lines 22 and 35.

Proceed to Sample Script 1- Stepping Into.

Sample Script 1- Stepping Into

After you insert breakpoints, you can step into the function code.

Step Into

To use the Step Into facility, 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 runs to the next breakpoint.

7 On the **SQL Editor** toolbar, click **Go** once more.

Embarcadero SQL concludes the debug session and displays the Debug Session Results box.

Proceed to Sample Script 1 - Viewing Debug Session Results.

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 Results

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

To debug session results, do the following:

1. Click **OK**.

   DBArtisan closes the Debug Session Results box and terminates your debug session.

Proceed to Debugging Sample Script 2.

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, use Embarcadero SQL Debugger for Oracle to identify the bug, and then correct the script so that it can execute successfully.

Debugging Sample Script 2 is divided into six parts:

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 ( ), 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.

   DBArtisan opens the COUNT_TIME_INTERVAL node.

2. Double-click the **Functions** node.

   DBArtisan opens the Functions node.

3. Click **WEEKEND_DAYS ( )**, then right-click it and click **Execute**.

   DBArtisan 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**.

DBArtisan attempts to execute the function but returns an error indicating that the character string buffer is too small.

Proceed to Sample Script 2 - Starting the Debug Session.

Sample Script 2 - Starting the Debug Session

After you unsuccessfully execute the function WEEKEND_DAYS( ) and determine 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.

DBArtisan opens the Function Execution dialog box.

Proceed to Sample Script 2 - 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 input parameters, do the following:

1 At the end of the **P_START_DATE** row, click the drop-down arrow.

DBArtisan opens the calendar.

2 On the calendar, click **Left Arrow** to set the month to **November 1999**.

3 Click **1**.

DBArtisan 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.

DBArtisan opens a new calendar.

5 On the calendar, click **Left Arrow** to set the month to **November 1999**.

6 Click **30**.

DBArtisan displays 11/08/1999 in the Value column of the P_END_DATE row.

7 Click **Continue**.

DBArtisan closes the Function Execution dialog box, and then opens the following five Embarcadero SQL Debugger for Oracle interface windows:

- DDL Editor Window, which displays the SQL code for the function
- Watch Window
- Variables Window
- Call Stack Window
- Dependency Tree Window, which displays the dependent objects.

Proceed to Sample Script 2- Inserting Breakpoints.

Sample Script 2- Inserting Breakpoints

After you enter input parameters, you can begin inserting breakpoints. In this example, you insert the breakpoints in the extracted dependent object code. After you extract this code, 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.

DBArtisan displays the SQL code for the package body in the SQL Editor.

2 On the **Edit** toolbar, click **Find**.

DBArtisan 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**.

DBArtisan closes the Find dialog box.

5 Click line 60, the first line of executable code:

6 On the **SQL Editor** toolbar, click **Breakpoint**.

DBArtisan 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.

Proceed to Sample Script 2- Stepping Into.

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, monitor the Variables window. As you step through the code, the Variables window updates with the value of the variables.

Step Into

1 On the **SQL Editor** toolbar, click **Step Into**.

DBArtisan 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**.

DBArtisan moves the yellow arrow to the next line of the code, line 66.

4 On the **SQL Editor** toolbar, click **Step Into**.

DBArtisan 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**.

DBArtisan 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**.

DBArtisan 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.

Proceed to Sample Script 2 - Correcting 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:

- 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, 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**.

   DBArtisan 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.

   DBArtisan 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**.

   DBArtisan 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**.

   DBArtisan successfully executes the corrected function.

DBArtisan PL/SQL Profiler

The DBArtisan PL/SQL Profiler module lets Oracle 8i developers 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. DBArtisan PL/SQL Profiler collects and stores data in database tables that let you identify and isolate performance problems and provide code coverage information. The DBArtisan 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

Requirements:

- Oracle 8.1.5 or later
- Oracle built-in package DBMS_PROFILER
- Oracle tables:
  - PLSQL_PROFILER_RUNS
  - PLSQL_PROFILER_UNITS
  - PLSQL_PROFILER_DATA (user's schema)

The table below describes the sections of this chapter:

Section	Description
Setting Up DBArtisan PL/SQL Profiler	Describes the process of setting up DBArtisan PL/SQL Profiler.
DBArtisan PL/SQL Profiler Functionality	Describes the functionality of DBArtisan PL/SQL Profiler.
Using DBArtisan PL/SQL Profiler	Describes how to run a profile session.

NOTE: The DBArtisan PL/SQL Profiler is an optional add-on module.

Setting Up DBArtisan PL/SQL Profiler

The Oracle profiling tables must be on the Oracle server before you can use the DBArtisan PL/SQL Profiler. The first time you open the PL/SQL Profiler, DBArtisan checks the server for the profiling tables. If the profiling tables are not on the server, DBArtisan automatically starts Oracle SQL*Plus, which installs profiling tables on the Oracle server.

For DBArtisan to run Oracle SQL*Plus (which installs the profiling tables) your Oracle server and client must meet the following conditions:

- The Oracle server and the client have the same version of Oracle.
- The client has the Oracle\BIN directory on the path.
- The client has the Oracle file, SQLPLUS.exe in the Oracle\BIN directory.

- The following Oracle files are in the Oracle\RDBMS\ADMIN directory:
  - DBMSPBP.sql
  - PROFLOAD.sql
  - PROFTAB.sql
  - PRVTPBP.blp

  NOTE: If the Oracle server and the client machines are running different versions of Oracle, after running SQL*Plus, DBArtisan displays the following error message: Version of package is incompatible.

DBArtisan PL/SQL Profiler Explorer

The DBArtisan 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 DBArtisan 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.
Run level	Contains all Unit, Unit Name, Unit Type, Run Date\time data for the specific Run ID.

DBArtisan PL/SQL Profiler Functionality

DBArtisan PL/SQL Profiler offers the following functionality:

- Start
- Flush
- Run Summary
- Run Detail
- Unit Summary
- Clear Profile Table
- Unit Detail
- Stop

Start

DBArtisan 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 profile session, do the following:

1 On the **Utilities** menu, click **PL/SQL Profiler**, and then click **Start**.

DBArtisan 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**.

DBArtisan begins profiling.

Starting an Existing Profile Session

1 On the **Utilities** menu, click **PL/SQL Profiler**, and then click **Start**.

DBArtisan opens the PL/SQL Profiler - Start dialog box.

2 Click the **Profile Label** list, and then click the existing profile.

3 Click **OK**.

DBArtisan begins profiling.

For more information, see DBArtisan PL/SQL Profiler Functionality.

Flush

DBArtisan PL/SQL Profiler lets you move the data from the dynamic tables into analysis tables with the flush command.

The table below describes the options and functionality on the PL/SQL Profiler - Flush dialog box:

Option	Description
Flush	Click to delete the data in a running profile.
Flush & Analyze	Click to open the PL/SQL Profiler Run Detail window. For details, see Run Detail.
Cancel Button	Click to abort the flush and continue the profiling session.

NOTE: You can only Flush a running Profile.

Flushing a Profile

To flush a profile, do the following:

1 On the **Utilities** menu, click **PL/SQL Profiler**, and then click **Flush**.

DBArtisan opens the PL/SQL Profiler - Flush dialog box.

For more information, see DBArtisan PL/SQL Profiler Functionality.

Run Summary

The DBArtisan PL/SQL Profiler Run Summary window lets you to 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 **Utilities** menu, click **PL/SQL Profiler**, and then click **Run Summary**.

   DBArtisan opens the PL/SQL Profiler - Run Summary window.

2. In the **PL/SQL Profiler - Run Summary** window, click the **Label** list, and then click the target profile to populate the table.

For more information, see DBArtisan PL/SQL Profiler Functionality.

Run Detail

The DBArtisan PL/SQL Profiler Run Detail window lets you to view the following information for each of your profiles:

- Run Number
- Run Date
- Run Time

The Run Detail Tab lets you:

- View the information for all runs or you can view profile information based on the unit type or unit owner.
- View results in milliseconds, seconds and minutes.
- View graphical displays of the profiling data that let you can navigate to the specific unit within the summary portion of the window.
- Specify the number of top lines to display in the graphical portion of the interface.

  TIP: Each graph is a working object. You can select data in a graph and the corresponding line of source displays in the lower pane of the interface.

Opening the Run Detail Window

To open the Run Detail Window, do the following:

1. On the **Utilities** menu, click **PL/SQL Profiler**, and then click **Run Detail**.

   DBArtisan opens the PL/SQL Profiler - Run Detail window.

2. 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.

For more information, see DBArtisan PL/SQL Profiler Functionality.

Unit Summary

The DBArtisan PL/SQL Profiler Unit Summary window lets you to view the following information for each of your profiles:

- Run ID
- Run Date
- Run Time
- Unit Time
- Percentage of Run Time

The DBArtisan 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 navigate 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 **Utilities** menu, click **PL/SQL Profiler**, and then click **Unit Summary**.

   DBArtisan 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.

For more information, see DBArtisan PL/SQL Profiler Functionality.

Clear Profile Table

DBArtisan 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 **Utilities** menu, click **PL/SQL Profiler**, and then click **Clear Profile Table**.

   DBArtisan clears the profile table.

2. In the DBArtisan dialog box, if you are sure that you want to clear out the profiler tables, click **Yes**.

For more information, see DBArtisan PL/SQL Profiler Functionality.

Unit Detail

The DBArtisan PL/SQL Profiler Unit Detail window lets you to view the following information for each of your profiles:

- Average Time
- Source

- PL/SQL Script

The DBArtisan 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). Additionally, also displays graphs of execution statistics for the top N run. You can use the graphical displays to navigate 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 DBArtisan 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 **Utilities** menu, click **PL/SQL Profiler**, and then click **Unit Detail**.

   DBArtisan opens the PL/SQL Profiler - Unit Detail window.

2. 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 Detai**l window, right-click, and then click **Advanced View** to populate the table with Advanced View information.

For more information, see DBArtisan PL/SQL Profiler Functionality.

Stop

DBArtisan PL/SQL Profiler Stop command pauses the data gathering operation. Stop & Analyze populates the summary tables so that you can view the Unit Detail and Run Summary windows.

The table below describes the options and functionality on the PL/SQL Profiler - Stop dialog box:

Option	Description
Stop	Click to stop the profiling session.
Stop & Analyze	For details, see Run Detail..
Cancel	Click to continue the profiling session.

Stopping a Profiling Session

To stop a Profiling Session, do the following:

1 On the **Utilities** menu, click **PL/SQL Profiler**, and then click **Stop**.

DBArtisan opens the PL/SQL Profiler - Stop dialog box.

For more information, see DBArtisan PL/SQL Profiler Functionality.

Using DBArtisan PL/SQL Profiler

The steps in this section provide a high level overview of running a profiling session, and cover the following processes:

- Starting the Session.
- Executing the Sample Script.
- Stopping and Analyzing the Session.

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.

Using the DBArtisan PL/SQL Profiler

1 On the **Utilities** menu, click **PL/SQL Profiler**, and then click **Start**.

DBArtisan 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**.

DBArtisan begins profiling.

4 On the **Datasource Explorer**, execute on 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 **Utilities** menu, click **PL/SQL Profiler**, and then click **Stop.**

DBArtisan opens the PL/SQL Profiler - Stop dialog box.

The table below describes the options and functionality on the PL/SQL Profiler - Stop dialog box:

Option	Description
Stop	Click to stop the profiling session.
Stop & Analyze	Click to open the **PL/SQL Profiler Run Detail** window. For details, see Run Detail. 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.
Cancel	Click to continue the profiling session.

6 Use the **Utilities** menu to open any of the following PL/SQL Profiler windows. For more information, see:

- Run Summary
- Unit Summary
- Unit Detail

For more information, see DBArtisan PL/SQL Profiler Functionality.

Sample Profiling Session

The DBArtisan 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 profile in the walk through, and the PROFILER_DEMO.SQL is what you profile during the walk through.

NOTE: To create the objects in the PROFILER_BUILD_DEMO.SQL script, you need CREATE privileges.

The sample script demonstrates the following features of the DBArtisan PL/SQL Profiler:

- Unit Detail
- Run Detail
- Show Only Hit Lines
- Advanced View

During the installation, DBArtisan places the scripts in the \UsrScrpt subfolder of the main DBArtisan installation folder.

Overview

Sample Profiling Session is divided into six parts:

- Getting Started
- Starting the Session

- Executing the Sample Script
- Stopping the Session
- Re-running & Re-executing the Session
- Stopping & Analyzing

Proceed to Sample Profiling Session - Getting Started.

Sample Profiling Session - Getting Started

In this step of Sample Profiling Session, you create the objects that you profile in the walk through.

Overview

The Getting Started section guides you through:

- Opening PROFILER_BUILD_DEMO.SQL.
- Changing the Datasource Explorer Display.
- Confirming the Creation of the Package.

Getting Started

1. Start **DBArtisan**.
2. On the **File** menu, click **Open**.

   DBArtisan opens the Open Files dialog box.
3. In the **Open Files** dialog box, navigate to the UsrScrpt subfolder of the main DBArtisan installation folder, press ENTER, and then double-click **PROFILER_BUILD_DEMO.SQL** to open the script in a SQL Editor window.

   DBArtisan opens the PROFILER_BUILD_DEMO.SQL script in an SQL Editor window.
4. On the **SQL Editor** window, click **Execute**.

   DBArtisan executes the script and create the package.
5. On the **Datasource Explorer** window list, click **Organize by Owner**.
6. On the **Datasource Explorer** window, click the node of your owner name.

   DBArtisan 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.

Proceed to Sample Profiling Session - Starting the Session.

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**.

   DBArtisan opens the Open Files dialog box.

2. In the **Open Files** dialog box, type the path to the UsrScrpt directory, press ENTER, and then double-click **PROFILER_DEMO.SQL**.

   DBArtisan opens the script in a SQL Editor window.

   DBArtisan opens the PROFILER_DEMO.SQL script in an SQL Editor window.

3. On the **Utilities** menu, click **PL/SQL Profiler**, and then click **Start**.

   DBArtisan opens the PL/SQL Profiler - Start dialog box.

4. In the **Profile Label** list, enter **DemoProfile**.

5. Click **OK**.

   DBArtisan begins the profiling session.

   NOTE: If this is the first time you start the PL/SQL Profiler, DBArtisan displays a dialog box.

   NOTE: Click Yes and DBArtisan opens SQL*Plus to create the tables. You need to start the profiling session again (see step 3 above.)

Proceed to Sample Profiling Session - Executing the Sample Script.

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**.

   DBArtisan executes the script and opens a Results Tab.

Proceed to Sample Profiling Session - Getting Started.

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 **Utilities** menu, click **PL/SQL Profiler**, and then click **Stop**.

   DBArtisan opens the PL/SQL Profiler - Stop dialog box.

2. Click **Stop**.

Proceed to Sample Profiling Session - Re-running & Re-executing the Session.

Sample Profiling Session - Re-running & Re-executing the Session

In this step of Sample Profiling Session, you run the same profile session and execute the DEMO script again.

Sample Profiling Session - Re-running & Re-executing the Session

To re-run and re-execute the session, do the following:

1. In the **SQL Editor,** click the **Query Tab**.
2. On the **Utilities** menu, click **PL/SQL Profiler**, and then click **Start**.

   DBArtisan opens the PL/SQL Profiler - Start dialog box again.
3. Click the down arrow on the **Profile Label** list, and then click **DemoProfile**.
4. Click **OK**.

   DBArtisan begins the profiling session.
5. On the **SQL Editor** toolbar, click **Execute**.

   DBArtisan executes the script again and opens the Results Tab.

Proceed to Sample Profiling Session - Stopping & Analyzing.

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 the sample profiling session, do the following:

1. On the **Utilities** menu, click **PL/SQL Profiler**, and then click **Stop**.

   DBArtisan opens the PL/SQL Profiler - Stop dialog box again.
2. Click **Stop & Analyze**.

   DBArtisan opens the PL/SQL Profiler - Run Detail window.
3. Click the **Run** list, and then click **Run#x**.

   NOTE: DBArtisan 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.

   DBArtisan PL/SQL Profiler populates the grid with information on the procedure, package body and package specification.

   NOTE: For the purposes of this walk though we have created this package under the account SCOTT.
4. Click the **Run** list again, and then click the **Run#x** 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**.

   DBArtisan 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 DBArtisan PL/SQL Profiler shows only the lines of code that executed.

9 Right-click, and then click **Advanced View.**

The DBArtisan 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 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_( ), YEARS_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

Capacity Analyst

Welcome to the Embarcadero Capacity Analyst. The Capacity Analyst collects critical database storage, object, and performance metrics at scheduled intervals and provides visual analysis and reporting mechanisms that allow for historical trend analysis and forecasting.

Capacity Analyst lets you:

- Create and maintain Capacity Collections that serve as containers for related statistics.
- Collect statistics in Capacity Collections.
- View graphic and report-styled presentations of trend and predictive analysis scenarios.
- Maintain functions regarding the handling of stored statistics.

The table below describes the major sections of Help.

Section	Description
Geting started with the Capacity Analyst	Provides information about installing, opening, and licensing Capacity Analyst.
Capacity Analyst tutorial	Helps you get started with Embarcadero Capacity Analyst. Provides the foundation you need to explore the many features and benefits of Capacity Analyst.
Using Capacity Analyst	Shows how to use each component of the application. Using Capacity Analyst breaks each process and procedure into simple step-by-step instructions.

Geting started with the Capacity Analyst

Welcome to the Embarcadero Capacity Analyst. The Capacity Analyst helps you analyze current usage trends in your key databases and lets you forecast where your databases are heading in the future.

For more information, see:

- What You Should Know Before Installation
- Capacity Analyst Repository Installation
- Uninstalling Capacity Analyst Repository
- Managing Repository Security
- Repository Information
- Licensing Capacity Analyst

What You Should Know Before Installation

When you install the Analyst, several Embarcadero-generated stored objects are pushed onto the database server:

These objects are as follows:

>

SQL SERVER

Procedures

eas_apply_collection_filter	eas_apply_collection_schedule	eas_apply_purge_condition
eas_apply_refresh_stats_condition	eas_create_collection	eas_delete_collection
eas_get_collection_metric_info	eas_get_metric_query	eas_get_metric_query_ex
eas_get_metric_query_ex1	eas_get_purge_condition	eas_get_purge_status
eas_get_refresh_stats_condition	eas_get_refresh_stats_status	eas_is_admin
eas_is_auto_purge_set	eas_is_collection_available	eas_is_collection_scheduled
eas_is_filter_all_set	eas_list_perf_stats	eas_list_wait_events
eas_purge_data	eas_refresh_stats	eas_resolve_collection_objects
eas_resolve_datastruct	eas_resolve_indexes	eas_resolve_perfstats
eas_resolve_tables	eas_resolve_waitevents	eas_run_collection
eas_run_datastruct	eas_run_perfstats	eas_run_waitevents
eas_set_auto_purge_status	eas_set_collection_status	eas_set_refresh_stats_status
eas_update_schedule_dates		

Foreign Keys

eas_collection_fk1	eas_collection_fk2	eas_collection_object_fk1
eas_collection_object_fk2	eas_filter_fk1	eas_filter_fk2
eas_group_recfk	eas_map_fk1	eas_object_fk
eas_object_type_fk	eas_option_fk1	eas_property_fk1
eas_property_fk2	eas_schedule_fk	eas_statistic_fk1
eas_statistic_fk2	eas_statistic_fk3	

Tables

eas_collection	eas_collection_object	eas_collection_sql
eas_database	eas_filter	eas_group
eas_map	eas_object	eas_object_type
eas_option	eas_option_type	eas_property
eas_schedule	eas_statistic	eas_version

Metadata

eas_column_map_insert	eas_database_insert	eas_group_insert
eas_object_type_insert	eas_option_insert	eas_option_type_insert
eas_version_insert		

DB2

Foreign Keys

eas_col_obj_fk1	eas_col_obj_fk2	eas_collection_fk1
eas_collection_fk2	eas_filter_fk1	eas_filter_fk2
eas_group_recfk	eas_map_fk1	eas_object_fk
eas_object_type_fk	eas_option_fk1	eas_property_fk1
eas_property_fk2	eas_schedule_fk	eas_statistic_fk1
eas_statistic_fk2	eas_statistic_fk3	

Index

eas_collection_idxs	eas_collection_object_idxs	eas_filter_idxs
eas_group_idxs	eas_map_idxs	eas_object_type_idxs
eas_option_idxs	eas_property_idxs	eas_schedule_idxs
eas_statistic_idxs		

Metadata

eas_column_map_insert	eas_group_insert	eas_object_type_insert
eas_option_insert	eas_option_type_insert	eas_version_insert

Tables

eas_collection_object	eas_collection_sql	eas_database
eas_group	eas_map	eas_object_type
eas_option	eas_option_type	eas_schedule
eas_version		

ORACLE

Foreign Keys

eas_collection_fk1	eas_collection_fk2	eas_collection_object_fk1
eas_collection_object_fk2	eas_collection_sql_fk1	eas_filter_fk1
eas_filter_fk2	eas_map_fk1	eas_object_fk
eas_option_fk1	eas_property_fk1	eas_property_fk2
eas_schedule_fk	eas_statistic_fk2	eas_statistic_fk3

Indexes

eas_collection_fk1	eas_collection_fk2	eas_filter_fk1
eas_filter_fk2	eas_object_fk	eas_object_uk
eas_option_fk1	eas_property_fk1	eas_property_fk2
eas_schedule_fk	eas_statistic_fk1	

>

Metadata

eas_column_map_insert	eas_object_type_insert	eas_option_insert
eas_option_type_insert	eas_version_insert	

Tables

eas_collection	eas_collection_object	eas_collection_sql
eas_database	eas_filter	eas_map
eas_object	eas_object_type	eas_option
eas_option_type	eas_property	eas_schedule
eas_statistic	eas_version	

SYBASE

Foreign Keys

eas_collection_fk1	eas_collection_fk2	eas_collection_object_fk1
eas_collection_object_fk2	eas_filter_fk1	eas_filter_fk2
eas_group_recfk	eas_map_fk1	eas_object_fk
eas_object_type_fk	eas_option_fk1	eas_property_fk1
eas_property_fk2	eas_schedule_fk	eas_statistic_fk1
eas_statistic_fk2	eas_statistic_fk3	

Indexes

eas_collection_idxs	eas_collection_object_idxs	eas_filter_idxs
eas_group_idxs	eas_map_idxs	eas_object_idxs
eas_object_type_idxs	eas_option_idxs	eas_property_idxs
eas_schedule_idxs	eas_statistic_idxs	

Metadata

eas_column_map_insert	eas_database_insert	eas_group_insert
eas_object_type_insert	eas_option_insert	eas_option_type_insert
eas_version_insert		

Procedures

sp_embt_ca_check_syb_version	sp_embt_ca_db_obj_space	sp_embt_ca_dbspace
sp_embt_ca_delete_collection	sp_embt_ca_get_col_metric_info	sp_embt_ca_get_metric_query_ex
sp_embt_ca_get_purge_cond	sp_embt_ca_get_purge_status	sp_embt_ca_get_refresh_cond
sp_embt_ca_get_refresh_status	sp_embt_ca_is_admin	sp_embt_ca_is_auto_purge_set
sp_embt_ca_is_col_available	sp_embt_ca_is_col_scheduled	sp_embt_ca_is_filter_all_set
sp_embt_ca_purge_data	sp_embt_ca_run_collection	sp_embt_ca_run_datastruct

sp_embt_ca_set_col_status	sp_embt_ca_upd_schedule_dates	

Table

eas_collection_object	eas_collection_sql	eas_database
eas_group	eas_map	eas_object_type
eas_option	eas_option_type	eas_schedule
eas_version		

Opening the Capacity Analyst

Capacity Analyst runs as a tab in the main DBArtisan window.

To open the Capacity Analyst:

1. Start DBArtisan and connect to the datasource that you want to run reports or analysis against. For details, see Datasource Management.
2. From the **Analyst** menu select **Capacity Analyst**.

   OR

   On the Analyst Series toolbar click the Capacity Analyst button.

Capacity Analyst opens, displaying any existing collections.

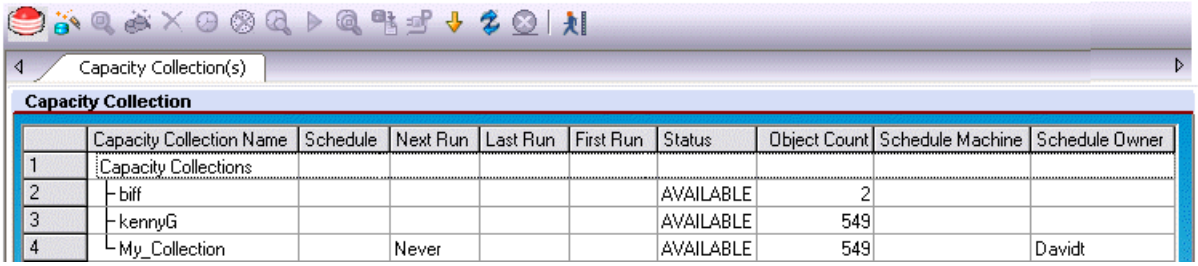

Capacity Analyst Repository Installation

The Capacity Analyst Repository Installation Wizard lets you install the repository on each database the Capacity Analyst runs against. The repository holds all the captured metrics that you want to collect. For example, this can be the space used in an Oracle tablespace or the number of physical reads on a Microsoft SQL Server.

The wizard installs a set of objects to manage the captured data, and a security role to manage access to the repository and objects.

The Analyst Repository Installation Wizard lets you:

- Create a new repository.
- Upgrade an existing repository.
- Assign a repository owner.

- Assign object placement in the database.
- View feedback on the repository install actions.

Completing the Analyst Repository Installation Wizard

To install the repository, do the following:

1 On the **DBArtisan** toolbar, click the **Capacity Analyst** icon.

DBArtisan opens Capacity Analyst. Capacity Analyst detects that no repository is present and opens the Analyst Repository Installation Wizard.

2 Use the following table as a guide to understanding and setting options in the Wizard,

Panel	Option	Description
1	Create a new repository	Lets you create a new repository under a new schema.
	Upgrade existing repository	Lets you upgrade the version of a current repository. Updates the stored procedures, and adds any new columns to tables. Also lets you use a an existing repository that was created for a different Analyst. **NOTE:** Available only if the wizard detects the presence of an earlier version.
	Use current login information	**IBM DB2 ONLY:** Select to use the default connectivity information as the login id/password.
	Login Name	Lets you type the server administrator ID. **NOTE:** An administrator ID is not required after the initial repository setup.
	Password	Lets you type the server administrator password.
	Login As	**ORACLE ONLY:** Lets you select one of three Oracle login modes: Default - Establishes connection with no system roles. SYSDBA - Lets you perform all DBA-related functions, like startup and shutdown. **NOTE:** SYSDBA is required for the SYS id when you log into a 9i instance. SYSOPER - Lets you perform many DBA-related functions, like startup and shutdown.
	Using	**MICROSOFT SQL SERVER:** Lets you specify authentication type, either SQL Server Authentication or Windows Authentication.
2	Database	MICROSOFT SQL SERVER ONLY: Lets you select a database. SYBASE ASE ONLY: The following system databases are excluded from the list of potential targets for Repository installation: 'Replication' 'model' 'sybsecurity' 'sybsystemdb' 'sybsystemprocs' 'tempdb' 'dbccdb' 'sybdiag'
	Login	**MICROSOFT SQL SERVER and SYBASE ASE ONLY:** Lets you select an owner. **NOTE:** Should have the CREATE PROCEDURE privilege on the storage device that is selected. For Oracle, the storage for tables and storage for indexes are tablespaces. The owner must have privileges for the tablespaces.
	Owner	**ORACLE ONLY:** Select a new user to own the repository or click New to open the User Wizard. For details, see User Wizard (Oracle). **NOTE:** Should have the CREATE PROCEDURE privilege on the storage device that is selected. For Oracle, the storage for tables and storage for indexes are tablespaces. The owner must have privileges for the tablespaces.
	Storage for tables (and its available space)	Lets you separate the storage placement of repository tables and indexes. For Oracle, the storage for tables are tablespaces. Select from the list. Or click New to open the creation wizard.

Panel	Option	Description
	Storage for indexes (and its available space)	Lets you separate the storage placement of repository tables and indexes. For Oracle, the storage for indexes are tablespaces. Select from the list. Or click New to open the creation wizard.
	Password	Lets you specify a password for the repository owner.
	Schema	**IBM DB2 ONLY:** Lets you select a schema. All the tables in the repository will be created for this schema. The default value is EMBTCA.
	Tablespace for Repository Tables	**IBM DB2 ONLY:** Lets you select the tablespace on which the repository tables are created.
	Tablespace for Repository Indexes	**IBM DB2 ONLY:** Lets you select the tablespace on which the repository tables are created.
	New	**IBM DB2 ONLY:** Click to open the Tablespace WIzard. FOr details, see Tablespace Wizard for IBM DB2 (IBM DB2 LUW).
3	Repository Login	Displays repository login
	Execution Log	Displays installation log.
	Execution Errors	Displays any errors if they occur.
	Finish	Click to complete the installation. Capacity Analyst installs the necessary files and displays the installation process. When the installation finishes, it displays a finished message. Click Finish to open the Capacity Analyst Tab. If the installation fails, Capacity Analyst lets you stop the repository installation. Capacity Analyst deletes the objects that it created.

3 When ready, click **Install**.

Capacity Analyst creates the repository. Capacity Analyst prompts you to grant other user accounts the privilege of using Capacity Analyst.

4 Select a basic security role:

- Administrator – Users granted this privilege can create, edit, remove, and view capacity collection information.
- User - Users granted this privilege can only view capacity collection information.

TIP: You can easily grant or remove users from these roles at repository installation time or afterwards. For details, see Managing Repository Security.

Uninstalling Capacity Analyst Repository

The Capacity Analyst Repository Uninstall Wizard reverses the process completed by the Analyst Repository Installation Wizard.

Completing the Capacity Analyst Repository Uninstall Wizard

1 In the DBArtisan **Datasource Explorer**, connect to a datasource that has repository installed.

2 On the **Analyst** toolbar, click the **Capacity Analyst** button.

3 On the **Capacity Collection Editor** toolbar, select **Command**, **Repo Management**, and then **Uninstall Analyst Repo**.

Capacity Analyst opens the Capacity Analyst Repository Uninstall Wizard.

4 Complete the wizard panels.

5 Click **Finish**.

Managing Repository Security

The Manage Repository Security dialog box lets you manage repository security.

The table below describes the options and functionality on the Manage Repository Security dialog box:

Option	Description
Grant	Click to grant privileges to target user.
Revoke	Click to revoke privileges from target user.

Completing the Manage Repository Security Dialog Box

1 In the DBArtisan **Datasource Explorer**, connect to a datasource that has repository installed.

2 On the **Analyst** toolbar, click the **Capacity Analyst** button.

3 On the **Capacity Collection Editor** toolbar, select **Command**, **Repo Management**, and then **Manage Analyst Repo**.

Capacity Analyst opens the Manage Repository Security dialog box. The Manage Repository Security dialog box lets you manage security for the Capacity Analyst.

4 To grant privileges, select the target user, and then click **Grant**.

Capacity Analyst opens the Grant Privilege(s) To dialog box. The Grant Privilege(s) To dialog box lets you select or clear the check boxes corresponding to the target privileges.

5 To revoke privileges, select the target user, and then click **Revoke**.

Capacity Analyst opens the Revoke Privilege(s) From dialog box. The Revoke Privileges (s) From dialog box lets you select or clear the check boxes corresponding to the target privileges.

6 Click **OK**.

Repository Information

The Repository Information dialog box displays a tree view to present the information about the repository. The table below describes the standard nodes all datasource platforms:

Node	Description
Repository Owner	Displays the login name of the repository owner.
Object Placement	Displays the database (for SQL Server and Sybase), tablespace (for Oracle), placement for tables and indexes (filegroup name - SQL Server, segment name - Sybase, tablespace - Oracle and DB2).
Repository Version	Schema and logic version for the tables/indexes and procedures respectively.

Microsoft SQL Server

The Repository Information dialog box displays the following information for Microsoft SQL Server:

- Repository owner
- Repository version
- Object placement
  - Database
  - Filegroup name where tables are stored
  - Filegroup name where indexes are stored.
- Dates
  - Date Created

Sybase ASE

The Repository Information dialog box displays the following information for Sybase ASE:

- Repository owner
- Repository version
- Object placement
  - Database
  - Segment name where tables are stored
  - Segment name where indexes are stored.
- Dates
  - Date Created

Oracle

The Repository Information dialog box displays the following information for Oracle:

- Repository owner
- Repository version
- Object placement
  - Tablespace name where indexes and tables are stored
- Dates
  - Date Created
  - Date Updated

IBM DB2

The Repository Information dialog box displays the following information for IBM DB2:

- Repository schema
- Repository owner
- Repository version
- Object placement
  - Tablespace name where indexes and tables are stored

- Dates
  - Date Created

Completing the Repository Information Dialog Box

1 In the DBArtisan **Datasource Explorer**, connect to a datasource that has repository installed.

2 On the **Analyst** toolbar, click the **Capacity Analyst** button.

3 On the **Capacity Collection Editor** toolbar, select **Command**, **Repo Management**, and then **Manage Analyst Repo**.

Capacity Analyst opens the Repository Information dialog box.

4 To save repository information as a report, click **Save**.

Capacity Analyst opens the Save Grid dialog box.

5 In **Save In**, select the location for the repository file.

6 In **File Name**, type the file name.

7 In **Save as Type**, select a file type.

8 Click **Save**.

Capacity Analyst generates the report and returns to the Repository Information dialog box.

9 Click **OK**.

Capacity Analyst closes the Repository Information dialog box.

Licensing Capacity Analyst

The Capacity Analyst is licensed per named server. You are issued one license key for the number of databases/servers you want to license. For example, if you purchase ten licenses for Capacity Analyst, you only need to obtain one license key from Embarcadero Technologies.

For more information, seeCapacity Analyst License Manager.

Capacity Analyst License Manager

The Analyst License Manager lets you manage licenses.

The table below describes the options and functionality on the Analyst License Manager:

Option	Description
License Summary	Displays basic licensing metadata for the product.
Available Datasources	Lets you select a datasource from the list of datasource available to license.
License	Click to add a license to a datasource.
Licensed Datasource(s)	Displays datasources already licensed.
Remove	Click to remove a license from a datasource.
Apply	Click to make you license choices permanent.

Option	Description
Cancel	Click to close the Licensing Dialog box. **NOTE:** If unsaved license changes have been made, the Analyst asks if you want to discard their changes before exiting.

Completing the Capacity Analyst License Manager

To complete the Analyst License Manager, do the following:

1. Open DBArtisan.
2. On the **Help** menu, click **Licensing**.

   DBArtisan opens the License Configuration Dialog box.
3. Click **License**.

   DBArtisan opens the Analyst License Manager.
4. Select options.
5. To make your license choices permanent, click **Apply**.
6. To close the Analyst License Manager, click **Cancel**.

Capacity Analyst Add an Analyst License Dialog Box

The Add an Analyst License Details dialog box lets you add a license.

The table below describes the options and functionality on the Add a Analyst License Details dialog box:

Option	Description
Select a DBMS platform	Lets you select the DBMS platform.

For more information, see Capacity Analyst License Manager.

Capacity Analyst tutorial

Embarcadero Capacity Analyst offers intelligent insight into the past, current, and future inner-workings of a database. By collecting historical database and performance metrics at user-defined intervals, Capacity Analyst provides historical trend analysis along with smart forecasting abilities that let you plan for the future needs of key databases. Capacity Analyst runs inside the Embarcadero DBArtisan console so you can perform all your database administration and analysis in one place.

Embarcadero Capacity Analyst lets you analyze current trends in your database storage and systems performance and helps you understand when corrective action is necessary. Capacity Analyst also helps you forecast the future needs of the your database so that your key systems can continue to meet the growing needs of its users.

Proceed to About this Capacity Analyst tutorial

About this Capacity Analyst tutorial

This short tutorial helps you get started with Embarcadero Capacity Analyst. After you complete this exercise. you will have the foundation you need to explore the many features and benefits of Capacity Analyst. You will learn how to set up Capacity Analyst and become familiar with creating and scheduling statistical collections. You will also understand how to analyze database trends and perform future forecasts of your database's needs.

Proceed to Summary of the tasks in this tutorial

Summary of the tasks in this tutorial

The main categories of Capacity Analyst activities are:

- **Creating and editing collections** - a collection is basically a specified set of metrics. Capacity Analyst lets you use a wizard to create collections, specifying DBMS-specific metrics as content. You can edit and clone collections as well.
- **Executing collections** - based on counters, timers, and statistics packages available from the DBMS you are working against, collections store calculated totals. Executing a collection updates the totals in the collection. Depending on your needs, you can execute collections in an ad hoc fashion or you can schedule collections to be executed at regular intervals.
- **Analysis and reporting** - Capacity Analyst provides prepackaged reports and a visual analysis tool.

This tutorial walks you through a sample of these activities. You will create a collection with a small variety of metrics, schedule the collection to be updated each hour, and use the analysis and reporting tools.

Proceed to Choosing an appropriate database for the tutorial exercises

Choosing an appropriate database for the tutorial exercises

The full power of the analysis and reporting tools can only be realized if you work through the exercises against an active database. For best results, use an active test bed or a production system.

Proceed to Choosing an appropriate database for the tutorial exercises

Rationale tor a two-part tutorial

While a collection stores historical statistics, totals cannot reflect database activities that occurred before the collection was created. In order to view meaningful analysis and reports, enough time must pass to allow a bare minimum of database activity. With that in mind, this tutorial is presented in two parts:

- Session 1: Creating and scheduling a collection
- Session 2: Using analysis and reporting tools

After completing the exercises in session 1, you should wait minimally four hours before starting the second session. THis should allow sufficient database activity to allow for meaningful analysis.

Session 1: Creating and scheduling a collection

In Session 1, you will first ensure that you are properly set up to evaluate the Capacity Analyst. Subsequently, you will run the Capacity Analyst, build a small statistics collection, and ensure that it is updated on an hourly basis.

Proceed to Getting Started with Capacity Analyst (optional)

Getting Started with Capacity Analyst (optional)

The following topics prepare you for first use of the Capacity Analyst:

- Downloading and Installing DBArtisan
- Starting DBArtisan and connecting to a datasource
- Installing the Capacity Analyst repository

When you have performed these tasks, or if you have already performed them, you can proceed to Creating and Scheduling a Capacity Collection Using Capacity Analyst.

Downloading and Installing DBArtisan

You can obtain the latest version of Capacity Analyst by downloading the DBArtisan software from the Embarcadero Web site at http://www.embarcadero.com. Capacity Analyst runs fully-contained within DBArtisan so you will need the DBArtisan product to use Capacity Analyst.

To download and install DBArtisan:

1. Locate the download area of the Web site and follow the steps indicated.
2. Save the file on your computer and then double-click to launch the self-extracting file that will guide you through the installation process.

When you first install an evaluation copy of DBArtisan, you can use the tool for the duration of the trial period. After that time, a permanent license is needed.

Proceed to Starting DBArtisan and connecting to a datasource

Starting DBArtisan and connecting to a datasource

After installing DBArtisan, you must set up datasources to establish reusable connections to your database servers. For now, you can establish a single connection to the database/server you will use to evaluate Capacity Analyst. DBArtisan can detect and let you register any datasource in your network.

To connect to a database/server:

1. Start DBArtisan. For a default installation, **Start** > **Programs** > **Embarcadero DBArtisan** *version number* > **DBArtisan** (where *version number* corresponds to the version of DBArtisan that you downloaded).
2. On the **Datasource** menu, select **Discover Datasources**.

   The **Discover Datasources** dialog opens. After searching your network it lists all datasources to which you can connect.

3 Select the check box corresponding to the target datasource and click **Register**.

The datasource is added to the left-hand pane of the datasource explorer, under the relevant DBMS folder. For example:

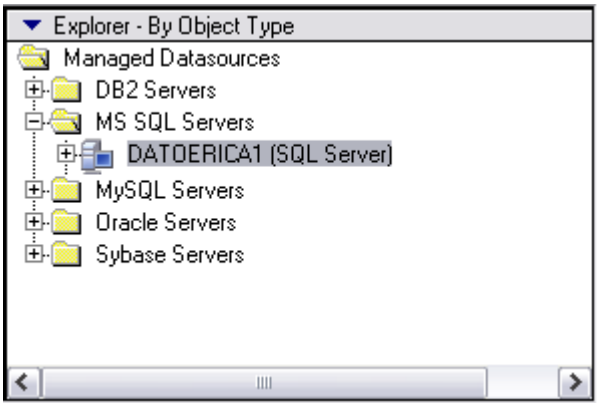

4 RIght-click the datasource and select **Connect** from the context menu.

A login dialog opens.

5 Provide valid user name and password credentials and click **OK**.

Proceed to Installing the Capacity Analyst repository

Installing the Capacity Analyst repository

Capacity Analyst makes use of a repository to store historical statistics about your database. A repository is required for each database/server that you use with Capacity Analyst.

To install a repository on the datasource you connected to:

1 From the **Analyst** menu, select **Capacity Analyst**.

OR

On the Analyst Series toolbar, click the **Capacity Analyst** button.

DBArtisan opens Capacity Analyst. Capacity Analyst detects that no repository is present and opens the Analyst Repository Installation Wizard.

Depending on the type of DBMS type of the datasource you connected to, specific options will differ. In general, the Wizard prompts for login/password credentials, as well as server, database, and storage details.

2 Provide the information that the Wizard prompts you for and when complete, click the **Install** button and provide additional information when prompted.

Capacity Analyst creates the necessary repository objects.

Proceed to Creating and Scheduling a Capacity Collection Using Capacity Analyst

Creating and Scheduling a Capacity Collection Using Capacity Analyst

When you open Capacity Analyst, Capacity Analyst displays a simple Capacity Collection Manager interface that lets you view existing collections, create new collections, and perform many other management functions.

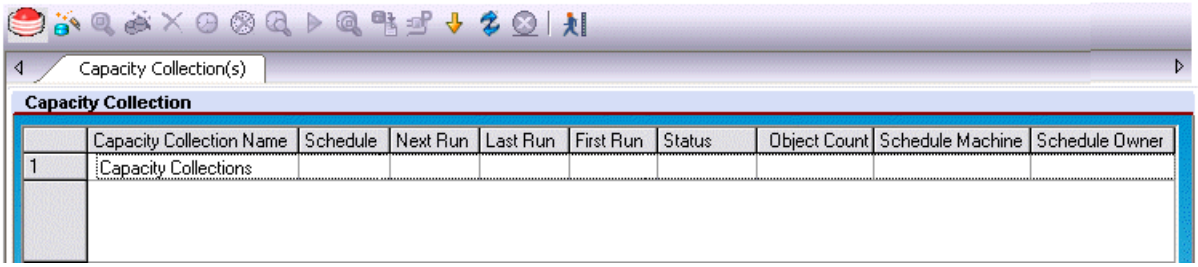

Proceed to Creating a New Capacity Collection Using Capacity Analyst

Creating a New Capacity Collection Using Capacity Analyst

To get started collecting information for your database, you first need to create a new collection. A collection is essentially a set of related metrics or statistics that you want to report on. As you begin to make use of Capacity Analyst, you will develop a strategy for metrics you will store in each collection. For the purposes of this tutorial, you can build an arbitrary set of statistics into a collection.

You create a collection using the Capacity Collection Wizard.

To open the Capacity Collection Wizard:

1 On the Capacity Analyst toolbar, click the **New Collection Wizard** button.

Capacity Analyst opens the Capacity Collection Wizard.

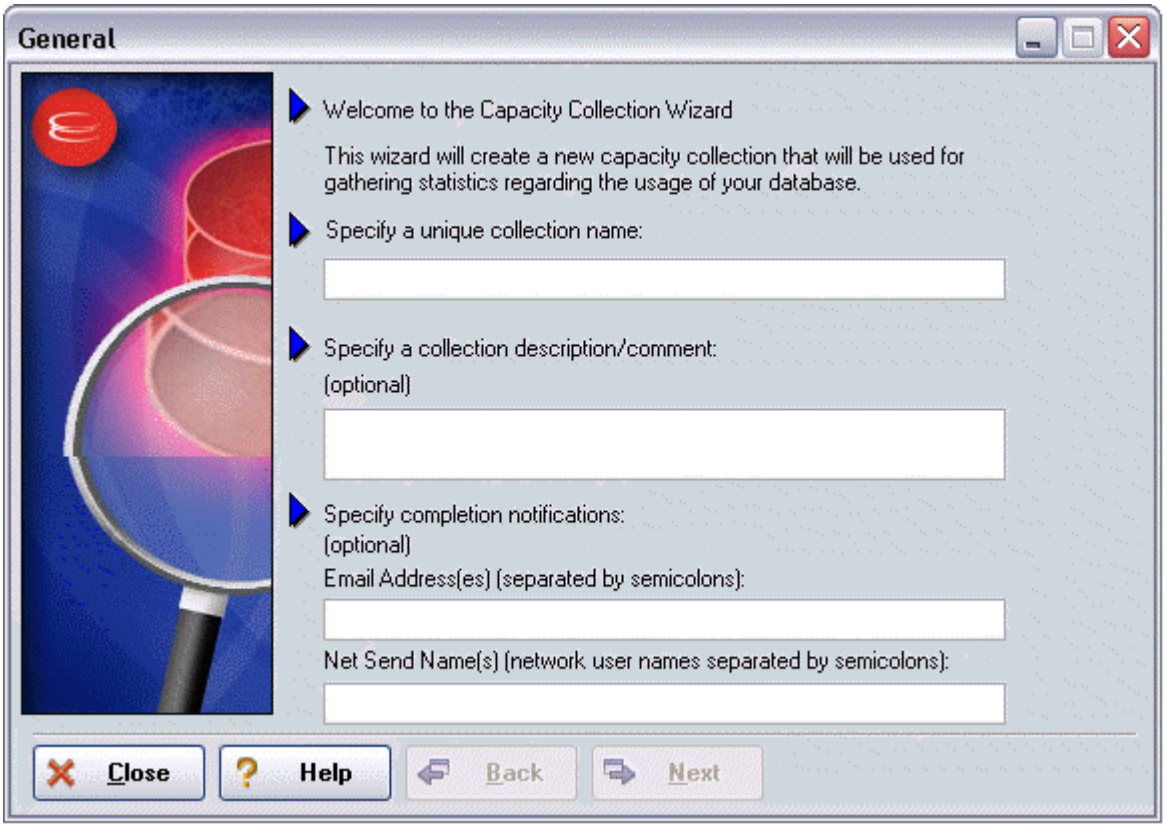

The first panel is devoted to naming your capacity collection and specifying various collection properties like descriptions and optional job completion notifications.

To complete the Wizard:

1 Provide the information in the fields provided and click the **Next** button to continue.

Capacity Analyst opens the **Selection Types** panel of the Capacity Collection Wizard. The second panel asks you what type of information you would like to collect for your database. While this panel prompts you to choose general categories, subsequent panels will let you provide more specific information about the choices you make here.

The specific options differ according to the type of DBMS datasource you are connected to, but in general the categories are:

- Storage structures - Depending on the DBMS you are working with, your analysis and reports can contain space and usage statistics on items such as databases and database devices, logs and backup space, partitions, tablespaces, filegroups, and files.
- DB objects - Your analysis and reports can contain space and usage statistics on tables, indexes, and for some supported DBMS types, other database objects.
- Performance-based statistics - if available for the DBMS you are connected to, this lets you use product-specific performance-based statistic packages.
- Wait events - if available for the DBMS you are connected to, this lets you work with wait-based event information provided by the DBMS.

2 For the purposes of this tutorial, select all options on the panel and then click **Next**.

The next set of panels let you select specific details. Each panel corresponds to one of the options you chose on the **Selection Types** panel.

Most panels offer categories that let you drill down by expanding the category to expose specific options. You can then select specific options and move them from the **Available** list to the **Selected** list and then click **Next** to move to the next panel.

3 Complete all panels up until the **Statistics** tab is displayed.

TIP: Since table and index statistics are supported for all DBMS, as an aid in subsequent sessions, select a variety of indexes and tables on the **DB Objects** panel.

The **Statistics** tab prompts you to update the statistics you chose. In order to use the analysis tools offered by Capacity Analyst, you must execute the collection. This performs calculations and rollups on the metrics in the collection, letting you produce up-to-date reports and charts.

4 Look over the options available on the **Statistics** panel, enable statistics update and select one or more update options, and then click **Next**.

The AutoPurge panel is displayed. It offers the option to delete older data, each time that you execute a collection.

5 Enable AutoPurge and click the **Finish** button.

Proceed to Scheduling a Collection in Capacity Analyst

Scheduling a Collection in Capacity Analyst

Immediately after your collection has been saved to the repository, Capacity Analyst prompts you if you would like to schedule the collection to run at specified intervals. It is not necessary that you schedule your collection to actually collect statistics in the repository. You can execute a collection in an ad-hoc manner at any time through the Capacity Collection Manager.

Depending on your requirements however, scheduling a collection may be the preferred route. For example, if you want to run packaged reports on a scheduled basis, scheduling collection of your collections will ensure up-to-date statistics in your scheduled reports.

For the purposes of this tutorial, you want the statistics in your collection updated every hour.

To schedule updates for your new collection:

1. Select **Yes** when asked if you would like to schedule your new collection

   The Microsoft Task Manager scheduler opens.

2. Click **Advanced** to open the **Advanced Schedule Options** dialog.

3. Select **Repeat task**, specify that the tasks is to be repeated every hour, ensure that the task is repeated until a time at least four hours from now, and click OK to close the **Advanced Schedule Options** dialog.

4. Click **OK** to schedule your collection and dismiss the dialog.

   TIP: You can add, edit, or remove any schedule for any collection at a later time by using the management functions supplied in the Capacity Collection Manager interface.

Proceed to Reviewing other collection management functions

Reviewing other collection management functions

While you have completed all required tasks for Session 1, it is worthwhile looking at some of the options Capacity Analyst offers in creating and updating collections. Later, as you develop a strategy for setting up the collections for your enterprise, you will use additional Capacity Analyst features in managing your collections. With respect to maintaining your collections you have the following options:

Editing collections	After creating a collection, you can modify general settings and the AutoPurge option. Since a collection stores totals and rolled up statistics that can be compromised, you cannot modify your choice of statistics or the objects you generate those statistics against.
Cloning collections	Cloning lets you create a collection based on an existing collection and modify the metrics used and options specified in the original.
Deleting collections	If you no longer have use for a collection or you are going to replace it, you can delete that collection.

Similarly, when you created your collection, you selected the option to update your statistics on an hourly basis. In some cases, you may require ad hoc reports or analysis. If so, you can update your collection manually.

To manually update a collection:

1. In the **Capacity Collection Manager**, select a collection.

2. On the Capacity Analyst toolbar, click the **Execute** button.

   

   Depending on the volume of statistics you specified when creating the collection, this may take a few seconds. When complete the **Last Run** field will reflect the update.

There are other functions contained inside the Capacity Collection Manager, however those described above are likely the ones you will use the most often when managing your capacity collections. You will gain experience using these features as you learn more about Capacity Analyst.

Proceed to Session 2: Using analysis and reporting tools

Session 2: Using analysis and reporting tools

After several hourly updates of the statistics in the collection you created in Session 1, you are ready to start using the analysis tools. In Session 2, you will experiment with the Collection Analysis Viewer and generate one of the Capacity Analyst predefined reports.

Proceed to Performing Trend Analysis Using Capacity Analyst

Performing Trend Analysis Using Capacity Analyst

The Collection Analysis Viewer presents selected statistics in a graphical format and lets you view historical trends and extrapolate that data into the future.

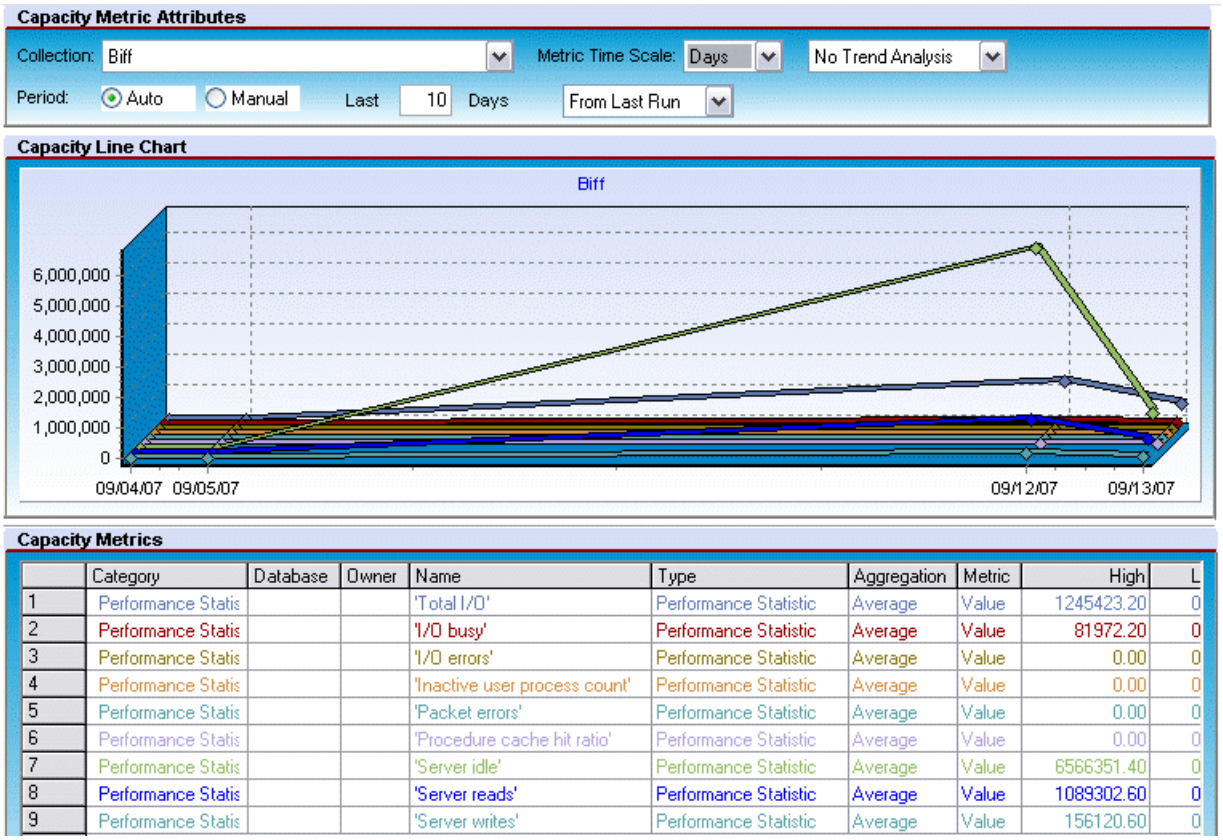

	Category	Database	Owner	Name	Type	Aggregation	Metric	High	L
1	Performance Statis			'Total I/O'	Performance Statistic	Average	Value	1245423.20	0
2	Performance Statis			'I/O busy'	Performance Statistic	Average	Value	81972.20	0
3	Performance Statis			'I/O errors'	Performance Statistic	Average	Value	0.00	0
4	Performance Statis			'Inactive user process count'	Performance Statistic	Average	Value	0.00	0
5	Performance Statis			'Packet errors'	Performance Statistic	Average	Value	0.00	0
6	Performance Statis			'Procedure cache hit ratio'	Performance Statistic	Average	Value	0.00	0
7	Performance Statis			'Server idle'	Performance Statistic	Average	Value	6566351.40	0
8	Performance Statis			'Server reads'	Performance Statistic	Average	Value	1089302.60	0
9	Performance Statis			'Server writes'	Performance Statistic	Average	Value	156120.60	0

To open the Collection Analysis Viewer, do the following:

1. Start DBArtisan again, and connect to the datasource you created earlier. If you need help, refer to Starting DBArtisan and connecting to a datasource.
2. On the **Analyst** menu, select **Capacity Analyst**.
3. Select the Capacity Collection that you recently created.
4. On the Capacity Analyst toolbar, click the **Collection Analysis Viewer** button.

Proceed to Choosing the charted metrics you want to view

Choosing the charted metrics you want to view

The first time that you open the Collection Analysis Viewer, it prompts you select the metrics you want to work with. You specify the metrics using the displayed **Capacity Metric** dialog. For example:

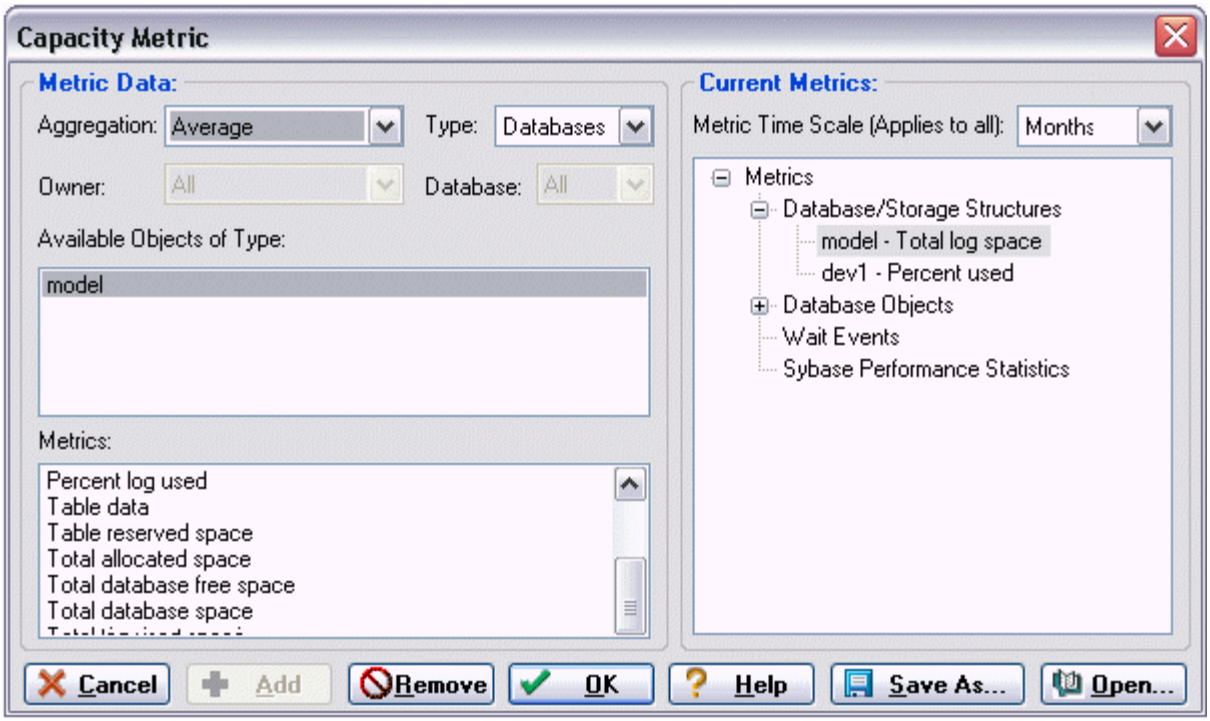

To specify the metrics that are to appear in the analysis:

1. From the **Aggregation** dropdown, select an aggregation type: **Average**, **Maximum**, or **Minimum**.
2. From the **Type** dropdown, select the type of metric you want to add, such as a table or index.

   NOTE: Certain Metric types allow you to further qualify using the **Owner** and **Database** dropdowns.

3. From the **Available Object of Type** list, choose a specific object or statistics type.

   The **Metrics** list is updated with the metrics available for the **Type** and **Available Objects of Type** selections you made.

4. Select a metric from the **Metrics** list and click **Add**.

   Capacity Analyst adds the metric to the **Current Metrics** list.

5. To add additional metrics, repeats step 1 through step 4. Keep in mind that the Collection Analysis VIew will calibrate its chart according to the metric with the largest value. If you are going to add multiple metrics, choose metrics that you expect will have similar ranges.
6. When finished adding metrics for the graphical portion of your analysis, click **OK**.

Proceed to Viewing Trends in Capacity Analyst

Viewing Trends in Capacity Analyst

When you finish adding metrics to the Collection Analysis Viewer, Capacity Analyst presents the data in graphical and grid form in the Collection Analysis Viewer. Your first step should be to ensure that the displayed timescale reflects the fact that you have been updating metrics for a matter of hours. You do this using the controls in the **Capacity Metric Attributes** area.

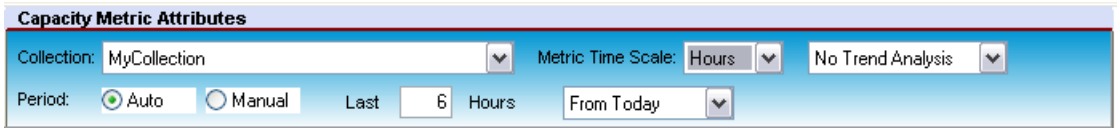

To optimize viewing of the metrics in your collection:

1 From the **Metric Time Scale** dropdown, choose **Hours**.

2 In the **Last** box, enter the number of hours since you scheduled your collection for hourly updates.

Having only statistics that cover a matter of hours, your viewing options are currently somewhat restrictive. However, during a later session you may want to experiment with the following viewing options:

- By default, the **No Trend Analysis** option is selected. This option shows historical data only. If you want statistics extrapolated into the future, you can use the **Trend to Date** or **Trend to Value** options to help you predict potential problems.
- You are currently using the **Auto** Period. It lets you quickly select an interval from the current day or the last statistics update. The **Manual** option lets you to have more granular control over the exact dates/times that Capacity Analyst uses to plot the data.
- To display the actual data values, position the pointer over the trend lines in the graph. Pop up values display the exact points on the chart.
- Each metric's statistic displays in the Analysis viewer's bottom grid with the high, low, and average value for the time period being shown for each metric.
- After your data has been displayed, you can continue to add or remove metrics from the Collection Analysis Viewer by using the Add or Remove toolbar buttons, and perform other functions using the various options available on the Collection Analysis viewer.

3 When ready to proceed to the next exercise, on the Capacity Analyst toolbar click the **Close the Active Tab** button.

Proceed to Using Capacity Analyst predefined reports

Using Capacity Analyst predefined reports

An alternative forecasting method is through use of predefined Analysis reports. Capacity Analyst provides reports on growth trends of devices, databases, tables, rows, and other objects, as well as reports on fragmentation details and failure prediction. Like the COllection Analysis View, the reports show historical data and can be configured to extrapolate to a future date.

To generate a report on the fastest growing database objects:

1. Ensure that the collection you created is selected.
2. Click the **Report** button.

   

   The Capacity Analyst opens the **Select Report Type** dialog box.
3. To create a pre-defined trend analysis report, select **Analysis Report** and click **OK**.

   The Capacity Analyst opens the **Capacity Analyst Report Wizard**.
4. Accept the default to create a **New Analysis Report**.
5. Click **Next**.
6. From the report drop down list, select **'N" Fastest Growing Objects**. Note that there are many other reports available.
7. Accept the default **Auto Entry** time period of 1 day.
8. Accept the **Day** default rollup.
9. Specify the future date you want Capacity Analyst to use for its future 'stop' date.
10. Click **Next**.
11. Specify what databases/objects to include in the report.
12. Click **Next**.
13. Specify a **Report File Name** and a **Report Title**.
14. Click **Execute**.

    Capacity Analyst opens the report in a new tab and stores an HTML-based version of the report in the location you specified in the **Report File Name** field.

The HTML report can be viewed in any web browser so you can build internal web sites of your forecasts for all interested parties, including project leads, and operation managers.

Using Capacity Analyst

The general categories of tasks performed using the Capacity Analyst are:

- Creating and editing collections - Each collection stores a set of statistics. How you set up your collections will depend on your own reporting and monitoring requirements. Capacity Analyst lets you choose the statistics in your collections and control options such as automatic purging and archiving of data.
- Updating statistics in collections for use in reports and charts - In order to produce up-to-date reports and charts, the rolled up data in collections must be updated. Capacity Analyst lets you update your collections manually or to schedule regular, automated updates.
- Viewing charts and reports on collections - The charts and reports available from the Capacity Analyst let you view trends and help you make forecasts regarding growth of your databases. The Collection Analysis Viewer lets you view graphical representation of the metrics in a collection as a function of time. Standard reports provide information on growth trends for devices, databases, tables, rows, and other objects, as well as fragmentation details and failure prediction.

Creating and editing collections

Typically, how you set up collections will be dictated by your own requirements. If your enterprise has several applications for example, each using multiple databases, you might be interested in viewing graphs and reports on an application-by-application basis, with each collection comprising statistics for several databases. At a lower level, you may be interested in viewing statistics on collections of table or indexes.

Similarly, how you set up collections can also be impacted by Capacity Analyst functionality. For example, a collection option lets you schedule statistics update at specified intervals. If your requirements include reports with different intervals, hourly as opposed to weekly for example, you could set up your collections accordingly. Or since the Collection Analysis Viewer calibrates a graph to the range of the highest valued statistic displayed, you might group statistics with similar expected values.

In short, you should consider all functionality available in Capacity Analyst before deciding on a collection setup. For details see Content and options in a collection and Viewing charts and reports on collections.

The following topics provide details on creating and editing Capacity Analyst collections:

- Creating a new collection
- Editing a capacity collection
- Cloning a collection
- Deleting a collection

Content and options in a collection

The reports and analysis you work with using Capacity Analyst include such statistics as space usage and access count details on database objects, databases, files and other storage mechanisms. They can also include performance-based and wait based statistics. The specific content and options available depend on the DBMS that DBArtisan is working with. The following topics provide summaries of available options and content on a product-by-product basis:

- IBM DB2 for Windows, Unix, and Linux capacity collection content and options
- Microsoft SQL Server capacity collection content and options
- Oracle capacity collection content and options
- Sybase capacity collection content and options

IBM DB2 for Windows, Unix, and Linux capacity collection content and options

When creating or editing a Capacity Analyst collection, you specify the statistics you want in that collection and specify additional notification and update options. The table below summarizes the available statistics and options on the tabs/panels of collection wizards and editors for this DBMS.

Content/Option Category	Editor Tab or Wizard panel	Description or specific DBMS feature support
Identification/Notification	**General**	Lets you provide a name and description for the collection and optionally, provide an email address or net send name for notifications.
Types of statistics collected	**Selection Types**	Lets you select the types of content and options you want to include in the collection. You provide specifics in subsequent panels/tabs.
Database/storage structure space statistics	**Storage Structures**	A DB2 capacity collection can contain usage statistics on database (partition) space, DMS tablespace space, and SMS tablespace space.

Content/Option Category	Editor Tab or Wizard panel	Description or specific DBMS feature support
Specific database object types for which to collect statistics	**DB Objects**	A DB2 capacity collection can contain usage statistics on indexes and tables
Update and rollup options	**Statistics**	A DB2 capacity collection offers the option to update statistics with additional specific choices for tables (distribution and column options), indexes (collection of extended and sample statistics), and access options (allowing read/write access during collection).
Automatic purging of old data	**AutoPurge**	Offers the option to automatically purge data older than a specified age, each time the collection is updated.

For information on how to create and edit collections, see:

- Creating a new collection
- Editing a capacity collection
- Cloning a collection

Microsoft SQL Server capacity collection content and options

When creating or editing a Capacity Analyst collection, you specify the statistics you want in that collection and specify additional notification and update options. The table below summarizes the available statistics and options on the tabs/panels of collection wizards and editors for this DBMS.

Content/Option Category	Editor Tab or Wizard panel	Description or specific DBMS feature support
Identification/Notification	**General**	Lets you provide a name and description for the collection and optionally, provide an email address or net send name for notifications.
Types of statistics collected	**Selection Types**	.Lets you select the types of content and options you want to include in the collection. You provide specifics in subsequent panels/tabs.
Database/storage structure space statistics	**Storage Structures**	An SQL Server capacity collection can contain size statistics on database backup space and database log backup space, and usage and size statistics on database space, filegroups, and files.
Specific database object types for which to collect statistics	**DB Objects**	An SQL Server capacity collection can contain usage and size statistics on indexes and tables.
Wait-event-based statistics	**Wait Events**	An SQL Server capacity collection can contain statistics on SQL Server Wait Types
Performance-based statistics	**Performance Statistics**	Provides statistics in the following performance counter categories: Access Methods, Buffer Manager, Buffer Partition, Databases, General Statistics, Latches, Locks, Memory Manager, SQL Statistics
Update and rollup options	**Statistics**	Provides the option to update statistics with additional options to specify a scan range, specify index or columns update, and recompute statistics
Automatic purging of old data	**AutoPurge**	Offers the option to automatically purge data older than a specified age, each time the collection is updated.

For information on how to create and edit collections, see:

>

- Creating a new collection
- Editing a capacity collection
- Cloning a collection

Oracle capacity collection content and options

When creating or editing a Capacity Analyst collection, you specify the statistics you want in that collection and specify additional notification and update options. The table below summarizes the available statistics and options on the tabs/panels of collection wizards and editors for this DBMS.

Content/Option Category	Editor Tab or Wizard panel	Description or specific DBMS feature support
Identification/Notification	**General**	Lets you provide a name and description for the collection and optionally, provide an email address or net send name for notifications.
Types of statistics collected	**Selection Types**	Lets you select the types of content and options you want to include in the collection. You provide specifics in subsequent panels/tabs.
Database/storage structure space statistics	**Storage Structures**	All tablespace, all non-temporary tablespaces, or select specific tablespaces
Specific database object types for which to collect statistics	**DB Objects**	Lets you include statistics on clusters, indexes, index partitions, rollback segments, tables, and table partitions.
Wait-event-based statistics	**Wait Events**	Lets you include statistics based on Oracle Wait Events.
Performance-based statistics	**Global Stats**	Lets you include Global Database Statistics.
Update and rollup options	**Statistics**	Offers the option to update statistics with additional options to compute or estimate statistics.
Automatic purging of old data	**AutoPurge**	Offers the option to automatically purge data older than a specified age, each time the collection is updated.

For information on how to create and edit collections, see:

- Creating a new collection
- Editing a capacity collection
- Cloning a collection

Sybase capacity collection content and options

When creating or editing a Capacity Analyst collection, you specify the statistics you want in that collection and specify additional notification and update options. The table below summarizes the available statistics and options on the tabs/panels of collection wizards and editors for this DBMS.

Content/Option Category	Editor Tab or Wizard panel	Description or specific DBMS feature support
Identification/Notification	**General**	Lets you provide a name and description for the collection and optionally, provide an email address or net send name for notifications.
Types of statistics collected	**Selection Types**	Lets you select the types of content and options you want to include in the collection. You provide specifics in subsequent panels/tabs.

>

Content/Option Category	Editor Tab or Wizard panel	Description or specific DBMS feature support
Database/storage structure space statistics	**Storage Structures**	Lets you use Global Space Summary statistics for the datasource or space and usage statistics for specific database devices and databases.
Specific database object types for which to collect statistics	**DB Objects**	Lets you include statistics on indexes and tables.
Wait-event-based statistics	**Wait Events**	Wait event statistics are only available if MDI performance views have been created.
Performance-based statistics	**Performance Statistics**	Lets you include Sybase Performance Statistics in a collection.
Update and rollup options	**Statistics**	Option to update statistics with additional options to update indexes or partitions
Automatic purging of old data	**AutoPurge**	Offers the option to automatically purge data older than a specified age, each time the collection is updated.

For information on how to create and edit collections, see:

- Creating a new collection
- Editing a capacity collection
- Cloning a collection

Creating a new collection

Capacity Analyst has a WIzard that walks you through specifying the statistics and options in a collection. Before creating a new collection, you should be familiar with the material in Content and options in a collection.

To create a new collection:

1. Open the Capacity Analyst. See Opening the Capacity Analyst.
2. On the Capacity Analyst toolbar click the **New Collection Wizard** button.

   

   The Capacity Collection Wizard opens.
3. Provide the information required on each panel of the Wizard. For more information on these options, see the relevant DBMS-specific topic under Content and options in a collection.
4. When done click **Finish**.

   Capacity Analyst offers the option to schedule updates for the statistics in this collection.
5. If you decide to schedule statistics updates, use the Windows Task Scheduler to provide update details.

   Capacity Analyst adds the new collection to the **Capacity Collection** list.

See the following tasks for information on tasks you can perform after creating a collection:

- Editing a capacity collection
- Updating statistics in collections for use in reports and charts
- Viewing charts and reports on collections

Editing a capacity collection

Because a collection stores rolled up statistics, all gathered and calculated over a specific interval, editing statistics in a collection could corrupt the data in a collection. Therefore, you cannot modify the statistics in a collection or the devices, databases or database objects that statistics are gathered against. You can only modify the basic identifiers for the collection and the refresh and autopurge options.

To edit an existing collection:

1. Open the Capacity Analyst. See Opening the Capacity Analyst.
2. Select the collection to be edited and then on the Capacity Analyst toolbar click the **Open Collection** button.

   

3. On the **General** tab, edit the name, description, email address, or net send name.
4. On the **Statistics** tab, enable or disable update of statistics when you save your changes.
5. On the **AutoPurge** tab, enable or disable autopurge of older data when you save your changes.
6. Click **OK**.

Cloning a collection

The Capacity Analyst lets you create a new collection from an existing collection. This lets you build a new collection starting from properties defined in the existing collection.

To make a copy of an existing collection:

1. Open the Capacity Analyst. See Opening the Capacity Analyst.
2. Select the collection to be cloned and then on the Capacity Analyst toolbar click the **Clone Collection** button.

   

   Capacity Analyst opens the Capacity Analyst Collection Wizard. With the exception of the collection name, all statistics and options specified for the new collection are identical to those for the existing collection.
3. Minimally, provide a name for the new collection. Complete each panel of the wizard, changing settings as required. For more information on these options, see the relevant DBMS-specific topic under Content and options in a collection.
4. When done click **Finish**.

   Capacity Analyst offers the option to schedule updates for the statistics in this collection.
5. If you decide to schedule statistics updates, use the Windows Task Scheduler to provide update details.

   Capacity Analyst adds the new collection to the **Capacity Collection** list.

See the following tasks for information on tasks you can perform after cloning a collection:

- Editing a capacity collection
- Updating statistics in collections for use in reports and charts
- Viewing charts and reports on collections

Deleting a collection

If you no longer use the reports or analysis associated with a collection, Capacity Analyst lets you delete the collection.

To delete a collection:

1. Open the Capacity Analyst. See Opening the Capacity Analyst.
2. Select a collection and then on the Capacity Analyst toolbar click the **Delete** button.

   

   Capacity Analyst prompts you to verify that you want to delete the collection.
3. Click **Yes**.

Updating statistics in collections for use in reports and charts

Once you have installed and set up Capacity Analyst on a server, raw statistics are gathered and stored in the repository on that server. In order to view up-to-date charts and reports, you must execute a collection. Executing the collection calculates totals and performs all rollups for that collection in preparation for reporting.

There are two methods of updating statistics in a collection:

- Manually updating collections, useful prior to running an ad hoc report
- Scheduling updates of collections to run regularly, useful when reporting is performed on a regular basis

Manually updating collections

You can manually execute a statistics collection. This calculates totals and statistical indicators such as averages used for reporting for that collection.

To manually execute a collection:

1. Open the Capacity Analyst. See Opening the Capacity Analyst.
2. Select a collection and then on the Capacity Analyst toolbar click the **Execute** button.

   

   When complete, the **Last Run** field reflects the time that the collection was updated.

After executing a collection, you can run reports or view charts for that collection, that are up-to-date as of the time that you executed the collection. For details, see Viewing charts and reports on collections.

Related topics:

- Creating and editing collections

Scheduling updates of collections

If you run reports or want to regularly archive statistics for intervals, you can schedule execution of statistics collections. This updates the totals, rollups, and statistical indicators such as averages, on a regular basis. You can schedule collection updates to occur regularly such as daily weekly, or monthly, or tie execution to events such as system startup, logon, or during idle time.

>

The following topics provide details on the scheduling tasks you can perform against a collection:

- Creating an update collection schedule
- Editing an update collection schedule
- Deleting an update collection schedule

Creating an update collection schedule

You can create an update collection schedule to have statistics in that collection updated at regular intervals or have updates triggered by events such as system startup or logon.

In setting up your schedule, keep the following points in mind:

- Some statistical collections depend on the underlying dynamics of the database objects being collected. For example, storage structure space metrics, table and index statistics, etc., are oftentimes best collected only once a day. Other statistics, such as wait events and global database metrics can be collected on a much more granular level (every hour, etc.)
- If you routinely stop and start your database at a certain time, you may want to ensure that your collection job is not scheduled to run during its downtime. Also realize that stopping and starting a database resets the counters used for wait events and global database metrics.
- If you use the Microsoft scheduler on your machine, remember to not turn your machine off during the times your collection is scheduled to run or the job will not run.

  TIP: You can add, edit, or remove any schedule for any collection at a later time by using the management functions supplied in the Capacity Collection Manager interface.

To schedule updates for a collection:

1. Open the Capacity Analyst. See Opening the Capacity Analyst.
2. Select a **Capacity Collection** that currently does not have an associated execution schedule.
3. Click the **Schedule** button.

   

   Capacity Analyst opens the Job Scheduler.
4. Use the Job Scheduler to specify the frequency of the report or the event that triggers the update, and other details of the schedule.
5. When finished, click **OK**.

Related topics:

- Editing an update collection schedule
- Viewing charts and reports on collections

Editing an update collection schedule

After creating an update schedule for a statistics collection, you can revise the schedule. You may wish to change the frequency of updates or the event which triggers an update, or more specific details of the update schedule.

>

To modify the update schedule for a collection:

1. Open the Capacity Analyst. See Opening the Capacity Analyst.
2. Select a **Capacity Collection** that currently has an associated execution schedule.
3. Click the **Edit Schedule** button.

   

   Capacity Analyst opens the Job Scheduler.
4. Use the Job Scheduler to change the frequency of the report or the event that triggers the update, and other details of the schedule.
5. When finished, click **OK**.

Related topics:

- Scheduling updates of collections
- For information on modifying the statistics in a collection, see Editing a capacity collection
- Viewing charts and reports on collections

Deleting an update collection schedule

If you no longer want statistics in a collection to be calculated on a regular basis or if you want to create a new, completely reworked schedule, you can delete the existing schedule for a collection.

To unschedule a Capacity Collection:

1. Open the Capacity Analyst. See Opening the Capacity Analyst.
2. Select a **Capacity Collection**.
3. On the Capacity Analyst toolbar, click the **Unschedule** button.

   

   Capacity Analyst unschedules the Capacity Collection.

Related topics:

- Scheduling updates of collections
- Creating an update collection schedule

Purging old data in a collection

Policy or practicality may dictate that the metrics in your analysis and reports not include data older than a specified date. Capacity Analyst lets you manually purge data older than a specified number of months or years.

NOTE: In addition to manually purging old data, when creating or editing a collection, you can set up the collection to automatically purge data older than a specified date on every update. For details, see Creating and editing collections.

To manually purge collection data older than a specified number of months or years:

1. Open the Capacity Analyst. See Opening the Capacity Analyst.
2. Select a **Capacity Collection**.
3. On the toolbar, click the **Purge** button.

   

   The **Perform Manual Data Purge** dialog opens.
4. Use the controls in the **Purge data older than** area to provide a specific number of months or years.
5. Click **Execute** and when prompted to confirm, click **Yes**.

Viewing charts and reports on collections

The preliminary steps in setting up Capacity Analyst are Creating and editing collections and Updating statistics in collections for use in reports and charts. Once your collections are defined and the statistics updated, Capacity Analysts lets you view forecasting/trends charts and reports on your collections. Key options are:

- Using the Collection Analysis Viewer - lets you view charted representations of collection metrics as a function of time.
- Working with Capacity Analyst reports - lets you work with predefined reports that focus on growth of databases, storage mechanisms, and specific database objects, as well as failure prediction and fragmentation.

Using the Collection Analysis Viewer

The Collection Analysis Viewer lets you perform ad hoc analysis of statistics in a collection. It presents the metrics in two formats: a graphical representation that shows values over time and a tabular format that includes metrics such as average, maximum, and minimum values.

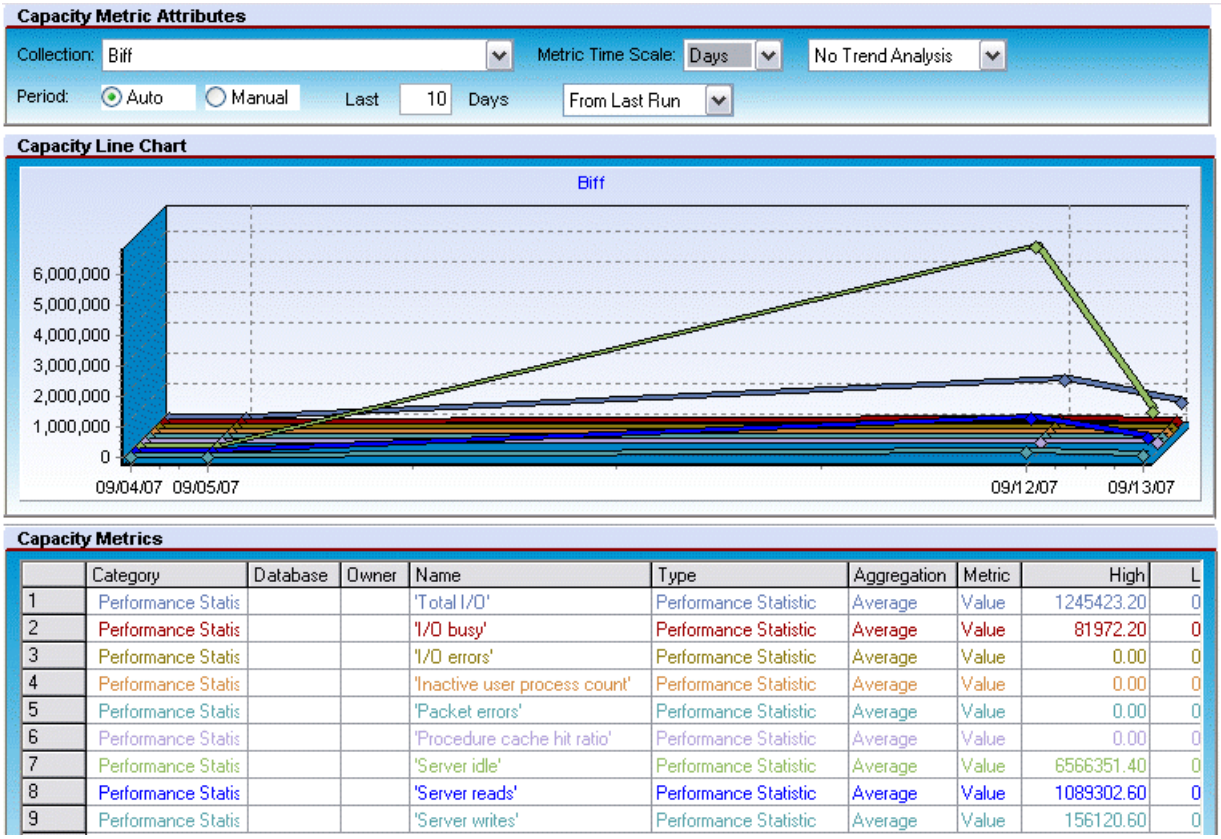

	Category	Database	Owner	Name	Type	Aggregation	Metric	High	L
1	Performance Statis			'Total I/O'	Performance Statistic	Average	Value	1245423.20	0
2	Performance Statis			'I/O busy'	Performance Statistic	Average	Value	81972.20	0
3	Performance Statis			'I/O errors'	Performance Statistic	Average	Value	0.00	0
4	Performance Statis			'Inactive user process count'	Performance Statistic	Average	Value	0.00	0
5	Performance Statis			'Packet errors'	Performance Statistic	Average	Value	0.00	0
6	Performance Statis			'Procedure cache hit ratio'	Performance Statistic	Average	Value	0.00	0
7	Performance Statis			'Server idle'	Performance Statistic	Average	Value	6566351.40	0
8	Performance Statis			'Server reads'	Performance Statistic	Average	Value	1089302.60	0
9	Performance Statis			'Server writes'	Performance Statistic	Average	Value	156120.60	0

The Collection Analysis Viewer lets you:

- View historical trends for a metric and extrapolate into the future
- Compare and contrast multiple metrics
- Vary the time scale for fine-grained or coarse-grained views

The first time that you open the Collection Analysis Viewer on a particular collection, you are prompted to select the metrics you want to view for that collection. Subsequently, when using the Collection Analysis Viewer against that collection, you can add metrics to and delete metrics from the collection.

To open the Collection Analysis Viewer on a collection:

1 Open the Capacity Analyst. See Opening the Capacity Analyst.

2 Select a Capacity Collection. For more information, see Creating and editing collections.

3 On the toolbar, click the **Collection Analysis Viewer** button.

NOTE: If this is the first time you have opened this particular collection, you are prompted to specify the statistics to display. For help with this step, see Adding metrics to the Collection Analysis Viewer for a collection.

Capacity Analyst opens the Collection Analysis View or the selected collection.

Viewing options available when working with the Collection Analysis Viewer include:

- Using the Delete Metric button to remove a statistic from the displayed metrics for a collection

- Using the zoom buttons on the Collection Analysis Viewer toolbar to zoom in and out on the report

- Using the refresh button on the Collection Analysis Viewer toolbar to update displayed values

- Using the **Collection** dropdown to open the Collection Analysis Viewer on a different collection
- Using the **Metric Time Scale** dropdown to choose between displayed time scales of hours, days, months, or years
- Using the trending and value dropdowns to extrapolate metric values out to specified values or specific dates
- Using the **Period** controls to change the date/time range displayed
- Viewing values for metrics at specific locations on the graph by floating the cursor over specific locations on the curve

Adding metrics to the Collection Analysis Viewer for a collection

There are two situations in which you can add metrics viewed using the Collection Analysis Viewer:

- The first time that you open the Collection Analysis Viewer for a given collection
- When you want to manually add metrics by clicking the Collection Analysis Viewer toolbar's Add Metric button:

>

In both cases, Capacity Analyst opens a **Capacity Metric** dialog to let you add metrics.

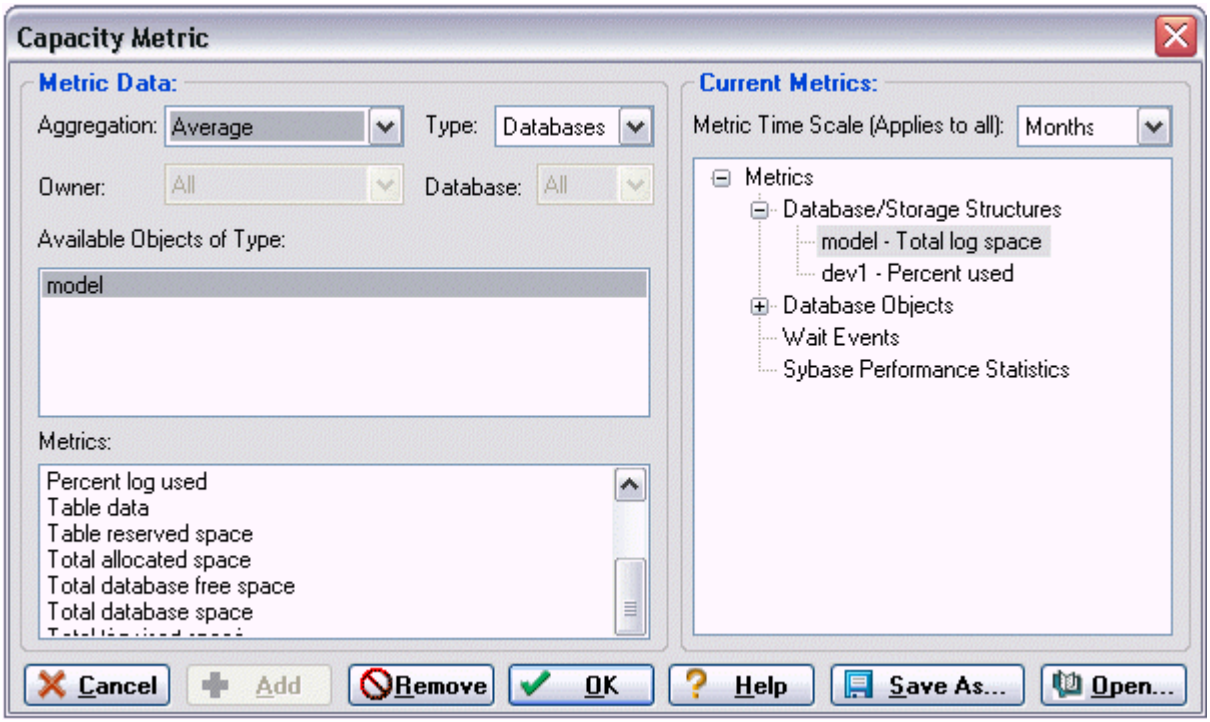

For each metric you want add to the Capacity Analyst Analysis Viewer:

1. From the **Aggregation** dropdown, select an aggregation type: **Average**, **Maximum**, or **Minimum**.
2. From the **Type** dropdown, select the type of metric you want to add.

   NOTE: Certain Metric types allow you to further qualify using the **Owner** and **Database** dropdowns.

3. From the **Available Object of Type** list, choose a specific object or statistics type.

   The **Metrics** list is updated with the metrics available for the **Type** and **Available Objects of Type** selections you made.

4. Select a metric from the **Metrics** list and click **Add**.

   Capacity Analyst adds the metric to the **Current Metrics** list.

5. To add additional metrics, repeats step 1 through step 4.
6. When finished adding metrics for the Capacity Analyst Collection Analysis Viewer, click **OK**.

Capacity Analyst updates the Collection Analyst Viewer with the new metrics.

Working with Capacity Analyst reports

Capacity Analyst provides two report options. For details, see the following topics:

- Running a Capacity Collection Summary Report - lets you view a summary of the metrics and options currently defined for a collection.
- Running an Analysis Report - lets you view one of several predefined reports focusing on growth trends

Running a Capacity Collection Summary Report

You can produce a high-level summary of the content and options associated with a collection. It includes:

- Basic identification information such as the name, associated datasource, and owner of the collection
- A listing of the options, refresh and auto-purge for example, chosen for that collection
- A summary of the types of statistics available for use in the collection
- A detailed listing of the objects for which statistics are collected

NOTE: For a detailed listing of the specific statistics and options available for the DBMS you are working with, see the DBMS-specific heading under Content and options in a collection.

To run a Capacity Collection Summary Report:

1. Open the Capacity Analyst. See Opening the Capacity Analyst.
2. Select a Capacity Collection. For more information, see Creating and editing collections.
3. On the Capacity Analyst toolbar, click the **Report** button.

   

   Capacity Analyst opens the **Select Report Type** dialog.
4. Select **Capacity Collection Summary Report** and click OK.

Capacity Analyst opens the summary report in a new tab.

Running an Analysis Report

Capacity Analyst lets you run a set of predefined reports against a collection. The reports and their availability by DBMS are:

- Database growth trends (Sybase, SQL Server, DB2)
- Tablespace failure prediction (Oracle)
- Tablespace growth trends (DB2, Oracle)
- Device growth trends (Sybase)
- Filegroup growth trends (SQL Server)
- File growth trends (SQL Server)
- Object growth trends (Sybase, SQL Server, DB2, Oracle)
- Object extent trends (Sybase, SQL Server)
- Object extent failure prediction
- Database backup growth trends (SQL Server)
- ‘N’ Fastest growing databases (Sybase, SQL Server, DB2)
- ‘N’ Fastest growing tablespaces (Oracle)
- ‘N’ Fastest growing objects (Sybase, SQL Server, DB2, Oracle)
- Table rows trend (Sybase, SQL Server, DB2, Oracle)
- Table chained/migrated rows trend (Oracle)

- Table forwarded records trend (Sybase, SQL Server)
- Index level trend (Oracle)
- Object fragmentation trends (Sybase, SQL Server, DB2)
- Tablespace fragmentation trend (Oracle)
- Database failure prediction (Sybase, SQL Server)
- DMS tablespace failure trends (Oracle)

To run an analysis report:

1 Open the Capacity Analyst. See Opening the Capacity Analyst.

2 Select a Capacity Collection. For more information, see Creating and editing collections.

3 On the Capacity Analyst toolbar, click the **Report** button.

Capacity Analyst opens the **Select Report Type** dialog.

4 Select **Analysis Report** and click OK.

Capacity Analyst opens the Capacity Analyst Report Wizard.

5 Use the following table as a guide to understanding the settings and options for this wizard

Panel	Tasks and settings	Description
Panel 1	**Do you wish to create a new analysis report or 'playback' a previously saved report definition?**	As the last step in generating a report, you are offered the option to save the report definition, for the purpose of running the report again. If you have already saved report definitions, you can use this panel to enable the **Previously Saved Report** option and select the existing report definition. Otherwise, to run an ad hoc report or start a new report definition, you can use the **New Analysis Report** option.
Panel 2	**What report would you like to create?**	Lets you select one of the predefined report types listed above.
	Enter the time period sample to use for the report	**Auto Entry** lets you specify a number of days, months or years while **Manual Entry** lets you provide specific start and end dates.
	Enter the time rollup to use for the report.	Lets your report display statistics rolled up by day, month, or year.
	Enter the date to use for forward predictions.	Lets you enter a future end date that the Capacity Analyst uses to forecast.
	Select options for the report.	The options offered differ by the type of report you selected and the DBMS that DBArtisan is working with.
Panel 3	**Include objects for report**.	Expand or collapse databases or owner lists to display specific object types and use the arrow buttons to move selected objects between the **Available** and **Selected** lists.
Panel 4	**Data will be presented in text/grid format.**	Lets you specify file name and path where the report will be stored and provide a title for the report.
	Save this report definition?	Lets you save the report definition, specifying the file name and path for the definition.

6 When complete, click **Schedule** or **Execute** to run the report.

>

When complete, the report opens in a new DBArtisan tab.

NOTE: When viewing the report in DBArtisan, you can use the toolbar's **Refresh** button to update values displayed.

A static, HTML version is also saved to with the file name and path you specified in the **Report File Name** field.

Performance Analyst

Performance Analyst is a robust client-side monitoring product that presents a complete picture of a database's performance. Performance Analyst is a quick, lightweight tool that lets you investigate performance bottlenecks.

The table below describes the major sections of Help.

Section	Description
Introduction	Provides information about installing, opening, and licensing Performance Analyst.
Application Basics	Provides conceptual information about Performance Analyst's product design and navigation.
Using Performance Analyst	Shows how to use each component of the application. Using Performance Analyst breaks each process and procedure into simple step-by-step instructions.
Performance Analyst for Oracle Tutorial	Helps you get started with Embarcadero Performance Analyst. Provides the foundation you need to explore the many features and benefits of Performance Analyst.
Oracle Performance Analyst statistics	Performance Analyst includes expert help for all supported database platforms.
Performance Analyst for Microsoft SQL Server Tutorial	Helps you get started with Embarcadero Performance Analyst. Provides the foundation you need to explore the many features and benefits of Performance Analyst.
Microsoft SQL Server Performance Analyst statistics	Performance Analyst includes expert help for all supported database platforms.
Performance Analyst for Sybase Tutorial	Helps you get started with Embarcadero Performance Analyst. Provides the foundation you need to explore the many features and benefits of Performance Analyst.
Sybase ASE Performance Analyst statistics	Performance Analyst includes expert help for all supported database platforms
Performance Analyst for IBM DB2 for Linux, Unix, and Windows Tutorial	Helps you get started with Embarcadero Performance Analyst. Provides the foundation you need to explore the many features and benefits of Performance Analyst.
IBM DB2 (W/U/L) Performance Analyst statistics	Performance Analyst includes expert help for all supported database platforms

>

Welcome to Performance Analyst

Performance Analyst is a client-side only visual database monitor that displays current database performance data, including key performance metrics, bottleneck analysis, and user-configurable thresholds.

Running inside the Embarcadero DBArtisan console, or stand-alone in a Web browser, Performance Analyst lets you identify performance problems that threaten database availability or response times.

Supported DBMS

Embarcadero Technologies currently offers Performance Analyst for Oracle 8i - 10g. Performance Analyst does not support Oracle Clients 7.3.4, 8.1.7, and 9.2. Performance Analyst also supports Microsoft SQL Server and DB2.

For more information, see:

- What You Should Know Before Installation
- Operating System Prerequisites
- Opening the Performance Analyst
- Licensing Performance Analyst

What You Should Know Before Installation

When you install the Analyst, several Embarcadero-generated stored objects are pushed onto the database server:

These objects are as follows:

ORACLE PROCEDURES

Procedures

ALL_TABLESPACES	ALL_USERS	HOTTESTDBFILES
INSTANCESTARTUP	IO_DD_ACCESSPATTERNDETAIL	IO_DD_ACTIVEROLLBACKDETAILS
IO_DD_ARCHIVELOGMODE	IO_DD_CURRENTDISKSORTS	IO_DD_DAILYARCHIVELOGSUMMARY
IO_DD_DATAFILEIO	IO_DD_DBWRDETAIL	IO_DD_ESTPHYSICALIOBYTABLE
IO_DD_JOBQUEUESUMMARY	IO_DD_LGWRDETAIL	IO_DD_PROBLEMJOBS
IO_DD_REDOWASTAGE	IO_DD_REDOWASTAGE2	IO_DD_ROLLBACKIO
IO_DD_RUNNINGJOBSDETAIL	IO_DD_SESSIONROLLBACKACTIVITY	IO_DD_SESSIONSORTDETAIL
IO_DD_TABLESPACEIO	MEMORY_DD_BUFFERCACHERATIO	MEMORY_DD_BUFFERCACHEUTIL
MEMORY_DD_BUFFERPOOLHITRAT	MEMORY_DD_BUFFERPOOLMEMCONF	MEMORY_DD_BUFFERPOOLOBJASS
MEMORY_DD_CACHEDBLOCKSBYTS	MEMORY_DD_CACHEDTABLES	MEMORY_DD_DATADICTDETAILS
MEMORY_DD_LATCHDETAILS	MEMORY_DD_LATCHRATIOS	MEMORY_DD_LIBCACHEHITRATIO
MEMORY_DD_LIBCACHEPERF	MEMORY_DD_PGADETAIL	MEMORY_DD_PGAOVERVIEW
MEMORY_DD_PGAWORKAREA	MEMORY_DD_PINNEDOBJECTS	MEMORY_DD_SHAREDPOOLFRAG
MEMORY_DD_SHAREDPOOLUSAGE	MEMORY_DD_SHAREDPOOLUTIL	NOT_WAITING_SESSIONS_COUNT
OBJECT_DD_GLOBALOBJSUM	OBJECT_DD_HOTCODE	OBJECT_DD_HOTTABLES
OBJECT_DD_INDEXES	OBJECT_DD_INVALIDOBJDET	OBJECT_DD_INVALIDOBJSUM

OBJECT_DD_REORGANALYSIS	OBJECT_DD_TABLESINSAMETS	OBJECT_DD_USEROBJSUM
OBJECTBUFFERPOOLPLACEMENT	RAC_LOCKS_CONVERT	RAC_LOCKS_DLM
RAC_LOCKS_LATCH	RAC_LOCKS_PCM	RAC_MEMORY_ANALYSIS
RAC_MEMORY_DATACACHE	RAC_MEMORY_KEYRATIO	RAC_MEMORY_LIBRARYCACHE
RAC_OVERVIEW_CONFLICTS	RAC_OVERVIEW_INSTANCE	RAC_OVERVIEW_RATIOS
RAC_OVERVIEW_WAITS	RAC_OVERVIEW_WORKLOAD	RAC_PING_DATAFILE
RAC_PING_INSTANCE	RAC_PING_OBJECT	RAC_SESSIONS_BLOCKS
RAC_SESSIONS_TOP	RAC_SESSIONS_USER	RAC_SESSIONS_USERWAITS
RESPONSETIME_CURRENT_RESTIME	RESPONSETIME_CURRENTWORKLOAD	RESPONSETIME_HISTLISTEVENT
RESPONSETIME_HISTORY	RESPONSETIME_HISTORYFILE	RESPONSETIME_HISTORYOBJECT
RESPONSETIME_HISTORYSESSION	RESPONSETIME_HISTORYSQL	RESPONSETIME_HISTORYSUMMARY
RESPONSETIME_LASTHOUR	RESPONSETIME_LASTHOURDETAIL	RESPONSETIME_LISTMETRIC
RESPONSETIME_SESBOTTLENECK	RESPONSETIME_SESBOTTLENECKDTL	RESPONSETIME_SESBOTTLENECKHIST
RESPONSETIME_SESSIONTIME	RESPONSETIME_SYSBOTTLENECKHIS	RESPONSETIME_SYSTEMBOTTLENECK
RESPONSETIME_SYSTEMTIME	RESPONSETIME_WAITHISTOGRAM	RESPONSETIME_WAITHISTOLIST
RETRIEVEHOMEPAGE	RETRIEVEHOMERATIOS	RETRIEVEIOPAGE
RETRIEVEMEMORYPAGE	RETRIEVEMEMORYRATIOS	RETRIEVEOBJECTPAGE
RETRIEVEOBJECTRATIOS	RETRIEVESPACEPAGE	RETRIEVEUSERPAGE
RETRIEVEUSERRATIOS	SESSION_DD_ALLLOCKS	SESSION_DD_BLKLOCKS
SESSION_DD_CRNTWAITS	SESSION_DD_CURRENTSQL	SESSION_DD_CURRENTSQLSTAT
SESSION_DD_HISTWAITS	SESSION_DD_OPENCURSORS	SESSION_DD_OTHERSQL
SESSION_DD_OVERVIEW	SESSION_DD_OVERVIEWSTAT	SESSION_DD_TRANSACTIONS
SESSION_LAST_TEN_WAIT	SESSION_TIME_SUMMARY	SPACE_DD_ACCOBJECTGRID
SPACE_DD_DATAFILEDETAIL	SPACE_DD_DATAFILEFRAG	SPACE_DD_EXTENTANALYSIS
SPACE_DD_SPACEBYOBJECT	SPACE_DD_SPACEBYUSER	SPACE_DD_TABLESPACEFRAG
SPACE_DD_TABLESPACEGROWTH	SPACE_DD_TABLESPACEMAP	SPACE_DD_TABLESPACEOBJSUM
SPACE_DD_TOPUSERS	SPACEANALYSIS	SQLANALYSIS_HOME
SQLANALYSIS_IO	SQLANALYSIS_MEMORY	SQLANALYSIS_OBJECT
SQLANALYSIS_SPACE	SQLANALYSIS_USER	TABLEHIGHWATERMARKPROBLEMCOUNT
TOP_SYSTEM_BOTTLENECKS_10G	TOPCPUPROCESS	TOPDBFILEWAITS
TOPDISKSORTPROCESS	TOPFRAGMENTEDDATAFILES	TOPIOHOGS
TOPIOPROCESS	TOPLARGESCANPROCESS	TOPLATCHMISSES
TOPLOGFILEWAITS	TOPMEMORYHOGS	TOPMEMORYPROCESS
TOPPARSEPROCESS	TOPSESSIONBOTTLENECKS	TOPSESSIONS
TOPSESSIONS_10G	TOPSPACEHOGS	TOPSQL_DETAILS
TOPSQL_RECENT_HISTORY	TOPSQL_SQLTEXT	TOPSQL_STATS

TOPSQL_SUMMARY_DETAILS	TOPSYSTEMBOTTLENECKS	USERACTIVITYANALYSIS
USERS_DD_ALL_LOCKS	USERS_DD_BLK_LOCKS	USERS_DD_BY_TIME_WAITED
USERS_DD_BY_TOTAL_WAITS	USERS_DD_CRNTWAITS	USERS_DD_HISTWAITS
USERS_DD_HISTWAITS_10G	USERS_DD_SYSTEMWAITS_10G	VERSION
WAITING_SESSIONS_COUNT	WORSTSCHEMAFORCHAINEDROWTABLES	WORSTSCHEMAFOREXTENTPROBLEMS
WORSTSCHEMAFORINVALIDOBJECTS		

SQL SERVER

Procedures

sp_embt_pa_all_database_enumerrorlog	sp_embt_pa_all_database_list	sp_embt_pa_current_sq
sp_embt_pa_database_analysis	sp_embt_pa_database_bottleneck	sp_embt_pa_database_drill_backups_db
sp_embt_pa_database_drill_backups_detail	sp_embt_pa_database_drill_backups_log	sp_embt_pa_database_drill_config_detail
sp_embt_pa_database_drill_config_server	sp_embt_pa_database_drill_error_log_content	sp_embt_pa_database_drill_error_log_detail
sp_embt_pa_database_drill_error_log_volume	sp_embt_pa_database_drill_obj_index_det	sp_embt_pa_database_drill_obj_summ
sp_embt_pa_database_drill_obj_table_det	sp_embt_pa_database_drill_overview	sp_embt_pa_database_drill_replication_agent
sp_embt_pa_database_drill_replication_thruput	sp_embt_pa_database_drill_sqlagent_alrtdet	sp_embt_pa_database_drill_sqlagent_alrtsum
sp_embt_pa_database_drill_sqlagent_jobdet	sp_embt_pa_database_drill_sqlagent_jobsum	sp_embt_pa_database_drill_sqlagent_status
sp_embt_pa_database_errorlog	sp_embt_pa_db_requiring_backup	sp_embt_pa_db_workload_analysis
sp_embt_pa_home_active_processes	sp_embt_pa_home_activity_metrics	sp_embt_pa_home_bottleneck_waits
sp_embt_pa_home_bottlenecks	sp_embt_pa_home_storage_database_count	sp_embt_pa_home_storage_totals
sp_embt_pa_instance_startup	sp_embt_pa_io_access_ratios	sp_embt_pa_io_activity_metrics
sp_embt_pa_io_bottlenecks	sp_embt_pa_io_drill_database	sp_embt_pa_io_drill_file
sp_embt_pa_io_drill_system_awe	sp_embt_pa_io_drill_system_io	sp_embt_pa_io_drill_system_physical
sp_embt_pa_io_drill_system_space	sp_embt_pa_io_drill_user_lock	sp_embt_pa_io_drill_user_obj
sp_embt_pa_io_drill_user_sql	sp_embt_pa_io_drill_user_temp	sp_embt_pa_io_get_file_stats
sp_embt_pa_io_hottest_files	sp_embt_pa_io_ratios	sp_embt_pa_io_top_processes
sp_embt_pa_memory_analysis	sp_embt_pa_memory_drill_bcache	sp_embt_pa_memory_drill_buffer
sp_embt_pa_memory_drill_buffer_page_activity	sp_embt_pa_memory_drill_log_details	sp_embt_pa_memory_drill_log_instances
sp_embt_pa_memory_drill_proc_detail	sp_embt_pa_memory_drill_proc_summary	sp_embt_pa_memory_latch_stats
sp_embt_pa_memory_ratio_mem	sp_embt_pa_memory_sql_analysis	sp_embt_pa_memory_top_latch_waits
sp_embt_pa_memory_workload_analysis	sp_embt_pa_space_bottlenecks	sp_embt_pa_space_drill_filegroups

>

sp_embt_pa_space_drill_files	sp_embt_pa_space_drill_object_summary	sp_embt_pa_space_drill_objects_index_det
sp_embt_pa_space_drill_objects_table_det	sp_embt_pa_space_drill_server_by_db	sp_embt_pa_space_drill_server_summary
sp_embt_pa_space_drill_vlogs	sp_embt_pa_space_free_disk_space	sp_embt_pa_space_main
sp_embt_pa_space_used_disk_space	sp_embt_pa_top_sql_database	sp_embt_pa_top_sql_details
sp_embt_pa_top_sql_home	sp_embt_pa_top_sql_query	sp_embt_pa_top_sql_refresh
sp_embt_pa_top_sql_space	sp_embt_pa_top_sql_start	sp_embt_pa_top_sql_stop
sp_embt_pa_top_sql_trace	sp_embt_pa_user_bottleneck_activities	sp_embt_pa_user_bottleneck_sessions
sp_embt_pa_user_drill_locks_by_db	sp_embt_pa_user_drill_locks_by_locktype	sp_embt_pa_user_drill_locks_detail
sp_embt_pa_user_drill_session_waits_by_time	sp_embt_pa_user_drill_session_waits_by_type	sp_embt_pa_user_drill_session_waits_detail
sp_embt_pa_user_drill_system_waits_by_time	sp_embt_pa_user_drill_system_waits_by_total_waits	sp_embt_pa_user_drill_system_waits_detail
sp_embt_pa_user_drill_top_sessions	sp_embt_pa_user_drill_top_sessions_detail	sp_embt_pa_user_ratio_login
sp_embt_pa_user_ratio_process	sp_embt_pa_user_workload	sp_embt_pa_version

SYBASE

Procedures

sp_embt_pa_actsql_drill	sp_embt_pa_actsql_drill_det	sp_embt_pa_current_sql
sp_embt_pa_db_backup	sp_embt_pa_db_bnecks	sp_embt_pa_db_drill_cb_filter
sp_embt_pa_db_drill_cb_idx	sp_embt_pa_db_drill_cb_tbl	sp_embt_pa_db_drill_cnfg
sp_embt_pa_db_drill_dbs	sp_embt_pa_db_drill_engines	sp_embt_pa_db_drill_obj_filter
sp_embt_pa_db_drill_obj_idx	sp_embt_pa_db_drill_obj_suspct	sp_embt_pa_db_drill_obj_tbl
sp_embt_pa_db_drill_opt	sp_embt_pa_db_drill_overview	sp_embt_pa_db_drill_sus_filter
sp_embt_pa_db_engines	sp_embt_pa_db_max_engines=	sp_embt_pa_db_metadata
sp_embt_pa_db_obj_space	sp_embt_pa_db_srv_bnecks	sp_embt_pa_db_workload
sp_embt_pa_dbspace	sp_embt_pa_home_act_metrics	sp_embt_pa_home_bottlenecks
sp_embt_pa_home_storage_tot	sp_embt_pa_instance_startup	sp_embt_pa_io_anlysis
sp_embt_pa_io_bottlenecks	sp_embt_pa_io_drill_dbs	sp_embt_pa_io_drill_device_rt
sp_embt_pa_io_drill_device_sum	sp_embt_pa_io_drill_device_typ	sp_embt_pa_io_drill_engine
sp_embt_pa_io_drill_net_eng	sp_embt_pa_io_drill_net_ses	sp_embt_pa_io_drill_object
sp_embt_pa_io_hottest_devices	sp_embt_pa_io_key_system	sp_embt_pa_io_network
sp_embt_pa_io_trans_log	sp_embt_pa_io_workload	sp_embt_pa_mem_activity
sp_embt_pa_mem_analysis	sp_embt_pa_mem_bottleneck	sp_embt_pa_mem_eff_ratios
sp_embt_pa_mem_mdcache	sp_embt_pa_mem_workload	sp_embt_pa_memory_drill_dc_act
sp_embt_pa_memory_drill_dc_cfg	sp_embt_pa_memory_drill_dc_eff	sp_embt_pa_memory_drill_dc_obj
sp_embt_pa_memory_drill_pc_act	sp_embt_pa_memory_drill_pc_det	sp_embt_pa_memory_drill_pc_use
sp_embt_pa_monitor	sp_embt_pa_space_bottlenecks	sp_embt_pa_space_db_analysis

sp_embt_pa_space_db_dev_cnt	sp_embt_pa_space_dev_analysis	sp_embt_pa_space_drill_db_sum
sp_embt_pa_space_drill_dev_det	sp_embt_pa_space_drill_dev_sum	sp_embt_pa_space_drill_indexes
sp_embt_pa_space_drill_tables	sp_embt_pa_space_indexs_filte	sp_embt_pa_space_tables_filter
sp_embt_pa_user_activity	sp_embt_pa_user_bottlenecks	sp_embt_pa_user_drill_lcks_db
sp_embt_pa_user_drill_lcks_tp	sp_embt_pa_user_drill_locks	sp_embt_pa_user_drill_seswait
sp_embt_pa_user_drill_syswait	sp_embt_pa_user_drill_trans	sp_embt_pa_user_drill_ts
sp_embt_pa_user_drill_ts_procs	sp_embt_pa_user_index_scans	sp_embt_pa_user_process
sp_embt_pa_user_transactions	sp_embt_pa_user_workload	sp_embt_pa_userdet_io
sp_embt_pa_userdet_locks	sp_embt_pa_userdet_objects	sp_embt_pa_userdet_sql
sp_embt_pa_userdet_waits	sp_embt_pa_version	

Operating System Prerequisites

Performance Analyst requires prerequisites for performance monitoring Windows 2000/XP and UNIX and Linux. The connectivity requirements for the various servers are minimum requirements. The other requirements are optional.

Windows 2000/XP Prerequisites

Connectivity

Performance Analyst requires the following minimum conductivity prerequisites:

- Remote Registry Service should be running on the server if Performance Analyst is connecting to a remote machine. For local machine this does not matter.
- Adequate permissions to connect to the remote computer's registry.

Other Requirements

Performance Analyst requires the following optional prerequisites:

- Diskperf should be running on the server.

  NOTE: You can start diskperf by executing diskperf -y from the command line.

UNIX and Linux Prerequisites

Connectivity

Performance Analyst requires the following minimum conductivity prerequisites:

- Ability to connect to the server and execute commands on the server using one of the following protocols:
  - Telnet
  - SSH
  - Rlogin
- Make sure that the appropriate daemon is running on the server and that the userid used for connecting has the permission to connect using the appropriate protocol.

Commands

The table below describes the commands that the connection Userid needs to have access to, to collect performance stats. The default path for the Userid on UNIX should be setup so you can execute the commands without qualifying them with a path.

Command Type	Description	Commands/Files
Common Commands	Required on all UNIX and Linux servers.	ps, wc, cat, grep, vmstat, uname, uptime, who, df, netstat
Linux Commands	User should be able to access data from these files	/proc/meminfo /proc/cpuinfo /proc/partitions or iostat
AIX Commands	Required for AIX	lsdev pmcycles usr/sbin/prtconf iostat lsps pagesize
Solaris Commands	Required for Solaris	/usr/sbin/psrinfo iostat /usr/sbin/prtconf /usr/sbin/swap pagesize
HP Commands	Required for HP	/etc/swapinfo sar permission to access the following: /dev/mem /dev/kmem /var/adm/syslog/syslog.log

Opening Performance Analyst for the First Time

For Performance Analyst to work, a platform-specific package needs to be installed by the SYS account. When Performance Analyst first starts, it checks for the presence of this package and if it has not been installed, Performance Analyst prompts you to install it. Once created, any account with access to this package can monitor databases using Performance Analyst. You do not need to install any other objects on a monitored database.

> **NOTE:** You must have SYS privileges to install the required packages. After the packages are installed, anyone can use them.

To open the Performance Analyst for the first time, connect to a database through DBArtisan and then select Performance Analyst from the Analyst drop-down menu.

Subsequently, you can connect to Performance Analyst by going to **Start>Programs>EmbarcaderoDBArtisan>Performance Analyst**.

When you open the Performance Analyst this way, you can only review one datasource at a time, as is the case when you open Performance Analyst remotely.

NOTE: If you are connecting through your client machine, it will save you time and give you greater flexibility to connect to Performance Analyst through the DBArtisan console. If you connect through the console, you can open multiple Performance Analyst windows for as many datasources as you like and switch between them by clicking the target tab.

You can also open Performance Analyst in stand-alone mode from your Web client. To do this, see: Opening Performance Analyst in Stand-Alone Mode from Your Web Client

Opening Performance Analyst in Stand-Alone Mode from Your Web Client

The Embarcadero Performance Analyst Login dialog box lets you open Performance Analyst in stand-alone mode from a Web client. This way you can check the performance of any or all datasources that are registered through DBArtisan when you are working remotely.

To connect to the Performance Analyst when you are not using the DBArtisan console, simply go to

<installation directory>:\PAnalyst\PerfAnalyst.html

When you've made the connection, a login dialog box opens.

If you are connecting to the Performance Analyst from your Web client, you can only see one database at a time.

NOTE: Keep in mind that when you connect to Performance Analyst remotely, the Auto Connect option you may have selected when you registered the target datasource in DBArtisan is not functional. You must supply a valid user ID and password to log in to Performance Analyst successfully.

Login As: is only an option when you are logging into an Oracle database. Sybase, SQL Server and DB2 log in as default. The properties for the login roles are as follows:

- **Default** - Establishes a connection with no system roles.
- **SYSDBA** - Lets you perform *all* DBA-related functions. SYSDBA is required for the SYS id when you log into a 9i or later instance.
- **SYSOPER** - Lets you perform many DBA-related functions, like startup and shutdown.

Licensing Performance Analyst

The Performance Analyst is licensed per named server. You are issued one license key for the number of databases/servers you want to license. For example, if you purchase ten licenses for Performance Analyst, you only need to obtain one license key from Embarcadero Technologies.

For more information, see Analyst License Manager.

Performance Analyst License Manager

The Analyst License Manager lets you manage licenses.

The table below describes the options and functionality on the Analyst License Manager:

Option	Description
License Summary	Displays basic licensing metadata for the product.
Available Host(s)	Lets you select a host from the list of hosts available to license.
License	Click to add a license to a host.
Licensed Host(s)	Displays hosts already licensed.
Remove	Click to remove a license from a host.
Apply	Click to make you license choices permanent.
Cancel	Click to close the Licensing Dialog box. **NOTE:** If unsaved license changes have been made, the Analyst asks if you want to discard their changes before exiting.

Completing the Performance Analyst License Manager

To complete the Analyst License Manager, do the following:

1. Open DBArtisan.
2. On the **Help** menu, click **Licensing**.

   DBArtisan opens the License Configuration Dialog box.
3. Click **License**.

   DBArtisan opens the Analyst License Manager.
4. Select options.
5. To make your license choices permanent, click **Apply**.
6. To close the Analyst License Manager, click **Close**.

For more information, see Analyst License Manager.

Performance Analyst Add an Analyst License Dialog Box

The Add an Analyst License Details dialog box lets you add a license.

The table below describes the options and functionality on the Add a Analyst License Details dialog box:

Option	Description
Select a DBMS platform	Lets you select the DBMS platform.

For more information, see Completing the Analyst License Manager.

>

Performance Analyst Application Basics

Performance Analyst displays a complete picture of a database's performance. Performance Analyst is a quick, lightweight tool that lets you investigate performance bottlenecks.

For more information, see Product Design.

Performance Analyst Product Design

The Performance Analyst product design includes Home Pages.

Home Pages for Oracle

Performance Analyst displays home pages of various performance categories with detailed performance data. The Performance Analyst product design includes the following home pages:

- Home
- Memory
- I/O
- Objects
- OS
- RAC Detail View
- Space
- Users

Home Pages for Microsoft SQL Server

Performance Analyst displays home pages of various performance categories with detailed performance data. The Performance Analyst product design includes the following home pages:

- Home
- Database
- Memory
- I/O
- Space
- Users
- OS

Home Pages for IBM DB2 for Linux, Unix, and Windows

Performance Analyst displays home pages of various performance categories with detailed performance data. The Performance Analyst product design includes the following home pages:

- Home
- Cfg (Database Configuration)
- I/O

- Instance
- Memory
- OS
- Space
- Users

Home Pages for Sybase ASE

Performance Analyst displays home pages of various performance categories with detailed performance data. The Performance Analyst product design includes the following home pages:

- Home
- Databases
- I/O
- Memory
- OS
- Space
- Users

Performance Analyst Product Navigation

Navigating Performance Analyst is easily accomplished through two main mechanisms:

1. The main Performance Analyst toolbar buttons open the major performance analysis views (the toolbar buttons you see are self-explanatory) within the product as well as frequently used actions like Refresh and Performance Threshold Customizations.
2. If you right-click any area within Performance Analyst, the application displays a shortcut menu that lets you easily move between any of the main Performance Analyst home pages and drill-down views. The shortcut menu also contains context sensitive options like drilling down further into a particular statistic, copying the display to a text file, and more.

For more information, see Toolbar Icons.

Toolbar Icons

Performance Analyst has several unique icons that take you directly to some functions and statistical groups that can help you see and sort through groupings of data. The functional icons you see on the Performance Analyst toolbar are as follows:

- **Top Sessions** -  - This takes you directly to the Top Sessions tab of the Users Detail view where you see the Top I/O, Memory, CPU, and Sessions statistics.
- **Top SQL** -  - Opens the Top SQL detail view and lets you move among the Summary, Criteria, SQL Details, and Recent History tabs.
- **Response Time (Oracle 10 g only)** -  - Oracle 10g collects response time statistics, which are displayed when you click the icon.

- **Alarms** -  - Goes directly to the Alarms log.
- **Threshold Maintenance** -  - Clicking here opens the Threshold Template Maintenance dialog box. Thresholds are explained in the Using chapter.
- **Refresh** -  - Click this to update the stats/information on whatever page you are currently scrutinizing.
- **Options** -  - To open the Performance Analyst options dialog box, where you can change the values for Oracle, SQL Server, Sybase, or DB2 options, click this cheerful checkerboard.
- **Exit** -  - This closes Performance Analyst for the database you are reviewing. If you have multiple Performance Analysts open, clicking this does not close all Analysts.
- **Maximize Window** -  - This icon appears on each pane of a multi-pane window and lets you expand the pane of choice to fill the screen.
- **Minimize Window** -  - When you have seen enough of an expanded pane, you can bring it back to its original size by clicking this icon.
- **Help** -  - Another in this series of intuitive icons, click this to open the context-sensitive Help.

Alarm Sound

When you register an alarm in Performance Analyst, it creates a periodic alarm sound. To turn off the alarm sound, do the following:

1. Click **Start**, **Programs**, **Accessories**, **Entertainment**, and then **Volume Control**.

   Windows opens the Volume Control dialog box.

2. Under **Wave**, select **Mute**.

   NOTE: Hiding the Performance Analyst icon in the task bar does not mute the sound.

Using Performance Analyst

Performance Analyst allows you to organize performance metrics in the particular way you want to see them. You can drill-down into the many layers of database performance and also be able to locate performance issues in your database in a fast and effective manner. You can set alarm triggers and thresholds for multiple databases across several platforms.

For example, key ratios in Performance Analyst are calculated using *delta* statistics, which is superior to computing ratios with cumulative values. Delta statistics, which are values derived by using the before and after measures in each statistical refresh done in Performance Analyst, lets you see changes in performance metrics that would otherwise be invisible if cumulative statistics were used. This chapter tells you how to set the refresh intervals that best suit your needs.

So, this Using chapter explains *how* to reveal the layers of information you want to see. The platform-specific Expert Guides explain *what* you are seeing as they discuss individual statistics, the associated metrics, and troubleshooting suggestions:

- Oracle Performance Analyst statistics
- Microsoft SQL Server Performance Analyst statistics
- IBM DB2 (W/U/L) Performance Analyst statistics
- Sybase ASE Performance Analyst statistics

Performance Analyst Tools

Performance Analyst offers tools that help you configure what you can see when you survey target databases. It's an easy thing to set up an alarm notification for individual statistic thresholds based on parameters you designate. And, when you work through DBArtisan, you can easily correct problems if you notice something is amiss.

Performance Analyst supports the following functionalities:

Alarm Log	Change Graph Type	Copy	Details	Explain SQL
Manage with DBArtisan	Options	Response Time (ORACLE 10g ONLY)	Save	Thresholds
Top Sessions	Top SQL			

Alarm Log

The Alarm Log displays every performance violation found in your currently monitored databases. Unlike some products that only show you alarms for a particular database, Performance Analyst gives you the option to see all alarms for each database you are monitoring in one central log so you can see every problem in a single place. And, of course, you have the ability to determine alarm levels by setting the Threshold Properties.

To open an Alarm Log:

- Open the Performance Analyst and then scroll down to **Alarms** on the drop-down menu.
- Move your mouse over the Performance Analyst toolbar and click the exclamation point (you'll see **Alarms** appear in a tool tip under the hand icon).

Organizing the Alarm Log

You can arrange the data in the Alarm Log in several ways:

> **NOTE:** Changes you make to one Alarm Log are not enacted globally. This allows you to customize each log based on your needs database by database if you want.

- **Move columns left or right** to place them in an order that makes the most sense to you. Do this by dragging-and-dropping a column heading to re-place it in the target location.
- **Sort column rows up and down or sideways**. Right-click a column heading and select:
  - **Sort Ascending**: For example, you can see the alarms as they move from Critical to Cleared.
  - **Sort Descending**: For example, you can see the alarms as they move from Cleared to Critical.
  - **Sort Multiple Columns**: Lets you group multiple columns in ascending or descending order. You cannot mix and match, however. That is, you cannot sort some columns to present data in ascending order and others to present data in descending order.
  - Sort Clear: Lets you clear the relationships you created if you sorted multiple columns.

The table below describes the columns of the Alarm Log:

Column	Description
Alarm Type	Displays alarm types: Cleared, Critical, Fair and Poor.
Datasource	Datasource where the alarm was identified.
DBMS Type	Datasource DBMS type.
Date/Time	Date/time of when the alarm was observed.
Category	Displays metric category: memory, I/O, space, objects, or users.
Statistic	Metric name.
Value	Value of the statistic that caused the violation.
Explanation	Displays a text explanation of the problem. These text explanations can be very generic in nature and just plug in certain threshold metrics, etc.

Alarm Log Activities

From within the Alarms Window, right-clicking any row gives you more opportunities to get the most out of what you see. The Menu selections are:

- **Applicable Page** - Takes you to the row-specific Category page where you can see the problem metric and what's happening with the surrounding statistics. For example, if you click a critical alarm in the Databases category, you connect immediately to the Databases home page.
- **Options** - Takes you to the Alarm Log Options dialog box where you adjust some general alarm log properties and alarm filters.
- **Clear List** - When you select this, the alarm log is cleared of all entries.
- **Help** - Opens the Help discussion for the target element.
- **Refresh** - Refreshes the display to show you any updates.
- **Copy** - Places the data in the alarm log at that point in time on the Windows clipboard. From there you can paste it into a Word or Excel document.

- **Save As** - Lets you name and save the Alarm Log to the location of your choice as a Microsoft Excel worksheet. You can then manipulate the worksheet as you would any other.
- **Print** - Enables you to print the alarm log.

Alarm Log Options Dialog Box

The Alarm Log Options dialog box is command central for how and where you display and manipulate previously configured alarms. Alarms themselves are configured in the Threshold Properties, but this is where you arrange the particulars of the Alarm Log your databases. Keep in mind that changes you make to this dialog box are reflected globally across all platforms.

The table below describes the columns of the Alarm Log Options dialog box:

Option	Description
General Properties	
Do not log alarms to the alarm log	When you select this option, alarms will not be saved to an alarm *log*. Depending on the filters you select, you will still see alarms in the alarm *window*.
Log alarms to the alarm log	This option lets you save alarms to an alarm log. You can type or browse to choose a directory where you want to save the log and you can also set the maximum alarm log size in KB.
Filters	
Show alarms for only current datasource	When you select this option, you see only alarms for the current datasource in the alarm window. So, if you have multiple datasources open and you go from one to another, the alarm window will display only the alarms for the database you are observing at that very moment.
Show alarms for all currently modified datasources	This option lets you see what alarms are occurring across all databases you are connected to at the moment.
Show only active alarms	If you select this option, you'll see only those alarms that are active at that moment in time. These are alarms that you decided to include as part of a threshold template.
Show only alarms where the Explanation contains the following text	Lets you fine-tune the alarms you see depending on what you are most interested in seeing. For example, if you are particularly interested in seeing Bottleneck alarms, type Bottlenecks. **NOTE**: The text search for this function is case-sensitive.
Show only alarms that are of the following severity	Select to display alarms that are critical, poor, fair, or cleared.

Changing Graph Type

When you are looking at a statistic that is displayed in graph form, you can very easily change how that data is displayed. When you choose to switch from one type to another, the changes are limited to the specific graph and its data that you are manipulating. So, for example, the same page can display a bar graph and a pie chart.

The graph types available to you are:

- Line Graph
- Area Graph
- Pie Graph
- Bar Graph
- Horizontal Bar Graph

To change the graph type:

- Right-click the target graph, and from the drop-down menu, select **Change Graph Type**.

Copy

The Copy functionality lets you copy data in a tabular grid display into the Windows clipboard. The data can then be saved and used for later comparisons as you continue to tweak your databases' performance.

To Copy data, do the following:

1. Hover the pointer over the target section or metric, and then right-click and select **Copy**.

   Performance Analyst copies data into the Windows clipboard.

2. Select an existing, or new, Word or Excel document where you want to **Paste** the data.

3. **Save**, or **Save As**, the document to the target location.

Details

The Details functionality lets you drill-down into the target metric data. Generally you are taken to the metric's category detail page.

To see section or metric details:

- Right-click in the target section or metric and select **Details** from the drop-down menu.

  After you are taken to your destination, right-click again to see if another level of detail is available to you.

Explain SQL

Performance Analyst lets you run Explain SQL Plans from the Top SQL window. If you are seeing problem metrics, you can take steps to correct the SQL statements.

Opening the Explain SQL/Explain Plan:

1. Click the **SQL** button on the Performance Analyst toolbar to open the Top SQL details page.

2. Click the **SQL Details** tab.

   You can also right-click anywhere and select **Navigate to...>Top SQL>SQL Details** from the drop-down menus.

3. Move down the list of SQL text in the grid and the SQL text appears in the Selected SQL Text window.

4. Press the **Explain SQL** button to open the Explain Plan.

Manage with DBArtisan

Manage with DBArtisan opens the DBArtisan editor for the object that has been selected for editing in a grid/tabular display of data.

By moving directly to DBArtisan from Performance Analyst, you have an opportunity to fix database problems as soon as you see them. This option is not available, however, when you open the Performance Analyst from a Web client.

Performance Analyst Options Dialog Box

The Performance Analyst Options dialog box is where you can specify various refresh intervals, threshold polling, and datasource overrides among other settings. You can set or change defaults that are both global and platform-specific.

NOTE: The settings on the Global and Datasource tabs apply to all datasource platforms. The SQL Server, Sybase, and DB2 tab settings are platform-specific. Some options on the platform-specific tabs may look as if they are available to you, but in reality they are not. For example, if you are connected to an Oracle datasource, changing the Monitoring Refresh interval on the Sybase tab is a hollow gesture.

To open the Performance Analyst Options dialog box:

- Open the Performance Analyst and then scroll down to **Options** on the drop-down menu.
- Right-click in the Performance Analyst window for the target datasource and select **Options**.
- Move your mouse over the Performance Analyst toolbar and click the multicolored square (you'll see Options appear in a tool tip under the hand icon).

The table below describes the tabs of the Performance Analyst Options dialog box:

Tab	Field	Option	Description
Global	Engine	SQL metric refresh interval	Lets you set the time, in seconds, between collect cycles for the SQL metrics reflected in Performance Analyst. SQL metrics include all non-operating system metrics.
		OS metric refresh interval	Lets you set the time, in seconds, between collect cycles for all operating system-related metrics reflected in Performance Analyst. OS metrics include all metrics shown on the OS page.
		Threshold polling	Lets you determine which metrics, if any, you want to test against your pre-set thresholds. The possible settings for the Threshold polling box are: All Metrics - Polls all metrics Current Page Only - Polls only those metrics reflected on the currently open page in Performance Analyst Off - Turns polling off. If polling is turned off, Performance Analyst does not activate Alarms.
		Threshold polling interval	Lets you set the time, in seconds, between collect cycles for threshold-related metrics. Once collected, the polled metrics are compared to the pre-set thresholds. If you select **Off** in the Threshold polling box, Performance Analyst does not poll these metrics, regardless of the interval.
		Run queries for all tabs	Lets you determine whether or not you want Performance Analyst to refresh all tabs according to the SQL metric refresh interval. If Run queries for all tabs is not checked, Performance Analyst runs queries only for the current tab.
	Visual	Show alarms icon in system tray	Lets you determine whether or not you want Performance Analyst alarms to appear in the Windows System Tray on the lower right-hand side of your screen. If you select Show alarms icon in system tray, Performance Analyst places an icon in the Windows System Tray each time an alarm is raised.
		Number of Historic Values in Graphs	Lets you specify the number of values in the graphs.

Tab	Field	Option	Description
Datasource	Operating System	OS type	Lets you select the operating system type for the current datasource. The options are: MS Windows - Select MS Windows for all supported types of Microsoft Windows Unix or Linux - Select Unix or Linux if you operating system is any flavor of Unix
		Host name	Lets you specify the name of the host for the current datasource, if necessary. Performance Analyst automatically fills in any detectable information, including the Host name box.
		Username	Lets you select a local or network username that Performance Analyst can use to connect to the datasource. If you select the Use Datasource Username and Password check box, the Username box is unavailable.
		Password	Lets you select a local or network password that Performance Analyst can use to connect to the datasource. If you select the Use Datasource Username and Password check box, the Password box is unavailable.
		Protocol	Lets you specify a protocol for communicating with datasources on Unix-based operating systems. The options are: SSH Telnet Rlogin
		Use Datasource Username and Password	Lets you specify whether or not you want to use the datasource username and password to connect to the datasource as opposed to a network or operating system username and password
		Save Username and Password	Lets you save the username and password that appear in the Username and Password boxes.
	Overrides	SQL metric refresh interval	Lets you override the SQL metric refresh interval in the Global Tab for the current datasource. The SQL metric refresh interval specifies the time, in seconds, between collect cycles for the SQL metrics reflected in Performance Analyst. SQL metrics include all non-operating system metrics.
		OS metric refresh interval	Lets you override the OS metric refresh interval in the Global Tab for the current datasource. The OS metric refresh interval specifies the time, in seconds, between collect cycles for all operating system-related metrics reflected in Performance Analyst. OS metrics include all metrics shown on the OS page.
SQL Server	Global Options	SQL Profiling Enabled	Performance Analyst requires some privileges that go along with SQL Server Profiler in order to return some statistics for SQL Server databases. Check SQL Profiling Enabled to ensure that Performance Analyst can collect all the necessary metrics.
	New Datasource Defaults	Start SQL Profiling	Lets you set new datasources to start SQL Profiling when they are monitored by Performance Analyst.
		SQL Profiling refresh interval	Lets you set the default time for new datasources, in seconds, between collect cycles for the SQL metrics that require the SQL Server Profiler.
	Current Datasource Settings	Start SQL Profiling	Lets you set current datasources to start SQL Profiling when they are monitored by Performance Analyst.
		SQL Profiling refresh interval	Lets you set the default time for current datasources, in seconds, between collect cycles for the SQL metrics that require the SQL Server Profiler.

Tab	Field	Option	Description
Sybase	Global Options	Monitoring Table refresh interval	Lets you specify, in seconds, the amount of time between refreshes for the Monitoring Table. The Monitoring Table is table installed by Performance Analyst used to store various metric-related information. The default is 15 seconds.
	Datasource Overrides	Monitoring Table refresh interval	When you select this option and enter a time, in seconds, this will override
	Sample Query	Use custom sample query	You can use this option to build a custom sampling query and check the database response time by reviewing the Sample Query Response Time graph on the Database home page.
DB2	Monitor Settings	Monitor Switches	You are able to set monitor switches ON for the current session. When check marks are visible, that means they were previously set as defaults. Check or uncheck the boxes as you wish. Monitor settings are available for Buffer Pool. Statement, Lock, Table, Sort, Unit of Work. Metrics for all checked elements are gathered in the current session.
		SQL Statement Cache	You can collect SQL statements on each refresh, collect SQL statement only when you are on the Top SQL pages, or choose not to collect any SQL statements.
		Database System Monitor Heap	The database system monitor heap size shows the current value, in pages, allocated for the database system monitor data. The amount of memory required for monitoring depends on the number of monitoring applications, switch settings, and the level of database activity. If you see frequent heap overflow errors in the output message window when monitoring DB2 through Performance Analyst, adjust the setting here.
		Save settings as default	Gives you the option to set the monitoring settings as a default for the current DB2 database or for the current session only.

Response Time for Oracle 10g

The Response Time functionality lets you analyze response times. You can review historical data and pinpoint who was on a database and what they were doing during a specific period of time.

To see Response Time statistics:

- Right-click in the target and select **Navigate to...>Response Time Analysis** from the drop-down menu.
- Hover your mouse over the Performance Analyst toolbar and click the magnifying glass (you'll see **Response Time** appear in a tool tip under the hand icon).

Save As...

The Save As option gives you an opportunity to save data in a tabular format in the location of your choice. What this means is that the current data will be saved to an existing or a new Microsoft Excel worksheet. You can then manipulate and compare data with all the functionality Microsoft Excel gives you.

To Save As

- Right-click anywhere on the metric or in the Alarm Log, for example, and choose **Save As...** from the drop-down menu.

Thresholds

Thresholds are the point that must be exceeded to begin producing a user-defined effect or result, or to elicit a response, such as an alarm. Performance Analyst provides default threshold templates for each database platform. The default template is applied unless you choose to apply a customized template. There are two methods to place thresholds on statistics in Performance Analyst:

- **Threshold Templates** - Threshold templates are logical containers used to define and deploy thresholds to datasources. You can apply many templates at a single time and you can apply a single threshold template to many datasources. To customize a template, use the Threshold Template Maintenance dialog box. As you add or edit thresholds for a customized template, you make your changes in the Threshold Template Properties dialog box. And finally, when you have customized a threshold template to your satisfaction, you Select Datasource(s) where you want the template to be applied.
- **Datasource Thresholds -** Datasource Thresholds are default or user-defined values that mark performance boundaries. You can apply Datasource Thresholds to specific statistics on a datasource. To make changes to a datasource threshold, use the Datasource Thresholds Maintenance dialog box. From here, the Threshold Templates dialog box gives you the opportunity to use a previously created template to override the target datasource's user thresholds.

Whenever you make changes to a threshold, whether through the Threshold Template functions or from the Datasource Thresholds, you are presented with the Threshold Properties dialog box where you can specify when, and what type of alarm you want to trigger.

Additionally, you can Copy a template.

Threshold Properties Dialog Box

Thresholds are the point that must be met to begin producing a user-defined effect or result, or to elicit a response, such as an alarm. The Threshold Properties dialog box lets you modify properties for any threshold.

To open the Threshold Properties dialog box:

- Choose **Threshold Maintenance...** from the Performance Analyst drop-down menu to open a Threshold Template Maintenance dialog box, and then click **Add** or **Edit**.
- You can also choose **Datasource Thresholds...** from the Performance Analyst drop-down menu to open the Datasource Thresholds Maintenance dialog box, and then click **Add** or **Edit**.

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

Option	Description
Statistic	Lets you choose the target statistic where you want to add a threshold. **NOTE**: If you are adding a statistic to a template or datasource threshold, use the drop-down list to find the statistic you want to add. When you select Edit for any statistic showing in the User Thresholds window, the target statistic is displayed.
Generate Alarm If	Lets you choose or change the target alarm condition. The seven options are: greater than, greater than or equal to, is equal to, is not equal to, is not equal to, less than, and less than or equal to.
Value amount	Lets you set or change the numerical value for the threshold. You cannot leave the value text box blank when you first add the threshold.
Severity	Lets you select the target severity level: Fair, Poor, or Critical.

Threshold Templates Dialog Box

Threshold Templates are logical containers used to define and deploy thresholds to datasources. Performance Analyst offers a default template for each platform. You can copy or customize and apply templates to your heart's content. You cannot, however delete the default template.

To open the Threshold Templates dialog box:

- Choose **Datasource Thresholds...** from the Performance Analyst drop-down menu to open the Datasource Thresholds Maintenance dialog box, and then click **Change**.

The Threshold Templates dialog box displays the available Threshold Templates. The table below describes the options and functionality of the Threshold Templates dialog box:

Option	Description
Name	Displays the name(s) of the available platform-specific Threshold Template(s).
Description	Displays the description(s), if any of available Threshold Template(s).
Override Datasource's User Thresholds	Lets you choose which thresholds to use on the target datasource - the thresholds in the Threshold Template or the user-defined thresholds on the target datasource. If you select the check box, Performance Analyst will override any Datasource Thresholds on the target datasource with the thresholds in the Thresholds Template.

Threshold Template Properties

Threshold Templates are logical containers used to define and deploy thresholds to datasources. The Threshold Template Properties dialog box is where you create a custom template or where you edit an existing one.

To open the Threshold Template Properties dialog box:

- Choose **Threshold Maintenance...** from the Performance Analyst drop-down menu to open a Threshold Template Maintenance dialog box, and then click **Add** or **Edit**.

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

Option	Description
Name	Lets you name the threshold template you are creating. If you are editing a template, target template's name appears here.
Description	Lets you type or review the target template's description.
Platform	Offers a drop-down list of supported platforms. The selection box displays the target template's platform, but you can use this list to see thresholds from other platforms.
Threshold/Severity/Condition/Template	Displays each threshold contained in each template for a given platform. Here you see the severity and condition of each template as well.
Arrow	Lets you select thresholds for the Threshold Template. When you click the arrow, Performance Analyst moves selected thresholds to the Threshold Template you are creating or editing.
Add	Use this function to add statistics to the threshold template you are creating/editing. When you click this button, Performance Analyst opens the Threshold Properties dialog box.
Edit	Lets you edit the selected Datasource Thresholds appearing in the Current Template window. When you click this button, Performance Analyst opens the Threshold Properties dialog box.
Delete	Lets you delete the selected Datasource Thresholds. When you click this button, Performance Analyst deletes the selected thresholds.
Save	Saves any changes you made in this dialog box to create the new or modified template.

Threshold Template Maintenance Dialog Box

Threshold Templates are logical containers used to define and deploy thresholds to datasources. The Threshold Template Maintenance dialog box lets you add or edit templates, as well as configure template maintenance options. For example, you can copy the original Embarcadero Threshold template and modify it to fit your needs. You can control which statistics are sampled for performance violations, as well as the levels of severity to report. When you are satisfied with your Template, you can save it, and then, in one action, apply it to one or more datasources, thereby immediately enforcing your performance standards.

To open the Threshold Template Maintenance dialog box:

- Choose **Threshold Maintenance...** from the Performance Analyst drop-down menu.
- Hover your mouse over the red and green cylinder on the Performance Analyst toolbar and look for the **Threshold Maintenance** tool-tip.

The table below describes the options and functionality on the Threshold Template Maintenance dialog box:

Option	Description
Platform	This is a drop-down list of the platforms. The default is the platform of the datasource you connected to most recently, but you can make another selection.
Name	Displays the name(s) of available Threshold Template(s).
Description	Displays the description(s) of available Threshold Template(s). For customized templates, the description is merely a repetition of the template name unless you wrote something when you created or edited the template.
Add	Lets you create Threshold Templates. When you click this button, Performance Analyst opens the Threshold Template Properties dialog box.
Edit	Lets you edit the selected Threshold Templates. When you click this button, Performance Analyst opens the Threshold Template Properties dialog box.
Delete	Lets you delete the selected Threshold Templates.
Copy	Lets you create copies of the selected Threshold Template. When you click this button, Performance Analyst opens the Copy Template dialog box.
Apply	Lets you select datasources to apply the selected Threshold Templates. When you click this button, Performance Analyst opens the Select Datasource(s) dialog box.
Select/Unselect All	Lets you make global additions to your templates. **NOTE**: You cannot edit or delete Default templates.

Datasource Thresholds Maintenance Dialog Box

Datasource Thresholds are default or user-defined values that mark performance boundaries. You can apply Datasource Thresholds to specific statistics on a datasource. You can also combine Datasource Thresholds to create a template.

Performance Analyst offers the Datasource Thresholds Maintenance dialog box to let you modify Datasource Thresholds. There are two panes in the Datasource Thresholds Maintenance dialog box:

- Applied Threshold Template
- User Thresholds

To open the Datasource Thresholds Maintenance dialog box:

- Choose **Datasource Thresholds...** from the Performance Analyst drop-down menu.

- Right-click in any open Performance Analyst window and choose **Datasource Thresholds...** from the drop-down list.

The table below describes the options and functionality in the Applied Threshold Template box of the Datasource Thresholds Maintenance dialog box:

Option	Description
Name	Displays the name of the target Threshold Template.
Description	Displays any description of the Threshold Template.
Change	Lets you view and edit all platform-specific Threshold Templates. When you click this button, Performance Analyst opens the Threshold Templates dialog box.
Remove	Removes all thresholds in the Applied Threshold Template grid.
Threshold/Severity/Condition/	Displays the statistic on which the threshold is placed, whether it is critical, fair, or poor, and the condition that triggers the alarm.
Block	Prevents that threshold statistic from being triggered in the template. It's almost as if that threshold never existed in the template. You can, however, reinstate the threshold trigger by highlighting and deleting it from the User Thresholds pane.
Override	Copies the threshold settings from the selected Threshold Template and applies the change to the user-defined Datasource Thresholds. Click this button if the template has the correct Severity and Threshold, but you want to modify the condition for triggering the alarm. When you create a "User Threshold," you are overriding the threshold specified by applied template or you are including an additional threshold not included in the applied template. This may be useful if you have a group of individuals all monitoring the same primary statistics (and thus using the same template) but want to be able to give individual users the ability to modify/add/block a threshold used by the template, without affecting the original template.

The table below describes the options and functionality in the User Threshold box of the Datasource Thresholds Maintenance dialog box:

Option	Description
Threshold	Displays the statistic on which the threshold is placed.
Severity	Displays the severity for the threshold.
Condition	Displays the condition for the threshold.
Add	Lets you create Datasource Thresholds. When you click this button, Performance Analyst opens the Threshold Properties dialog box.
Edit	Lets you edit the selected Datasource Thresholds. When you click this button, Performance Analyst opens the Threshold Properties dialog box. If you Edit a Blocked threshold, it loses its blocked status.
Delete	Lets you delete the selected Datasource Thresholds. When you apply this to a blocked threshold, it is removed from the user thresholds.

Select Datasource(s)

When you have created or edited a threshold template, and you are ready to distribute it, you have the opportunity to apply it to one or more target datasources using the Select Datasources dialog box.

To open the Select Datasource(s) dialog box:

- Choose **Threshold Maintenance...** from the Performance Analyst drop-down menu, and then click **Apply**.

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

Option	Description
Datasources	You can apply the target threshold template to one or more datasources. You can only apply the target threshold template to the platform where it was created.
Override Datasource's User Thresholds	Lets you choose which thresholds to use on the target datasource - the thresholds in the Threshold Template or the user-defined thresholds on the target datasource. If you select the check box, Performance Analyst will override any Datasource Thresholds on the target datasource with the thresholds in the Thresholds Template.

Copy Template

The Copy Template dialog box lets you copy an existing Threshold Template, rename it, and save it for later use.

To open the Copy Template dialog box:

1. Choose **Threshold Maintenance...** from the Performance Analyst drop-down menu, to open the Threshold Template Maintenance dialog box.
2. Select the template you want to replicate, and then click **Copy**.

The table below describes the options and functionality of the Copy Template dialog box:

Option	Functionality
Name	Lets you type the name for the new Threshold Template. The default name is Copy of Frankenstein, or whatever it is you named your template.
Description	Lets you type the description of the new Threshold Template. The default description is simply a repetition of the template name.

Top Sessions

The Top Sessions functionality lets you drill-down into the target metric data and quickly pinpoint the heaviest users of I/O and memory. Generally you are taken to the database-specific Expert Guide when you need an explanation of what you are seeing.

To open Top Sessions:

- Hover your mouse over the little head on the Performance Analyst toolbar and look for the **Top Sessions** tool-tip.

To open the Session Details View

- Double-click any entry in the Top Sessions grid of the Top Sessions view, and the Session Details view opens.

Top SQL

The Top SQL view displays all SQL running or executed in the current datasource. You can use the lists and boxes to reorder the display according to your specific needs.

To view Top SQL, do the following:

- On the Performance Analyst toolbar, click the **SQL** button.

>

Performance Analyst for Oracle Tutorial

Embarcadero Performance Analyst provides key insight into the complex world of database performance. Running inside the Embarcadero DBArtisan console, or stand-alone in a Web browser, Performance Analyst helps you pinpoint performance problems that threaten a database's availability or response times. The powerful analysis facilities will satisfy advanced experts, yet novices will find Performance Analyst's interface easy to use and the expert help a welcome aid to troubleshooting difficult performance issues.

About This Performance Analyst for Oracle Tutorial

This tutorial is intended to help you get started with Embarcadero Performance Analyst.

After completion of this tutorial, you will have the foundation you need to explore the many features and benefits of Performance Analyst. You will have learned how to accurately monitor your database using the most popular performance methodologies in use today. You will know how to drill down into the many layers of database performance and also be able to locate performance issues in your database in a fast and effective manner. Finally, you'll understand how to fix performance problems using the integration that exists between Performance Analyst and Embarcadero DBArtisan.

This tutorial is divided into 8 sections.

You can use this basic tutorial as a roadmap of product highlights.

After starting this tutorial, you can select **Help** from the menu to find many additional resources that complement and build on many of the activities shown in this brief tutorial.

Proceed to Session 1: Getting Started with Performance Analyst for Oracle

Session 1: Getting Started with Performance Analyst for Oracle

Download and Install

You can obtain the latest version of Performance Analyst by downloading the DBArtisan software from the Embarcadero Web site at http://www.embarcadero.com/downloads/downloaddbartisan.asp. Performance Analyst runs fully contained within DBArtisan so you will need the DBArtisan product to use Performance Analyst (although Performance Analyst can also run in standalone fashion within a Microsoft Internet Explorer Web browser).

Click **Download**, and follow the steps indicated. Save the file on your computer and then double-click to launch the self-extracting file that will guide you through the installation process.

When you first install an evaluation copy of DBArtisan, you can use the tool for 14 days. After that time, a permanent license is needed.

Proceed to Overview for Performance Analyst for Oracle

Overview for Performance Analyst for Oracle

The graphic below illustrates all the main elements of the DBArtisan User Interface.

When DBArtisan is connected to an Oracle database, you can open Performance Analyst two ways:

1. On the toolbar, click the **Performance Analyst** icon.
2. On **Utilities** menu, select **Performance Analyst**.

   DBArtisan opens Performance Analyst.

Proceed to Session 2: Using Ratio-Based Performance Analysis in Performance Analyst for Oracle

Session 2: Using Ratio-Based Performance Analysis in Performance Analyst for Oracle

Ratio-based analysis has been around for many years, and used to be the only technique database administrators used when they were called on to diagnose the cause of a database slowdown. Performance ratios are very good roll-up mechanisms for busy DBAs, making possible the analysis-at-a-glance approach, however ratios must be computed correctly for them to be meaningful. Key ratios in Performance Analyst are calculated using *delta* statistics, which is superior to computing ratios with cumulative values. Delta statistics, which are values derived by using the before and after measures in each statistical refresh done in Performance Analyst, lets you see changes in performance metrics that would otherwise be invisible if cumulative statistics were used. Performance Analyst makes it easy to see the ratios that are computed with delta statistics vs. cumulative. A triangle symbol denotes a delta-based statistic, while square symbols represent a cumulative-based measure.

Ratio-based analysis is available for all performance categories (memory, etc.) where ratios can assist in determining the performance level of a database. For example, if you would like to see memory-specific ratios, on the Performance Analyst toolbar click **Memory**. Performance Analyst displays the Memory home page that showcases all important memory-related ratios.

Proceed to Obtaining Detail Information for a Ratio Using Performance Analyst for Oracle

Obtaining Detail Information for a Ratio Using Performance Analyst for Oracle

A ratio gives you a global picture for a particular performance measure. However, you often need more information regarding a statistical measure so you can make an informed decision. Performance Analyst has easy-to-use drill down features that make obtaining detailed information for various ratios a breeze.

For example, to view more information for the Buffer Cache Hit Ratio statistic, do the following:

1. On either the **Main** Home page or the **Memory** Home page, locate the **Buffer Cache Hit Ratio** statistic.
2. Position the pointer over the statistic.

   The pointer changes to a hand symbol.
3. Double-click the left mouse button.

   Performance Analyst opens the Buffer Cache Tab of the memory drill down view.
4. Or right-click the statistic, select **Details** or **Navigate to… > Memory > Buffer Cache**.

   Performance Analyst opens the Buffer Cache Tab of the memory drill down view.

Proceed to Session 3: Using Bottleneck/Wait-Based Performance Analysis in Performance Analyst for Oracle

Session 3: Using Bottleneck/Wait-Based Performance Analysis in Performance Analyst for Oracle

When an Oracle database is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting may mean nothing in the overall scheme of things, or it can be an indicator that a database bottleneck exists. Bottleneck analysis is a valid method of measuring performance because it helps you track where a database and user sessions have been spending their time. If latch contention or heavy table-scan activity has been dragging a database's performance down, you can use bottleneck analysis to confirm the actual root cause.

> **NOTE:** For Performance Analyst to show wait times for Oracle wait events, the database configuration parameter TIMED_STATISTICS must be set to TRUE.

Performance Analyst uses two broad areas of metrics to form its bottleneck analysis. The first is the Oracle wait event interface, which displays where the database and user sessions have been bottlenecked. But, when using bottleneck analysis you can not rely *only* on the information contained in the wait event views that Oracle provides. For example, an object may attempt to extend into another extent of space in a tablespace and yet, be denied if no such free space exists. Such a failure will not be reflected in any wait event, but still represents a very real bottleneck to the database. For this reason, Performance Analyst also shows non-wait event based bottlenecks that represent a very real threat to a database's availability and performance.

Performance Analyst prominently displays bottleneck analysis on every performance home page. The Main home page, displays a summary of all key bottlenecks threatening the database. You can also see specific bottlenecks for memory, I/O, space, objects, and users by navigating to each home page and viewing the Bottleneck Analysis section for that home page. For example, to see bottlenecks that specifically relate to I/O, do the following

1. On the toolbar, click the **I/O** button.
2. Locate the Bottleneck Analysis section at the lower-left portion of the screen

   The I/O bottleneck analysis section displays the top database and log-related Oracle wait events (ranking them by wait *time* and not actual wait count), along with the most active database files.

Proceed to Obtaining Detail on System Bottlenecks Using Performance Analyst for Oracle

Obtaining Detail on System Bottlenecks Using Performance Analyst for Oracle

You can easily obtain more information regarding global bottlenecks that are plaguing a database by drilling down into a particular bottleneck metric. For example, if you would like to see more data regarding the current top system bottlenecks, do the following:

1. Open the **Main** home page.
2. Double-click the **Top System Bottlenecks** area.

   Performance Analyst opens the System Waits view.
3. Or right-click anywhere in Performance Analyst, point to **Users**, and then click **System Waits**.

   Performance Analyst opens the System Waits view.

The System Waits view displays critical details regarding all the wait events that Oracle has recorded. You can easily see what events that occurred the most often as well as the events that have accumulated the most wait time.

Proceed to Viewing Detail for Session Bottlenecks in Performance Analyst for Oracle

Viewing Detail for Session Bottlenecks in Performance Analyst for Oracle

Not only can you obtain detail information for system bottlenecks, but you can also use Performance Analyst to get drill down information regarding session-level bottlenecks. There are several layers to session-level details you can view.

To see more information regarding session-level waits, do the following:

1. Open the **Main** home page.
2. Double-click the **Top Session Bottlenecks** area.

   Performance Analyst opens the Session Waits drill down view.
3. Or right-click anywhere in Performance Analyst, point to **Users**, and then click **Session Waits**.

   Performance Analyst opens the Session Waits drill down view.

When you are in the Session Waits drill down view, to drill further into a session's details, do the following:

1. In the Historical Waits or Current Waits grid, select a session that appears to be experiencing critical waits.
2. Double-click the session or right-click, and then select **Details**.

   Performance Analyst opens a drill down view that shows key statistics for the selected session, including its current work activities, wait events, and SQL statements.

Proceed to Session 4: Using Workload Performance Analysis in Performance Analyst for Oracle

Session 4: Using Workload Performance Analysis in Performance Analyst for Oracle

Key ratios give you a global perspective on database activity. Bottleneck analysis gives you insight into things that are holding up user activity and throughput. But another technique is necessary if a database professional is to really get a handle on what's occurring inside a badly performing database.

Workload analysis involves the investigation of two critical areas of a database's performance:

- Session resource consumption and activity
- SQL execution

Without looking at these two key performance categories, you can miss a lot of what could be responsible for perceived performance problems. When performance on a database takes a sudden nosedive, it is not uncommon to find one or two sessions that are causing the bulk of the workload. Finding these processes can be easily accomplished by viewing session metadata coupled with resource consumption and statistical execution statistics.

Proceed to Pinpointing Top Resource Consumers in Performance Analyst for Oracle

Pinpointing Top Resource Consumers in Performance Analyst for Oracle

Performance Analyst identifies top resource consumers in a couple of different ways. On the main home page, Performance Analyst highlights the top resource sessions across physical and logical I/O, memory usage, and CPU consumption.

The percentage used across all statistical categories are displayed so you can immediately pinpoint a session that is using all or most of a particular resource. The top resource sessions are also displayed on the memory, I/O, and other home pages as well with a listing for the top sessions for that category. For example, the top memory users appear on the memory home page.

Proceed to Getting Workload Details Using Performance Analyst for Oracle

Getting Workload Details Using Performance Analyst for Oracle

Performance Analyst offers several details on top resource sessions in addition to the summaries provided on each home page. A detailed and flexible top sessions view lets you view a summary of all the major resource sessions on the database as well as every critical detail regarding each session connected to the database.

Further drill downs are available that highlights specific details regarding each session's activities and submitted SQL. To see detailed information regarding a session identified in one of Performance Analyst's Workload Analysis sections, do the following:

1. In a **Workload Analysis** display, position the pointer over one of the processes.
2. Double-click the process or right-click, and then select **Details**.

If you want like to see all key workload analysis metadata for sessions, do the following:

1. On the toolbar, click the **Top Sessions** button.
2. Or right-click anywhere in Performance Analyst, point to **Users**, and then click **Session Waits**.

Proceed to Session 5: Using SQL Performance Analysis in Performance Analyst for Oracle

Session 5: Using SQL Performance Analysis in Performance Analyst for Oracle

SQL-based performance analysis is really a subset of Workload Analysis. Understanding current and historical SQL execution patterns lets you to have the second set of data points necessary to properly perform workload analysis. You may find that optimizing SQL code produces some of the best performance-enhancing boosts available for a database.

Proceed to Uncovering Resource-Intensive SQL in Performance Analyst for Oracle

Uncovering Resource-Intensive SQL in Performance Analyst for Oracle

Performance Analyst offers plenty of insight into SQL performance and execution patterns. Each one of the home pages depicts the most costly SQL statements as determined by a variety of different metrics for that performance category (most physical, logical I/O, etc.) For example, Performance Analyst might show that a single SQL statement is responsible for almost 60 percent of all the physical I/O on the database.

Proceed to Getting Details on Top SQL with Performance Analyst for Oracle

Getting Details on Top SQL with Performance Analyst for Oracle

Getting details on identified SQL statements is easy in Performance Analyst. For example, if you would like to see further information on a SQL statement identified on the Performance Analyst Home page, do the following:

1. In the **SQL Analysis** section, position the pointer over **Hash Values**.
2. Double-click or right-click, and then select **Details**.

   Performance Analyst opens the Top SQL summary page that shows a breakdown of all top SQL statements along with their execution statistics

If you would like to see the full SQL text of a SQL statement, do the following

1 In the **Top SQL Details** section, double-click a statement.

Performance Analyst opens the SQL Details Tab of the Top SQL view to display the entire SQL statement and the execution metrics. While on this pane, you can also perform an EXPLAIN of the SQL statement and examine its access path information.

NOTE: To automatically tune a problem SQL statement if you have Embarcadero SQL Tuner installed on your workstation, do the following:

2 On the **SQL Details Tab**, select a problem statement, and then click the **Tune SQL** button.

SQL Tuner opens and you can begin a tuning session.

Proceed to Session 6: Using Operating System Performance Analysis in Performance Analyst for Sybase

Session 6: Using Operating System Performance Analysis in Performance Analyst for Sybase

In many scenarios, an optimally tuned database may not perform well because there are constraints imposed by the system where the database is running. These constraints may include processes competing with the database sever for resources (CPU, I/O or Memory), a slow CPU, insufficient I/O devices, and insufficient memory. Performance Analyst allows you to examine the relevant operating system metrics for Windows Platforms (NT /2000/XP/2003) as they relate to Oracle.

To view the operating system metrics of the machine hosting the Sybase server you are monitoring, do the following:

1 Click the **OS** Performance Analyst toolbar button.

OR

2 Navigate to the OS Home Page by going to **OS**, and then **Home**.

NOTE: The first time you invoke the OS feature you will be prompted for an operating system login. Please specify the correct operating system, host name, username and password, and protocol.

The Operating System home page has been designed in a similar fashion as the datasource home page. On this page you can see the pertinent operating system metrics displayed in five quadrants broken down into the following categories:

- Resource Usage
- Memory Analysis
- Disk Analysis
- Workload Analysis
- Bottleneck Analysis

Proceed to Getting Details on Top SQL with Performance Analyst for Oracle

Obtaining Details on the Operating System Using Performance Analyst for Sybase

You can easily obtain more information regarding operating system problems that are degrading database performance by drilling down into a particular analysis quadrant. For example if you'd like to see more information regarding the processes running on the host machine, do the following:

1. Navigate to the Performance Analyst **Operating System** home page.
2. Double-click the **Top Memory Process** in the **Workload Analysis** quadrant.

   Performance Analyst opens the Processes detail view.

   OR
3. Right-click anywhere in Performance Analyst and select the **Navigate to**, **OS**, and then **Processes**.

The Processes detail tab shows all the processes on the host machine. From this view you can easily determine which processes are showing high marks in I/O or CPU, for example, by simply double clicking on the column headers and sorting the processes in ascending or descending order.

Proceed to Session 7: Identifying and Getting Help with Performance Problems Using Performance Analyst for Oracle

Session 7: Identifying and Getting Help with Performance Problems Using Performance Analyst for Oracle

You do not need to be an advanced expert in performance tuning to spot trouble in a database if you use Performance Analyst. There are a number of mechanisms in the tool that help you quickly locate areas of concern within a database.

Proceed to Using the Alarm Log

Using the Alarm Log

Performance Analyst has a central and flexible alarm log that lists every performance violation found in all monitored databases. Unlike some products that only show you alarms for a particular database, Performance Analyst lists all alarms for each database you are monitoring in one central log so you can see every problem in one place. To view the Performance Analyst Alarm log, do the following:

1. On the toolbar, click the **Alarms** button.

Once in the Alarm log, you can perform a number of actions including:

- Selecting an alarm and viewing the application page that shows details regarding the metric in violation.
- Filtering the types of alarms shown through the Options dialog.
- Clearing all current alarms.
- Copying or saving the Alarm log into a different format.

To select any of these options, do the following:

1. Position the pointer over any area of the **Alarm Log**.
2. In the shortcut menu, select an option.

Proceed to Recognizing Performance Violations

Recognizing Performance Violations

Performance Analyst makes it easy to visually pinpoint performance violations in a database. Any statistic that does not conform to your predefined performance standards is colorized in any view that it appears so you can easily see potential troubleshoots.

The colors used to highlight a particular performance violation depend on the severity of the problem. Performance Analyst uses several layers of performance standards/threshold violations with each layer being represented by a different color. You can control everything regarding threshold violations, including how thresholds are determined, their color representation in Performance Analyst's display, and more by using the Threshold Template feature of Performance Analyst.

How can Performance Analyst inform you of performance violations in your database if you minimize the product to your desktop and are working on other things? Even if Performance Analyst isn't right in front of you, it can still let you know that a problem has been found in your database through the sister alarm icon. When Performance Analyst detects a problem in one of your databases, one of the things the product does is enable an alarm icon in your workstation's systray area that communicates the presence of alarms in your database. So even if the Performance Analyst interface isn't in sight, the product can let you know that something in your database requires attention.

Proceed to Creating Threshold Templates

Creating Threshold Templates

How does Performance Analyst know what to signal as a problem in a database? It uses predefined performance standards that you define and enforce. All of this is accomplished by using Performance Analyst's Threshold templates. A template is a set of performance metrics and their thresholds that can be reused to enforce performance standards across any number of databases. Embarcadero provides a default performance template that contains nearly all of the most-watched performance statistics so you can begin monitoring with confidence right out of the box. However, you can create your own Threshold templates so you only are alerted to performance violations that are important to you and your particular environment.

To view or create Threshold templates, do the following:

1 On the toolbar, click the **Thresholds** button.

  Performance Analyst opens the Threshold maintenance dialog box.

The Threshold Template Properties dialog box lets you select the options to add or edit templates, as well as other maintenance options. For example, you can copy the original Embarcadero Threshold template and modify it to fit your own needs. You can control what statistics are sampled for performance violations as well as the levels of severity you would like reported. Once you are satisfied with your template, you can save it and then apply it to one or more datasources in one action to immediately begin enforcing your performance standards.

Proceed to Getting Help with Performance Problems

Getting Help with Performance Problems

As databases become more complex, so do the problems that plague them. Knowing what to do for a particular performance problem can be hard in some cases, but Performance Analyst helps by providing a complete set of expert help that you can use as a guide when performance problems occur. To obtain expert help for any statistic shown in Performance Analyst, do the following:

1 Position the pointer over a statistic that you need help.

2 Right-click and select **Help**.

The Expert Help gives you a general description of the statistic, the range of values that are considered optimal, and some suggestions on how to fix a problem situation. Expert help can also be invoked for any of the Ratio, Bottleneck, SQL, or Workload Analysis sections shown on every home page by clicking on the help icon for that section, which is located at the top right corner of the section.

Proceed to Session 8: Fixing Performance Problems Using Performance Analyst for Oracle

Session 8: Fixing Performance Problems Using Performance Analyst for Oracle

Finding performance problems is one thing, but fixing them is another. Many performance monitors point out a particular problem in a database, but then offer no way to actually correct the situation. Such is not the case with Performance Analyst. Because of its tight integration with DBArtisan, certain problems found in Performance Analyst can be visually corrected by DBArtisan.

If an Oracle tablespace is about to run out of free space, Performance Analyst can issue an alarm that a space shortage exists. Once the space problem has been identified, do the following:

1 On the **Space** home page toolbar, click the **Space** button.

2 In the **Space Analyst** grid, right-click a tablespace, and then select **Manage with DBArtisan**.

DBArtisan opens an editor where you can perform corrective actions, such as enabling Autoextend for the tablespace, resizing the tablespace's existing datafile(s), or adding a new datafile for the tablespace.

>

Oracle Performance Analyst statistics

This section includes expert help for all Oracle categories and statistics in Performance Analyst views. For detailed information on using the application, see Using Performance Analyst. This guide includes the following sections:

Home	I/O	Memory
Objects	OS	RAC Detail View
Response Time Analysis View (Oracle 10G Only)	Session Details	Space
Users	Top SQL	

Home

The Home view lets you review availability and overall performance of all monitored databases from a single window. The Home page includes the following sections:

Bottleneck Analysis pane	Key Ratio Analysis pane	SQL Analysis pane
Storage Analysis pane	Workload Analysis pane	

Key Ratio Analysis pane

Database performance analysts typically use one of two methods for examining the performance levels of a database: Ratio-based or wait/bottleneck-based. Ratio-based analysis involves examining a number of key database ratios that can be used to indicate how well a database is running. Performance ratios serve as roll-up mechanisms for at-a-glance performance analysis. Succinctly presented performance ratios can be quickly scanned to see if any database needs immediate attention.

While there are certainly many opinions as to which rules to follow, there are some standards to adhere to. To start with, many of the formulas that make up ratio-based analysis must be derived from delta measurements instead of cumulative statistics. Many of the global ratios a you examine come from the v$sysstat performance view. The performance view maintains a count of all the occurrences (in the VALUE column) of a particular database incident (in the NAME column) since the database was brought up. For databases that are kept up for long periods of time, these values can grow quite large and impacts how a particular ratio is interpreted. However, if delta statistics are used (taking, for a specified sampling period, the before and after counts of each statistic that make up a ratio), then an accurate and current portrayal of the various ratios can be had.

While there are a number of rules of thumb that can be used as starting points to begin the evaluation of database performance, DBAs must determine each database's individual 'personality' with respect to the key performance ratios. Some hard and fast rules simply do not apply to every database. The danger in using blanket ratio standards is that they can lead to haphazardly take action, which can at times contribute nothing to the situation, and sometimes even degrade performance.

The following ratios are used on the Performance Analyst Home page to succinctly communicate the general overall performance levels of the monitored database:

Buffer Cache Hit Ratio	Library Cache Hit Ratio	Dictionary Cache Hit Ratio
Memory Sort Ratio	Parse/Execute Ratio	Buffer Busy Waits
Rollback Contention Ratio	Latch Miss Ratio	Parallel Query Busy Ratio
Free Shared Pool Percent		

Buffer Cache Hit Ratio

Data read from memory will produce end-user response times many times faster than when that same data is read from disk. Keeping physical I/O's to an absolute minimum is one of the purposes of the Oracle buffer/data cache.

The buffer cache hit ratio is a terrific indicator of how often user requests for data are satisfied through memory vs. being physically read from disk. A number of key counters in Oracle are used to arrive at this statistic. The table below lists them:

Counter	Description
DB BLOCK GETS	Data read from memory for DML operations.
CONSISTENT GETS	Data read from rollback segments in memory.
PHYSICAL READS	Data read physically from disk.
DIRECT READS	Data read physically from disk that bypasses the buffer cache. Direct reads are filtered out of overall physical reads so an accurate cache hit ratio can be determined.

Dividing the data read from memory by data read from disk yields the cache hit ratio.

Location	Oracle Performance Analyst statistics > Home > Key Ratio Analysis pane Oracle Performance Analyst statistics > Memory > Key Ratio Analysis pane
Metrics	To help ensure excellent performance, you want to keep your cache hit ratio in the neighborhood of 90% or higher. However, you should be aware that every database has its own 'personality' and may exhibit excellent performance with below average readings for the cache hit ratio. You should also be aware that excessive logical I/O activity can produce a very high cache hit ratio while actually degrading overall database performance. Consistently viewed low readings (60% or less) may require tuning attention on the part of the DBA. **Note:** For Oracle8i or earlier, the adjustment of the db_block_buffers tuning parameter is required. For Oracle9i and later, the db_cache_size parameter is the parameter that needs attention. For any increases in db_block_buffers to take effect on Oracle8i or earlier, the database must be cycled. The db_cache_size parameter in Oracle9i and later, however, is dynamic and may be altered without stopping and starting the database instance.
Correction	If a problem is found in Oracle8i or earlier, then: *Edit the Init.ora file for the database.* *Increase the amount of db_block_buffers to a higher value (take caution to not over-allocate; ensure enough free memory exists on server before increasing value).* *Cycle the Oracle server when possible to allow the new value to take effect.* *Monitor the new value to see if performance improves.* If a problem is found in Oracle9i or later, then: *Increase the size of the db_cache_size parameter through use of the ALTER SYSTEM SET db_cache_size command value (take caution to not over-allocate; ensure enough free memory exists on server before increasing value).* *Monitor the new value to see if performance improves.* If using an SPFILE, save the new configuration values so Oracle will reuse them each time the database is stopped and re-started.

Library Cache Hit Ratio

Oracle's shared pool is made up of two main memory areas - the Library Cache and the Data Dictionary Cache, and offers a number of tuning possibilities.

The library cache holds commonly used SQL statements - basically database code objects. A method for improving performance in Oracle is to encourage the reuse of SQL statements so expensive parse operations may be avoided. The library cache assists this tuning effort.

The library cache hit ratio offers a key indicator in determining the performance of the shared pool. It shows how often SQL code is being reused by other database users vs. the number of times a SQL statement is broken down, parsed, and then loaded (or reloaded) into the shared pool.

Location	Oracle Performance Analyst statistics > Home > Key Ratio Analysis pane Oracle Performance Analyst statistics > Memory > Key Ratio Analysis pane
Metrics	A high library cache hit ratio is a desirable thing. You should strive for a hit ratio between 95-100%, with 99% being a good performance benchmark for code reuse. Note that when a database is first started, the library cache hit ratio will not be at an optimal level because all code being used will be relatively new, and as such, must be parsed and placed into the shared pool. If, however, after a solid hour or two of steady database time, the library cache hit ratio has not increased to desirable levels, you should look into the possibility of increasing the shared_pool_size parameter. Other red flags that can indicate a too small shared pool include: *A wait count for the event 'latch free' of ten or greater* *The library cache wait count of two or greater.* These indicators can be tracked with Performance Analyst's Bottleneck and Wait detail views. A way of improving the library cache hit ratio is by encouraging code reuse through the implementation of bind variables. Discouraging hard coded literals in application code and instead making use of variables bound at run time aids in the reuse of SQL code that is maintained in Oracle's shared pool. Note that bind variables can have an affect on the cost-based optimizer though. A second way is to pin frequently used code objects in memory so they will always be there when needed. This can be accomplished through the use of the system supplied DBMS_SHARED_POOL package. You can use Performance Analyst to view objects in the shared pool that are always present and/or have increasing reload numbers to help identify objects that are good candidates for pinning.
Correction	If a problem is found in Oracle8i or earlier, then: *Edit the Init.ora file for the database.* *Increase the amount of shared_pool_size to a higher value (take caution to not over-allocate; ensure enough free memory exists on server before increasing value).* *Cycle the Oracle server when possible to allow the new value to take effect.* *Monitor the new value to see if performance improves.* If a problem is found in Oracle9i or later, then: *Increase the size of the shared_pool_size parameter through use of the ALTER SYSTEM SET shared_pool_size command value (take caution to not over-allocate; ensure enough free memory exists on server before increasing value).* *Monitor the new value to see if performance improves.* *If using an SPFILE, save the new configuration values so Oracle will reuse them each time the database is stopped and re-started.* If you determine that SQL literals are causing SQL to not be reused, you can (in Oracle 8.1.6 and later): *Change the cursor_sharing parameter to FORCE by using the ALTER SYSTEM SET cursor_sharing=FORCE command.* *Monitor database to see if parse activity is reduced and library cache reloads shrink.* If using an SPFILE, save the new configuration values so Oracle will reuse them each time the database is stopped and re-started. If using an Init.ora file, add the cursor_sharing=FORCE parameter to the file.

Dictionary Cache Hit Ratio

Oracle's shared pool offers an number of tuning possibilities and is made up of two main memory areas:

1 Library Cache

2 Data Dictionary Cache

The dictionary cache hit ratio offers another key indicator in determining the performance of the shared pool. It shows how often object definitions are found in memory vs. having to read them in from disk. Because Oracle references the data dictionary many times when an SQL statement is processed, it is imperative that as much of this vital reference information be kept in RAM as possible.

Location	Oracle Performance Analyst statistics > Home > Key Ratio Analysis pane Oracle Performance Analyst statistics > Memory > Key Ratio Analysis pane
Metrics	Just as with the library cache, a high data dictionary cache hit ratio is desirable. Strive for a hit ratio between 90-100%, with 95% being a good performance benchmark. **NOTE:** When a database is first started, the data dictionary cache hit ratio is not at an optimal level because all references to object definitions are relatively new, and as such, must be placed into the shared pool. Look for hit ratios in the eighty's for new database startups. If, however, after a solid hour or two of steady database time, the data dictionary cache hit ratio has not increased to desirable levels, increase the shared_pool_size parameter. Databases supporting applications that involve large number of objects (such as an Oracle Financials installation) should have larger than normal shared pools to support the required object definitions. Although each parameter is not individually tunable (it was in Oracle6), you can see which area of the dictionary cache could be pulling the overall hit ratio down.
Correction	If a problem is found in Oracle8i or earlier, do the following: *Edit the Init.ora file for the database.* *Increase the amount of shared_pool_size to a higher value. Take caution to not over-allocate; ensure enough free memory exists on server before increasing value.* *Cycle the Oracle server when possible to allow the new value to take effect.* *Monitor new value to see if performance improves.* If a problem is found in Oracle9i or later, do the following: *Increase the size of the shared_pool_size parameter through use of the ALTER SYSTEM SET shared_pool_size command value. Take caution to not over-allocate; ensure enough free memory exists on server before increasing value.* *Monitor the new value to see if performance improves.* *If using an SPFILE, save the new configuration values so Oracle reuses them each time the database is stopped and re-started.*

Memory Sort Ratio

Oracle's SGA is not the only memory structure used by Oracle for database work. One of the other memory areas used by Oracle8i and earlier for normal activity is an area set aside for sort actions. When a sort operation occurs, Oracle attempts to perform the sort in a memory space that exists at the operating system level. If the sort is too large to be contained within this space, it continues the sort on disk - specifically, in the user's assigned TEMPORARY TABLESPACE. Oracle records the overall number of sorts that are satisfied in memory as well as those that end up being finalized on disk. Using these numbers, you can calculate the percentage of memory sorts vs. disk sorts and get a feel for how fast your sort activity is being resolved. Obviously, memory sorts completes many times faster than sorts forced to use physical I/O to accomplish the task at hand.

Oracle9i or later has the option of running automatic PGA memory management. Oracle has introduced an Oracle parameter called pga_aggregate_target. When this parameter is set and you are using dedicated Oracle connections, Oracle ignores all of the PGA parameters in the Oracle file, including sort_area_size, hash_area_size and sort_area_retained_size. Oracle recommends that the value of pga_aggregate_target be set to the amount of remaining memory (less a 10% overhead for other server tasks) on a server after the instance has been started.

Location	Oracle Performance Analyst statistics > Home > Key Ratio Analysis pane Oracle Performance Analyst statistics > Memory > Key Ratio Analysis pane Oracle Performance Analyst statistics > Users > Home > Key User Analysis pane
Metrics	If the memory sort ratio falls below 90%, and you are on Oracle8i or earlier, increase the parameters devoted to memory sorts - sort_area_size and sort_area_retained_size. For Oracle9i or later, investigate the use of pga_aggregate_target. Once the pga_aggregate_target has been set, Oracle automatically manages PGA memory allocation, based on the individual needs of each Oracle connection. Oracle9i or later allows the pga_aggregate_target parameter to be modified at the instance level with the alter system command, thereby lets you dynamically adjust the total RAM region available to Oracle9i. Oracle9i also introduced a new parameter called workarea_size_policy. When this parameter is set to automatic, all Oracle connections benefits from the shared PGA memory. When workarea_size_policy is set to manual, connections allocates memory according to the values for the sort_area_size parameter. Under the automatic mode, Oracle tries to maximize the number of work areas that are using optimal memory and uses one-pass memory for the others.
Correction	If you find a problem, do the following: *Edit the Init.ora or SPFILE file for the database.* *Increase the amount of sort_area_size to a higher value. Take caution to not over-allocate; ensure enough free memory exists on server before increasing value. EVERY user receives this amount for sorting.* *Cycle the Oracle server when possible to allow the new value to take effect.* *Monitor new value to see if performance improves.* In addition to increasing the amount of memory devoted to sorting, find inefficient SQL that cause needless sorts. For example, UNION ALL does not cause a sort whereas UNION does in an SQL query (to eliminate duplicate rows). DISTINCT is frequently misapplied to SQL statements and causes unnecessary sort actions. There are times you simply cannot stop sort activity. This being the case, try to keep it in memory whenever possible. However, large data warehousing systems oftentimes simply exhaust RAM sort allotments, so if disk sorts must occur, try to ensure three things: *Your user's TEMPORARY TABLESPACE assignment is not the SYSTEM tablespace, which is the default assignment.* **NOTE:** *For Oracle9i or later, you can specify a default tablespace other than SYSTEM for every user account that is created.* *The TEMPORARY TABLESPACE assigned to your users is placed on a fast disk.* *The TEMPORARY TABLESPACE has the tablespace parameter TEMPORARY assigned to it, which allows sort activity to be performed in a more efficient manner.*

Parse/Execute Ratio

Each time a new SQL statement is submitted to Oracle, the kernel must 'parse' the statement, which involves syntax checks, security checks, and object validations. The Parse/Execute Ratio shows the percentage of SQL executed that did not incur a hard parse.

NOTE: This statistic is available on the Home page, Memory home page, and the Users home page.

Location	Oracle Performance Analyst statistics > Home > Key Ratio Analysis pane Oracle Performance Analyst statistics > Memory > Key Ratio Analysis pane Oracle Performance Analyst statistics > Users > Key User Analysis pane
Metrics	Seeing low values might indicate that users are firing SQL with many hard-coded literals instead of using bind variables within an application. High values (90% and greater) generally indicate Oracle is saving precious CPU by avoiding heavy parse tasks.
Correction	The best way to reduce unnecessary parse activity is to encourage SQL statement reuse. This can be done by promoting SQL execution through the use of stored procedures or applications where bind variables can be used. Oftentimes, literals in otherwise identical SQL statements can cause unneeded parse work for Oracle. The use of bind variables can counter that problem. If you determine that SQL literals are causing SQL to not be reused, do the following (in Oracle 8.1.6 and later): *Change the cursor_sharing parameter to FORCE by using the ALTER SYSTEM SET cursor_sharing=FORCE command.* *Monitor database to see if parse activity is reduced and library cache reloads shrink.* *If using an SPFILE, save the new configuration values so Oracle reuses them each time the database is stopped and re-started. If using an Init.ora file, add the cursor_sharing=FORCE parameter to the file.*

Buffer Busy Waits

Buffer busy waits occur when a process needs to access a data block in the buffer cache, but cannot because it is being used by another process. A wait event generally happens because a buffer is being read into the buffer cache by another process or the buffer is in the buffer cache, but cannot be switched to a compatible mode immediately.

Location	Oracle Performance Analyst statistics > Home > Key Ratio Analysis pane Oracle Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	Buffer busy waits normally center around contention for rollback segments, too small an INITRANS setting for tables, or insufficient free lists for tables.
Correction	On Oracle8i or earlier, the remedy for each situation would be increasing the number of rollback segments, or altering tables to have larger settings for INITRANS to allow for more transactions per data block, and more free lists. For Oracle9i or later, you can use the automatic segment management feature in Oracle9i locally-managed tablespaces to help make free list problems a thing of the past. Using an UNDO tablespace in 9i or later can help remedy any rollback contention problem. You can also obtain which objects have actually experienced buffer busy waits in Oracle9i or later by querying the sys.v_$segment_statistics. This view is not populated unless the configuration parameter statistics_level is set to TYPICAL or ALL.

Rollback Contention Ratio

Rollback segments are used by Oracle to hold data needed to rollback any changes made through inserts, updates, or deletes to various Oracle objects. They also allow Oracle to have read consistency for long running queries, are used for recovery purposes, and play a role during exports of database information. In a heavy transaction processing environment, rollback segments are accessed continuously and therefore are subject to contention problems. The Rollback Contention Ratio helps identify contention occurring on the system relating to rollbacks.

Location	Oracle Performance Analyst statistics > Home > Key Ratio Analysis pane Oracle Performance Analyst statistics > Objects > Home > Key Ratio Analysis pane
Metrics	Overall, if the rollback contention ratio approaches 1% or more, create more rollback segments. Also consider creating a specialized, larger, rollback segment to be used by long running transactions. Doing so alleviates dynamic rollback extensions and cuts down heavily on ORA-01555 Snapshot Too Old errors.
Correction	Begin by creating new rollback segments and altering them to be online for use. Then monitor the overall contention ratio to see if it begins to drop. If you are using Oracle9i or later, consider using an UNDO tablespace and allowing Oracle to automatically control rollback segment management.

Latch Miss Ratio

Protecting the many memory structures in Oracle's SGA are latches. They ensure that one and only one process at a time can run or modify any memory structure at the same instant. Much more restrictive than locks (which at least allow for some collective user interaction), latches have no queuing mechanism - so either you get it or you do not and are forced to continually retry.

The latch miss ratio defines the number of times a process obtained a willing-to-wait latch vs. missing the attempt.

Location	Oracle Performance Analyst statistics > Home > Key Ratio Analysis pane
Metrics	If the latch miss ratio exceeds 1%, you should take action to resolve the amount of latch contention.

Correction

Examine the details regarding the latch contention. Increasing the shared_pool_size can assist in latch problems also. The table below describes latches:

Latch	Description
Cache buffer chain latch	Protects paths to database block buffers in the buffer cache. High I/O loads tend to cause contention for this latch. You can alleviate contention somewhat by adding more buffers to the cache (through the db_block_buffers or db_cache_size parameter) or by adding more LRU latch chain latches with the db_block_lru_latches parameter.
Library cache latches	Protects cached SQL statements in the library cache area of the Oracle shared pool. Contention for this latch is the usual result of literals being used for SQL statements instead of bind variables.

Other routine latch contention problems used to include the redo allocation and redo copy latches, but these have pretty much been made obsolete in Oracle 8.1.5 and later.

Parallel Query Busy Ratio

Oracle's parallel query feature, when used properly, allows for terrific increases in performance for many SQL and utility operations. Parallel operations can be introduced through SQL hints, specified in an object's DDL, or used in command line arguments for utility programs (like SQL*Loader). To effectively service parallel query requests, ensure that enough query servers exist in the database instance. The Parallel Query Busy Ratio is an indicator of how busy all the servers are on the database in question.

Location	Oracle Performance Analyst statistics > Home > Key Ratio Analysis pane
Metrics	If the Parallel Query Busy Ratio approaches 80-90%, add more query servers to the database, or examining parallel requests to ensure they are being used in an efficient and necessary manner.
Correction	To fix, do the following: *Edit the Init.ora file or SPFILE for the database.* *Increase the amount of parallel_max_servers to a higher value.* *Cycle the Oracle server when possible to allow the new value to take effect.* *Monitor new value to see if performance improves.* You can also investigate the use of the parallel_automatic_tuning parameter in Oracle 8.1 and later.

Free Shared Pool Percent

Oracle's shared pool need not use all of the memory given to it through the shared_pool_size parameter. If the database does not have many object and code definitions to reference, then the shared pool can contain an abundance of free memory that is not being used.

Location	Oracle Performance Analyst statistics > Home > Key Ratio Analysis pane Oracle Performance Analyst statistics > Memory > Key Ratio Analysis pane
Metrics	Under-allocating the shared pool size can have a serious impact on your database's performance, but over-allocating the shared pool can have run time ramifications as well. If you have a good chunk of memory allocated to the Oracle shared pool that is never used, it might be more of a performance enhancement to reduce the shared pool amount and instead give the memory to the buffer/data cache, or even back to the operating system itself. In terms of knowing when to reduce the shared pool, a good benchmark is continually seeing 2-3MB of free memory. On the other hand, if after an hour or so of beginning database operation, you see that virtually no free memory is left in the shared pool, or you are seeing ORA-4031 errors (that indicate definitions cannot find enough contiguous free space in the shared pool), increase the pool by 10% or more.
Correction	If you continuously see little or no free memory in the shared pool, do the following: For Oracle8i or earlier: *Edit the Init.ora file for the database.* *Increase the amount of shared_pool_size to a higher value. Take caution to not over-allocate; ensure enough free memory exists on server before increasing value.* *Cycle the Oracle server when possible to allow the new value to take effect.* *Monitor the new value to see if performance improves.* For Oracle9i or later: *Increase the size of the shared_pool_size parameter through use of the ALTER SYSTEM SET shared_pool_size command value. Take caution to not over-allocate; ensure enough free memory exists on server before increasing value.* *Monitor the new value to see if performance improves.* *If using an SPFILE, save the new configuration values so Oracle reuses them each time the database is stopped and re-started.* You can also use the ALTER SYSTEM FLUSHED SHARED_POOL command to remove all objects from the shared pool and start with a clean slate.

Bottleneck Analysis pane

When an Oracle database is unning, every connected process is either busy doing work or waiting to perform work. A process that is waiting can mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. And this is where wait-based or bottleneck analysis comes into play. DBAs used this form of performance analysis to determine if perceived bottlenecks in a database are contributing to a performance problem.

Bottleneck analysis is a valid method of measuring performance because it helps you track where a database has been spending its time. If latch contention or heavy table scan activity has been dragging a database's performance down, you can use bottleneck analysis to confirm the actual root cause. Once one or more wait events or other bottlenecks have been pinpointed as possible performance vampires, you can determine which sessions and objects are causing the problem.

For wait event analysis to be properly carried out, it is imperative that the timed_statistics initialization parameter be set to TRUE. By default this parameter is set to FALSE, which disallows the collection of wait times for each wait event defined to the Oracle engine. For one to really understand the impact of wait events on database performance, you need to discover what the database is or has been waiting on, and the durations of the waits. Having both allows a complete picture to be formed regarding the magnitude of wait-initiated performance degradations. Almost all Oracle experts now agree that allowing the collection of timing statistics adds little if anything to database overhead, so setting timed_statistics to TRUE should not be a worry. The Performance Analyst Home page identifies the top system and session waits that can decrease performance.

When using bottleneck analysis, you cannot rely only on the information contained in the wait event views that Oracle provides. For example, an object can attempt to extend into another extent of space in a tablespace and yet be denied if no such free space exists. Such a failure is not reflected in any wait event, but still represents a very real bottleneck to the database. In the same way that you cannot depend on only a few ratios to properly carry out ratio-based performance analysis, an administrator must include other statistical metrics in their overall bottleneck analysis framework to obtain an accurate performance risk assessment. Performance Analyst works to identify bottlenecks in your database that fall outside of pure wait events so you can get a total picture of all stoppages in your system. The Performance Analyst Home page displays space-related bottlenecks as well as other sources of contention/bottlenecks that can decrease the overall performance of your system.

The following bottleneck indicators are used on the Performance Analyst Home page to succinctly communicate the general overall performance level of the monitored database:

Top System Bottlenecks	.Top Session Bottlenecks	Problem Tablespaces
Problem Objects	Current Object Blocks	Enqueue Waits
Free List Waits	Storage Analysis pane	

Top System Bottlenecks

When viewing wait statistics, there are many levels of detail that you can view. The first level is the system view, which provides a global, cumulative snapshot of all the waits that have occurred on a system. Viewing these numbers can help you determine which wait events caused the most commotion in a database so far. The Top System Bottlenecks section identifies the top waits that have occurred on the Oracle database based on the number of waits per event.

For Oracle 10g only you see a graph that provides a historical look at the last hour and allows you to see what wait classes experienced the most wait activity in terms of time waited. To get detail into which wait events have accumulated the most time, you can drill down into the System Waits view.

Location	Oracle Performance Analyst statistics > Home > Bottleneck Analysis pane

Correction	Appendix A in the Oracle Reference manual contains a listing and description of every current wait event defined in Oracle. DBAs unfamiliar with what each event represents should keep this listing close by as they examine wait-based event metrics. For example, a 'db file scattered read' event is typically indicative of table scan operations. If you see many of these events, then you can begin to see if large table scans are occurring in the database. Like the 'db file scattered read' event, each wait event has its own meaning and individual end-resolution diagnosis. After looking at system-level wait activity, you can discover which current connections are responsible for waits at the system level. Performance Analyst reports on historical and current wait events at the session level, making this investigation easy to accomplish

.Top Session Bottlenecks

When viewing wait statistics, there are many levels of detail that you can view. The first level is the system view, which provides a global, cumulative snapshot of all the waits that have occurred on a system. The second level is a historical look at waits from a session level. The third is which sessions are currently experiencing waits. The Top Session Bottlenecks section identifies the top sessions that are currently waiting, based on their wait time in seconds.

Location	Oracle Performance Analyst statistics > Home > Bottleneck Analysis pane
Correction	Oracle documentation contains a listing and description of every current wait event defined in Oracle. DBAs unfamiliar with what each event represents should keep this listing close by as they examine wait-based event metrics. The most common wait viewed at the session level is an 'enqueue' wait, which typically identifies lock contention. If enqueue waits are observed, then you can check the "Current Object Blocks" count on the Performance Analyst Home page, as well as the Users which displays locks and blocking locks detail. As with an enqueue event, each wait event has its own meaning and individual end-resolution diagnosis.

Problem Tablespaces

The Problem Tablespaces metric is a count of all tablespaces in the database whose free space percentage amount has fallen below a user-defined limit.

Location	Oracle Performance Analyst statistics > Home > Bottleneck Analysis pane
Metrics	A rule of thumb for dynamic tablespaces (those with growing objects) is to keep a minimum of 10-15% free space available for object growth.
Correction	There are two ways to prevent a tablespace from running out of available free space: *Turn AUTOEXTEND on for the underlying tablespace's datafiles. This allows them to automatically grow when free space in the datafile has been exhausted.* *Using the ALTER TABLESPACE … ADD DATAFILE… command, you can manually add a new datafile to a tablespace that is about to run out of available free space.*

Problem Objects

The Problem Objects statistic is a count of all objects in the database that are in danger of reaching their maximum extent limit or cannot allocate their next extent of free space because of a lack of overall or contiguous free space in their parent tablespace.

Location	Oracle Performance Analyst statistics > Home > Bottleneck Analysis pane
Metrics	Modify any object approaching their maximum extent limit or unable to allocate a new extent of space so they can continue to grow in size.

Correction	Depending on the situation, there are a number of things you can do to prevent object extent problems: *Turn AUTOEXTEND on for the underlying parent tablespace's datafiles. This allows a tablespace to automatically grow when free space in the datafile has been exhausted, and allows an object to extend even when little or no current free space is available.* *Using the ALTER TABLESPACE … ADD DATAFILE… command, you can manually add a new datafile to a tablespace that is about to run out of available free space.* *You can alter an object that is at or near their maximum extent limit so that the object has unlimited extents.* *With Oracle 8.1.5 and later, you can use locally-managed tablespaces to ensure that no object ever reaches its maximum extent limit, because all objects are allowed unlimited extents.* *An object can be reorganized into another tablespace or reorganized in general to reduce the number of extents the object currently takes up.*

Current Object Blocks

A single blocking user has the potential to stop work for nearly all other processes on a small system, and can cause major headaches even on large systems. Although Oracle supports unlimited row-level locking, blocking lock situations do crop up. Blocks are most often caused by user processes holding exclusive locks and not releasing them via a proper COMMIT frequency.

NOTE: This statistic is also called Sessions Blocked on the Users and Session Blocks on the Objects Home page.

Location	Oracle Performance Analyst statistics > Home > Bottleneck Analysis pane
Metrics	Investigate any indicator above zero immediately before the situation has a chance to mushroom.
Correction	Once discovered, a blocking lock situation can normally be quickly remedied. You can issue a KILL against the offending process, which eliminates the user's stranglehold on the objects they were accessing. Other user processes then nearly almost always complete in an instant. Discovering the blocked lock situation is made easier by using tools like Performance Analyst, but preventing the blocking lock situation in the first place is where it gets tricky. You can look at the Users detail and view all current blocking locks to see exactly which sessions are holding the currently restrictive locks. The culprit of blocking lock scenarios is usually the application design, or the SQL being used within the application itself. Properly coding an application to reference database objects in an efficient order, and then using the right SQL to get the job done, is an art. Most DBAs who have had to face Oracle Forms applications have suffered through the dreaded SELECT … FOR UPDATE statements that place unnecessary restrictive locks on nearly every read operation, and know all too well that good coding practice is important. The key to avoiding lock contention is to process user transactions in the quickest and most efficient manner possible - something not always easy to do.

Enqueue Waits

An enqueue is an advanced locking device that allows multiple database processes to share certain resources. Enqueue waits typically occur when sessions wait to be granted a requested lock. Sometimes these locks are internal Oracle locks while other times they could be locks for rows of data. Enqueues are issued implicitly by Oracle.

Location	Oracle Performance Analyst statistics > Home > .Top Session Bottlenecks Oracle Performance Analyst statistics > Objects > Bottleneck Analysis pane
Metrics	Investigate any enqueue waits that read consistently above one or more (delta statistics).

Correction	Removing contention for enqueues is almost always an application design issue. If many enqueue waits are seen, this normally indicates either contention for certain rows in the database, or the result of database-initiated lock escalation. Examine the use of indexes to make sure all referencing foreign keys are indexes and that SQL is tuned to not tarry over rows in the database during modification operations. Enqueue waits can also be the result of space management tasks (such as objects extending) and disk sorts (mainly in tablespaces that do not make use of the TEMPORARY tablespace parameter).

Free List Waits

Free lists are lists of Oracle data blocks that contain free space for an Oracle object. Every table has at least one free list. Free lists are used to locate free blocks of space when a request is made of a table for the insertion of a row. Free list contention can reduce the performance of applications when many processes are involved in the insertion of data to the same table.

Location	Oracle Performance Analyst statistics > Home > .Top Session Bottlenecks Oracle Performance Analyst statistics > Objects > Bottleneck Analysis pane
Metrics	If consistent numbers for free list waits continues to appear, add additional free lists to the most dynamic database objects through the use of the STORAGE parameter. Another indicator of insufficient free lists are consistent, non-zero numbers for the buffer busy wait event.
Correction	If free list waits are observed, add more free lists to tables and indexes with high insert rates. For Oracle9i or later, objects can be created inside of locally-managed tablespaces that use the automatic segment management feature, which entirely does away with free lists.

Storage Analysis pane

While DBAs focus on memory settings and tuning SQL, storage problems can play a major role in wrecking an otherwise well-running database.

Storage problems generally take one of two forms:

1. Hit-the-wall, which can bring things to a complete standstill.
2. Performance vampire, which slowly drains the performance of a database over time.

Storage problems have the capability to bring the curtain down on a database very quickly, as in the case of an archive log destination running out of free space. But storage problems can also silently work behind the scenes to slowly but surely rob a database of its performance. For example, you can access a hub table in a busy database when an application is first given life, but over time, if it develops a heavy migrated row problem, it can cause things to run very differently.

The Storage Analysis section of the Performance Analyst Home page reports on the total used and free space of all tablespaces in the database as well as whether the database is running in Archive Log mode. For more detailed information on database space, use the Performance Analyst Space.

Total Used Space/Total Free Space

These statistics represent the total used and free space available in all tablespaces/datafiles in the database. Although good to know, a more detailed listing by tablespace is needed to determine where any actual space shortages exist in the database.

Location	Oracle Performance Analyst statistics > Home > Storage Analysis pane
Metrics	If any one tablespace begins to approach 90% used (and the objects contained within it are dynamic and growing as opposed to static), take action to prevent any future space allocation errors.

Correction	Here are a some things you can do to prevent a tablespace from running out of available free space: *Turn AUTOEXTEND on for the underlying tablespace's datafiles. This allows them to automatically grow when free space in the datafile has been exhausted.* *Using the ALTER TABLESPACE … ADD DATAFILE… command, you can manually add a new datafile to a tablespace that is about to run out of available free space.*

Archive Log

Oracle can be placed into archivelog mode, which tells the database to make copies of the online redo log files for point-in-time recovery purposes. The Archive Log statistic displays whether the database is running in archivelog mode or not. This information is displayed in the Storage Analysis section of the Performance Analyst Home page. For more information, see Storage Analysis pane.

Location	Oracle Performance Analyst statistics > Home > Storage Analysis pane

SQL Analysis pane

Much of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured database down in terms of end user response times. SQL statements that use much memory can also cause a problem in a database.

Before you can identify problem SQL in your database, you have to ask the question, “What is bad SQL?” What criteria do you use when you begin the hunt for problem SQL in your critical systems? Understand that even the seasoned experts disagree on what constitutes efficient and inefficient SQL; so there is no way to sufficiently answer this question to every Oracle professional's satisfaction. The table below lists some general criteria you can use when evaluating the output from various database monitors or personal diagnostic scripts:

Criteria	Description
Overall Response (Elapsed) Time	The amount of time the query took to parse, execute, and fetch the data needed to satisfy the query. It should not include the network time needed to make the round trip from the requesting client workstation to the database server. This statistic is available in Oracle9i or later.
CPU Time	The amount of CPU time the query took to parse, execute, and fetch the data needed to satisfy the query.
Physical I/O	This is often used as the major statistic in terms of identifying good vs. bad SQL, this is a measure of how many disk reads the query caused to satisfy the user's request. While you certainly want to control disk I/O where possible, it is important that you not focus solely on physical I/O as the single benchmark of inefficient SQL. Make no mistake, disk access is slower than memory access and also consumes processing time making the physical to logical transition, but you need to look at the entire I/O picture of a SQL statement, which includes looking at a statements' logical I/O as well.
Logical I/O	the number of memory reads the query took to satisfy the user's request. The goal of tuning I/O for a query should be to examine both logical and physical I/O, and use appropriate mechanisms to keep both to a minimum.
Repetition	The number of times the query has been executed. A problem in this area is not as easy to spot as the others unless you know your application well. A query that takes a fraction of a second to execute can still cause a headache on your system if it is executed erroneously (for example, a query that executes in a runaway PL/SQL loop) over and over again.

There are other criteria that you can examine like sort activity or access plan statistics (that show items like Cartesian joins and the like), but more often than not, these measures are reflected in the criteria listed above. Fortunately, Oracle records all the above measures (some only in 9i), which makes tracking the SQL that has been submitted against an Oracle database much easier.

>

Metrics

When you begin to look for inefficient SQL in a database, there are two primary questions you want answered:

1 What HAS been the worst SQL that has historically run in my database?

2 What IS the worst SQL that is running right now in my database?

When troubleshooting a slow system, you should be on the lookout for any query that shows an execution count that is significantly larger than any other query on the system. It can be that the query is in an inefficient PL/SQL loop, or other problematic programming construct. Only by bringing the query to the attention of the application developers will you know if the query is being mishandled from a programming standpoint.

There is the possibility that the SQL statement just is not tuned well. To determine that, you can go further into Performance Analyst's Top SQL view and, if you have Embarcadero SQL Tuner installed, you can port the SQL over to SQL Tuner to better optimize the statement.

Workload Analysis pane

When your database experiences performance slowdowns, you should answer the following:

- Who is currently connected to the database?
- What resources are they using?
- What are they currently executing?

The Workload Analysis section of the Performance Analyst Home page provides insight into the leading resource hogs of a database as well as a count of total active and inactive processes. You can view detailed information into what each leading session is currently involved with.

The Workload Analysis section of the Home page includes:

Top Processes	Active User Processes	Inactive User Processes

The buffer cache hit ratio is an indicator of how often user requests for data are satisfied through memory vs. being physically read from disk. Data read from memory produces user response times many times faster than when that same data is read from disk. Keeping physical I/Os to an absolute minimum is one of the Oracle buffer cache's purposes in life.

The table below describes the key counters Performance Analyst uses to calculate the buffer cache hit ratio:

Key Counter	Description
DB BLOCK GETS	Data read from memory for DML operations.
CONSISTENT GETS	Data read from rollback segments in memory.
PHYSICAL READS	Data read physically from disk.
Direct Reads	Data read physically from disk that bypasses the buffer cache. Direct reads are filtered out of overall physical reads so an accurate cache hit ratio can be determined.

Dividing the data read from memory by data read from disk yields the cache hit ratio.

Metrics

To help ensure excellent performance, you want to keep your cache hit ratio in the neighborhood of 90% or higher. However, every database has its own ‘personality’ and can exhibit excellent performance with below average readings for the cache hit ratio. Excessive logical I/O activity can produce a very high cache hit ratio while actually degrading overall database performance.

Investigate consistent low readings of 60% or less.

NOTE: For Oracle8i or earlier, the adjustment of the db_block_buffers tuning parameter is required. For Oracle9i and later, the db_cache_size parameter is the parameter that needs attention. Any increases in db_block_buffers to take effect on Oracle8i or earlier, the database must be cycled. The db_cache_size parameter in Oracle9i or later, however, is dynamic and can be altered without stopping and starting the database instance.

Correction

If a problem is found in Oracle8i or earlier, do the following:

- Edit the Init.ora file for the database.
- Increase the amount of db_block_buffers to a higher value. Take caution to not over-allocate; ensure enough free memory exists on server before increasing value.
- Cycle the Oracle server when possible to allow the new value to take effect.
- Monitor the new value to see if performance improves.

If a problem is found in Oracle9i or later, do the following:

- Increase the size of the db_cache_size parameter through use of the ALTER SYSTEM SET db_cache_size command value (take caution to not over-allocate; ensure enough free memory exists on server before increasing value).
- Monitor the new value to see if performance improves.
- If using an SPFILE, save the new configuration values so Oracle reuses them each time the database is stopped and re-started.

Top Processes

When the database population as a whole experiences a system slowdown, it is not uncommon to find one or two users who are responsible for bringing the system to its knees. In the best of worlds, users have an evenly divided amount of memory usage, disk I/O, CPU utilization, and parse activity. Unfortunately, this is not usually the case. Users submit large batch jobs during peak OLTP activity, or when sessions that are firing off untuned queries on a critical system.

If you are seeing a slowdown in your database, and cannot seem to find a root cause, one thing to examine is the resource consumption of the leading sessions on a system. A handful of users can overwhelm the I/O capabilities of Oracle (through untuned queries or runaway batch jobs) or hammer the CPU or memory structures of the database and server.

Performance Analyst makes it easy to pinpoint the top sessions by showing the leading processes at the physical I/O, logical I/O, memory, and CPU usage levels.

Location	Oracle Performance Analyst statistics > Home > Workload Analysis pane
Metrics	If any one session uses more than 50% of a total resource (CPU, memory, etc.) go into the session to find out what they are currently executing.

Active User Processes

The Active User Processes statistic is the total number of active and open threads reported in the database. Active Sessions displays the number of processes actively performing work.

>

For more information, see Inactive User Processes.

Location	Oracle Performance Analyst statistics > Home > Workload Analysis pane

Inactive User Processes

This metric is the total number of threads logged on to the database that are idle at the current time.

Location	Oracle Performance Analyst statistics > Home > Workload Analysis pane
Metrics	A large number of inactive users could indicate user sessions that have mistakenly been left logged on. Because each user thread consumes a portion of memory on the Oracle server, sever any sessions not needing a connection to reduce resource usage.
Correction	Double-click this statistic to open the Top Sessions Tab of the Users Detail page. On this tab you can check sessions that have many seconds idle and/or that have been logged on for very long periods of time as indicated by the logon time column. After verifying that the session is no longer necessary, you can KILL the session. For more information, see Active User Processes.

Memory

In addition to a Home page, the Memory category of Oracle Performance Analyst includes the following tabbed pages:

Buffer Cache Tab	Buffer Pools Tab	Dictionary Cache Tab
Latch Activity Tab	Library Cache Tab	Shared Pool Tab
PGA Tab		

Home

The Memory home page includes the following sections:

Bottleneck Analysis pane	Key Ratio Analysis pane	SGA Analysis pane
SQL Analysis pane	Workload Analysis pane	

Key Ratio Analysis pane

Object-related database activity can be examined using both ratio-based and wait/bottleneck-based analysis. Ratio-based analysis involves examining a number of key database ratios and statistical readings that can be used to indicate how active certain object types are. Performance ratios serve as roll-up mechanisms for busy DBAs to use for at-a-glance performance analysis.

When using ratio-based analysis, there are some standards to adhere to. To start with, many of the formulas that make up ratio-based analysis must be derived from delta measurements instead of cumulative statistics. Many of the global ratios that you examines come from the v$sysstat performance view. The performance view maintains a count of all the occurrences (in the VALUE column) of a particular database incident (in the NAME column) since the database was brought up. For databases that are kept up for long periods of time, these values can grow quite large and impacts how a particular ratio is interpreted. However, if delta statistics are used (taking, for a specified sampling period, the before and after counts of each statistic that make up a ratio), then an accurate and current portrayal of the various ratios can be had.

A final thing to remember about using ratio-based analysis is that, while there are a number of rules of thumb that can be used as starting points to begin the evaluation of database performance, DBAs must determine each database's individual 'personality' with respect to the key performance ratios. Some hard and fast rules simply do not apply to every database. The danger in using blanket ratio standards is that they can lead you to haphazardly take action, which can at times contribute nothing to the situation, and sometimes even degrade performance.

The following memory ratios are used on the Performance Analyst Memory home page to succinctly communicate the general overall memory performance levels of the monitored database. They duplicate statistics shown on the Home page. For detailed descriptions of those statistics, see Key Ratio Analysis pane.

Buffer Cache Hit Ratio	Library Cache Hit Ratio	Dictionary Cache Hit Ratio
Memory Sort Ratio	Free Shared Pool Percent	Parse/Execute Ratio

Bottleneck Analysis pane

When an Oracle database is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting may mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. You can use Bottleneck Analysis to determine if perceived bottlenecks in a database are contributing to a performance problem.

Memory bottlenecks can definitely cause performance degradation in an otherwise well-running database. Typically, these bottlenecks center around Oracle's buffer/data cache, library cache, and occasionally log buffer memory regions. To help you identify such problems, the following statistics are presented on the Memory home page:

Free Buffer Wait Average	Object Reloads	Redo Log Space Waits
Redo Log Space Wait Time	Top Latch Misses	

The following statistic, available on this pane, duplicates a statistic available on the Key Ratio Analysis pane of the Home page:

- Buffer Busy Waits

Free Buffer Wait Average

There are times when a user must wait to obtain a clean buffer in the buffer/data cache. The free buffer wait average provides a metric that communicates the average number of tree buffer waits..

Location	Oracle Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	Seeing consistent non-zero numbers for the free buffer wait average might indicate an undersized data/buffer cache.

Correction	If the Object Reloads statistic, along with other shared pool metrics, indicates an undersized shared pool, you can do the following depending on the version of Oracle you are running: If a problem is found in Oracle8i or earlier: *Edit the Init.ora file for the database.* *Increase the amount of shared_pool_size to a higher value. Take caution to not over-allocate; ensure enough free memory exists on server before increasing value.* *Cycle the Oracle server when possible to allow the new value to take effect.* *Monitor the new value to see if performance improves.* If a problem is found in Oracle9i or later: *Increase the size of the shared_pool_size parameter through use of the ALTER SYSTEM SET shared_pool_size command value. Take caution to not over-allocate; ensure enough free memory exists on server before increasing value.* *Monitor the new value to see if performance improves.* *If using an SPFILE, save the new configuration values so Oracle reuses them each time the database is stopped and re-started.* If you suspect that there are hot code objects that would benefit from being pinned into the shared pool, then you can go into shared pool details using Performance Analyst's monitors. You can easily spot the general code object types or specific objects that should be pinned in the pool from these monitors. Once identified, you can use the dbms_shared_pool package to pin the hot code objects into RAM.

Object Reloads

When a SQL statement or code object (stored procedure, trigger, etc.) is executed for the first time, Oracle loads the code definition into the shared pool. If the statement or code object is executed again, Oracle can reuse the object definition in memory, and thereby shorten execution time.

If a code object is aged out of the shared pool, and then re-requested, Oracle must load the object back into the pool. Such activity can lead to reduced response times. The Object Reloads statistic provides you with a count of how many objects have been reloaded back into the shared pool.

Location	Oracle Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	Seeing consistent non-zero numbers for object reloads could indicate a too small shared pool or that help identify that there are hot code objects that should be pinned in the shared pool so they cannot be removed.

Correction	If the Object Reloads statistic, along with other shared pool metrics, indicates an undersized shared pool, you can do the following depending on the version of Oracle you are running: If a problem is found in Oracle8i or earlier: *Edit the Init.ora file for the database.* *Increase the amount of shared_pool_size to a higher value. Take caution to not over-allocate; ensure enough free memory exists on server before increasing value.* *Cycle the Oracle server when possible to allow the new value to take effect.* *Monitor the new value to see if performance improves.* If a problem is found in Oracle9i or later: *Increase the size of the shared_pool_size parameter through use of the ALTER SYSTEM SET shared_pool_size command value. Take caution to not over-allocate; ensure enough free memory exists on server before increasing value.* *Monitor the new value to see if performance improves.* *If using an SPFILE, save the new configuration values so Oracle reuses them each time the database is stopped and re-started.* If you suspect that there are hot code objects that would benefit from being pinned into the shared pool, then you can go into shared pool details using Performance Analyst's monitors. You can easily spot the general code object types or specific objects that should be pinned in the pool from these monitors. Once identified, you can use the dbms_shared_pool package to pin the hot code objects into RAM.

Redo Log Space Waits

The Oracle database is able to manage recovery by recording all changes made to a database through the use of redo log files. Oracle writes modifications made to a database to these files that have the capability to be archived off to another medium for disaster recovery. The background process that performs these operations is Oracle's Log Writer (LGWR). There is a buffer area in Oracle's System Global Area (SGA) that is used to reduce redo log file I/O, whose size, or lack thereof, can affect performance in a busy system. Sometimes a user process must wait for space in this redo log buffer. Oracle uses the log buffer area to cache redo entries prior to writing them to disk, and if the buffer is not large enough for the redo entry load, waits can occur..

Location	Oracle Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	The two main numbers to watch are the redo log space requests and the redo log wait time. If either statistic strays too far from zero, increase the log_buffer parameter and add more memory to the redo log buffer.
Correction	To fix, do the following: *Edit the Init.ora file or SPFILE for the database.* *Increase the amount of log_buffer to a higher value (take caution to not over-allocate; ensure enough free memory exists on server before increasing value.* *Cycle the Oracle server when possible to allow the new value to take effect.* *Monitor new value to see if performance improves.* On some Oracle platforms, when adjusting the log_buffer parameter, make sure you make the amount a multiple of the block size. Otherwise, on database startup, Oracle returns an error stating that you have entered an invalid amount for the redo log buffer. **NOTE:** Oracle increases the log_buffer parameter if you make it smaller than its default size for a given platform.

Redo Log Space Wait Time

The Oracle database is able to manage recovery by recording all changes made to a database through the use of redo log files. Oracle writes modifications made to a database to these files that have the capability to be archived off to another medium for disaster recovery. The background process that performs these operations is Oracle's Log Writer (LGWR). There is a buffer area in Oracle's System Global Area (SGA) that is used to reduce redo log file I/O, whose size, or lack thereof, can affect performance in a busy system. Sometimes a user process must wait for space in this redo log buffer. Oracle uses the log buffer area to cache redo entries prior to writing them to disk, and if the buffer is not large enough for the redo entry load, waits can occur.

Location	Oracle Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	The two main numbers to watch are the redo log space requests and the redo log wait time. If either statistic strays too far from zero, increase the log_buffer parameter and add more memory to the redo log buffer.
Correction	To fix, do the following: *Edit the Init.ora file or SPFILE for the database.* *Increase the amount of log_buffer to a higher value (take caution to not over-allocate; ensure enough free memory exists on server before increasing value.* *Cycle the Oracle server when possible to allow the new value to take effect.* *Monitor new value to see if performance improves.* On some Oracle platforms, when adjusting the log_buffer parameter, make sure you make the amount a multiple of the block size. Otherwise, on database startup, Oracle returns an error stating that you have entered an invalid amount for the redo log buffer. **NOTE:** Oracle increases the log_buffer parameter if you make it smaller than its default size for a given platform.

Top Latch Misses

Latches protect the many memory structures in Oracle's SGA. They ensure that one and only one process at a time can run or modify any memory structure at the same instant. More restrictive than locks (which at least allow for some collective user interaction), latches have no queuing mechanism.

The Top Latch Misses chart indicates the top latch waits as ranked by total misses.

Location	Oracle Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	Seeing increasing numbers for the same latch miss can require further accessing session latch details, which can be accomplished inside Performance Analyst.
Correction	The remedy for latch misses depends on the latch itself. Buffer chain latches, for example, are indicative of high I/O rates in databases. Increasing the buffer/data cache can help as might adding more buffer LRU chain latches by increasing the db_block_lru_latches parameter. Library cache latch misses are oftentimes produced by non-reused SQL in the shared pool. Increasing the shared_pool_size can help, but a better remedy might be changing SQL to use bind variables or setting the cursor_sharing parameter to FORCE (Oracle 8.1.6 and later).

SQL Analysis pane

Much of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured database down in terms of end user response times. SQL statements that use much memory can also cause a problem in a database. The SQL Analysis for memory shows what SQL statements have consumed the largest percentages of shareable, persistent, and runtime memory

Before you can identify problem SQL in your database, you have to ask the question, “What is bad SQL?” What criteria do you use when you begin the hunt for problem SQL in your critical systems? Understand that even the seasoned experts disagree on what constitutes efficient and inefficient SQL; so there is no way to sufficiently answer this question to every Oracle professional's satisfaction. The table lists some general criteria you can use when evaluating the output from various database monitors or personal diagnostic scripts:

Criteria	Description
Overall Response (Elapsed) Time	The time the query took to parse, execute, and fetch the data needed to satisfy the query. It should not include the network time needed to make the round trip from the requesting client workstation to the database server. This statistic is available in Oracle9i or later.
CPU Time	The CPU time the query took to parse, execute, and fetch the data needed to satisfy the query.
Physical I/O	This is often used as the major statistic in terms of identifying good vs. bad SQL, this is a measure of how many disk reads the query caused to satisfy the user's request. While you certainly want to control disk I/O where possible, it is important that you not focus solely on physical I/O as the single benchmark of inefficient SQL. Make no mistake, disk access is slower than memory access and also consumes processing time making the physical to logical transition, but you need to look at the entire I/O picture of a SQL statement, which includes looking at a statements' logical I/O as well.
Logical I/O	The number of memory reads the query took to satisfy the user's request. The goal of tuning I/O for a query should be to examine both logical and physical I/O, and use appropriate mechanisms to keep both to a minimum.
Repetition	The number of times the query has been executed. A problem in this area is not as easy to spot as the others unless you know your application well. A query that takes a fraction of a second to execute can be a headache on your system if it has executed erroneously (for example, a query that executes in a runaway PL/SQL loop) over and over again.

There are other criteria that you can examine like sort activity or access plan statistics (that show items like Cartesian joins and the like), but more often than not, these measures are reflected in the criteria listed above. Fortunately, Oracle records all the above measures (some only in 9i), which makes tracking the SQL that has been submitted against an Oracle database much easier.

Metrics

When you begin to look for inefficient SQL in a database, there are two primary questions you want answered:

1 What HAS been the worst SQL that has historically been run in my database?

2 What IS the worst SQL that is running right now in my database?

When troubleshooting a slow system, you should be on the lookout for any query that shows an execution count that is significantly larger than any other query on the system. It could be that the query is in an inefficient PL/SQL loop, or other problematic programming construct. Only by bringing the query to the attention of the application developers will you know if the query is being mishandled from a programming standpoint.

There is the possibility that the SQL statement just is not tuned well. You can view Performance Analyst's Top SQL view and, if you have Embarcadero SQL Tuner installed, you can port the SQL over to SQL Tuner to better optimize the statement.

SGA Analysis pane

Most DBAs know all about the Oracle System Global Area (SGA). The SGA is Oracle's memory structural area devoted to facilitating the transfer of data and information between clients and the Oracle database. The table below describes Oracle memory structures:

Memory Structure	Description
Default buffer cache	Maintains data blocks when they are read from the database. If you do not specifically place objects in another data cache, then any data requested by clients from the database is placed into this cache. The memory area is controlled by the db_block_buffers parameter in Oracle8i and earlier and db_cache_size in Oracle9i or later.
Keep buffer cache	For Oracle 8 or later, you can assign various objects to a special cache that retains those object's requested blocks in RAM for as long as the database remains up. The keep cache's main is for often-referenced lookup tables that should be kept in memory at all times for fast access. The buffer_pool_keep parameter controls the size of this cache in Oracle8, while the db_keep_cache_size parameter handles the cache in Oracle9i or later. The keep pool is a sub-pool of the default buffer cache.
Recycle buffer cache	The opposite of the keep cache. When large table scans occur, the data that fills a memory cache will likely not be needed again and should be quickly discarded from RAM so that they do not occupy memory space and prevent needed blocks from assuming their place. Objects containing such data can be assigned to the recycle pool to ensure that such a thing does indeed occur. The buffer_pool_recycle parameter controls the size of this cache in Oracle8 and earlier, while the db_recycle_cache_size parameter handles the cache in Oracle9i or later.
Specific block size caches	For Oracle 8 or later, you can create tablespaces whose blocksize differs from the overall database blocksize. When data is read into the SGA from these tablespaces, their data has to be placed into memory regions that can accommodate their special block size. Oracle9i or later has memory settings for 2K, 4K, 8K, 16K, and 32K caches. The configuration parameter names are in the pattern of db_nk_cache_size.
Shared pool	Holds object structure as well as code definitions, and other metadata. Setting the proper amount of memory in the shared pool assists a great deal in improving overall performance with respect to code execution and object references. The shared_pool_size parameter controls the memory region.
Large pool	For Oracle 8 or later, you can configure an optional, specialized memory region called the large pool that holds items for shared server operations, backup and restore tasks, and other miscellaneous things. The large_pool_size parameter controls the memory region. The large pool is also used for sorting when the multi-threaded server (MTS) is implemented.
Java pool	Handles the memory for Java methods, class definitions, etc. The java_pool_size parameter controls the amount of memory for this area.
Redo log buffer	Buffers modifications that are made to the database before they are physically written to the redo log files. The log_buffer configuration parameter controls the memory area.

Oracle also maintains a fixed area in the SGA that contains a number of atomic variables, pointers, and other miscellaneous structures that reference areas of the SGA.

Workload Analysis pane

It is not uncommon for one or two users to cause the majority of runtime problems that plague a database. The problem could be a runaway process, an untuned batch procedure, or other user-initiated operation. Oftentimes, user connections can get out of hand with memory consumption, and extreme cases have caused headaches at both the database and operating system level (ORA-4030 errors).

If your server does not have an overabundance of memory, periodically check to see who your heavy memory users are along with the total percentage of memory each takes up. If you see one or two users who have more than 25-50% of the total memory usage, investigate the sessions further to see the activities they are performing.

>

Buffer Cache Tab

The Buffer Cache tab includes the following sections:

Buffer Cache Hit Ratio	Buffer Cache Utilization	Cached Blocks by Tablespace
Cached Tables		

Cached Blocks by Tablespace

Oracle's buffer/data cache contains blocks read in from disk at the user's request. The Cached Blocks by Tablespace section displays the amount of data held currently in the cache by tablespace.

Location	Oracle Performance Analyst statistics > Memory > Buffer Cache Tab

Cached Tables

One of the ways you can help data from a table stay in the buffer/data cache as long as possible is to enable the table's CACHE property. CACHE specifies that blocks retrieved for a table be placed at the most recently used end of the LRU list in the buffer cache when a full table scan is performed. This is particularly useful for small lookup tables.

The Cached Tables section displays the data for all tables that have their CACHE property enabled. The table below lists the information available in this section:

Column	Description
Owner	The owner of the table.
Table Name	The name of the table.
Table Size	The size of the table in bytes.

Location	Oracle Performance Analyst statistics > Memory > Buffer Cache Tab
Metrics	CACHE is best suited for small lookup tables. Therefore, if you see a large table that has been cached, you may want to reevaluate whether it is smart to cache the contents of that table in Oracle's buffer/data cache.
Correction	If you wish to disable the CACHE property of a table, you can alter the table and set the property to false.

Buffer Cache Utilization

The Buffer Cache Utilization section shows the state and use of the buffers currently in the buffer cache along with a count of how many there are for each state. Typically, the statuses of the buffers are: being read, free, read and modified, and read and not modified.

Location	Oracle Performance Analyst statistics > Memory > Buffer Cache Tab
Metrics	This grouping of statistics is quite helpful in determining if you have an overabundance of block buffers allocated. Consistently seeing large number of free buffers will clue you in to the fact that you can reduce the amount of block buffers in the cache and give memory back to the operating system. If, however, you see no free buffers within the first hour of bringing your Oracle database up, then you may consider adding more buffers to the cache.

Correction	If a problem is found, then: *Edit the Init.ora file or spfile for the database.* *Increase the amount of db_block_buffers (Oracle 8i and earlier) or db_cache_size (Oracle9i and later) to a higher value if free buffers are not found (take caution to not over-allocate; ensure enough free memory exists on server before increasing value). Reduce the number of an overabundance of FREE buffers are present.* *Cycle the Oracle server (for Oracle8i and earlier) when possible to allow the new value to take effect.* *Monitor new value to see if more numbers of free buffers show up for the FREE status.*

Dictionary Cache Tab

Oracle's data dictionary is a component of the shared pool that contains system elements necessary for the parsing of SQL statements, security resolution, object definitions, and more.

The overall data dictionary cache hit ratio provides a key indicator in determining the performance of the shared pool, and shows how often object definitions are found in memory vs. having to read them in from disk. Because Oracle references the data dictionary many times when an SQL statement is processed, it is imperative that as much as possible of this vital reference information be kept in RAM.

The Dictionary Cache tab shows the individual elements of the data dictionary along with their associated hit ratios. In versions 6.x of Oracle, these individual elements could be tuned, but in versions 7.x and later, the only main method for tuning them involves the adjustment of the entire shared pool setting. Although not tunable from an individual parameter level, each displayed element gives insight into which area of the data dictionary is either adding to or detracting from overall performance.

The table below describes the information available on this tab:

Column	Description
Parameter	The name of the individual data dictionary element.
Usage	The number of cache entries that contain valid data.
Gets	The number of requests for this element.
Get Misses	The number of requests resulting in a cache miss.
Hit Ratio	The ratio of cache hits versus misses of total requests. The maximum is 100%.
Scans	The number of scan requests.
Scan Misses	The number of times that a scan failed to find the needed data in the cache.
Scan Completes	The number of times the list was scanned completely.
Modifications	The number of insert, update, and delete actions.
Flushes	The number of disk flushes.

>

Metrics

An overall high data dictionary cache hit ratio is desirable, as are high hit ratios in each individual parameter. You should strive for a hit ratio between 90-100%, with 95% being a good performance benchmark.

NOTE: When a database is first started, the data dictionary cache hit ratio is not at an optimal level because all references to object definitions are relatively new, and as such, must be placed into the shared pool. Look for hit ratios between 80-90% for new database startups. If, however, after an hour or two of steady database time, the data dictionary cache hit ratio and individual hit ratios, have not increased to desirable levels, you should look into the possibility of increasing the shared_pool_size parameter.

NOTE: Databases supporting applications that involve large number of objects (such as an Oracle Financials installation) should have larger than normal shared pools to support the required object definitions.

Buffer Pools Tab

The Buffer Pools tab includes the following sections:

Buffer Pool Hit Ratio	Buffer Pool Memory Configuration	Buffer Pool Object Assignments

Buffer Pool Hit Ratio

Because data objects can be referenced with different usage patterns, Oracle8 and later offers the option to intelligently place objects into one of three buffer caches. The table below describes the types of buffer caches:

Cache	Description
KEEP	Designed to minimize misses in the buffer cache. Small objects that are frequently referenced are candidates for the KEEP buffer pool.
RECYCLE	Designed to avoid having large numbers of infrequently accessed blocks of data crowd out objects that need to be referenced in RAM, the RECYCLE pool is good for large objects that are scanned from beginning to end without the need for keeping all their data in RAM.
DEFAULT	The traditional cache for all other data objects.

NOTE: Unless you specify the KEEP or RECYCLE buffer cache, Oracle automatically places objects into the DEFAULT buffer cache.

The Buffer Pool Hit Ratio section the hit ratios for each of the Oracle8 and later buffer caches so you can easily see how often the objects placed into the various caches are being referenced in memory. Examining how often data is satisfied from memory vs. disk will help you determine if the caches are large enough and if they are being used in an optimal manner. The table below describes the information available in this section:

Column	Description
Buffer pool name	The name of the Oracle buffer pool.
Buffer pool hit ratio	The overall hit ratio for the particular cache.

NOTE: The DEFAULT buffer cache will only be shown for those installations not using the specialized caches available in Oracle8 and later.

Location	Oracle Performance Analyst statistics > Memory > Buffer Pools Tab

Metrics	The KEEP buffer pool should maintain a hit ratio as close to 100% as possible. However, the buffer pool hit ratio is not near 100% until the database has been up and running in a typical state for a short time. A poor hit ratio for the RECYCLE buffer pool may not be a bad thing since there is little chance of reusing a block stored in the buffer pool before it is aged out. **NOTE:** If you place objects into the KEEP buffer pool, you should periodically reexamine their object sizes to ensure that they are not growing to a physical state that will jeopardize the performance of the KEEP pool.

Buffer Pool Object Assignments

Because data objects can be referenced with different types of usage patterns, Oracle8 and later offers the option to intelligently place objects into one of three buffer caches. The table below describes these three buffer caches:

Cache	Description
KEEP	Designed to minimize misses in the buffer cache. Small objects that are frequently referenced are candidates for the KEEP buffer pool.
RECYCLE	Designed to avoid having large numbers of infrequently accessed blocks of data crowd out objects that need to be referenced in RAM, the RECYCLE pool is good for large objects that are scanned from beginning to end without the need for keeping all their data in RAM.
DEFAULT	The traditional cache for all other data objects.

Objects not specified to go into either the KEEP or RECYCLE caches will automatically be placed into the DEFAULT buffer cache.

The Buffer Pool Object Assignments section provides a summary with respect to which types of objects have been assigned to the various buffer caches. The table below describes the information available in this section:

Column	Description
Object Type	The type of object (table, index, etc.).
Buffer Pool	The buffer cache name (KEEP, etc.).
Total Objects Assigned	The number of objects that have been assigned to the particular cache.

Location	Oracle Performance Analyst statistics > Memory > Buffer Pools Tab
Metrics	When looking at overall database I/O activity, you should keep an eye out for objects that may lend themselves to being placed into a particular buffer cache. Consider using a KEEP cache for relatively small, frequently accessed tables that require fast response times. Large tables with random I/O activity and are scanned from beginning to end a lot are good candidates for a RECYCLE cache.

Buffer Pool Memory Configuration

Because data objects can be referenced with different types of usage patterns, Oracle8 and later offers the option to intelligently place objects into one of three buffer caches. The table below describes these three buffer caches:

Cache	Description
KEEP	Designed to minimize misses in the buffer cache. Small objects that are frequently referenced are candidates for the KEEP buffer pool.
RECYCLE	Designed to avoid having large numbers of infrequently accessed blocks of data crowd out objects that need to be referenced in RAM, the RECYCLE pool is good for large objects that are scanned from beginning to end without the need for keeping all their data in RAM.

>

Cache	Description
DEFAULT	The traditional cache for all other data objects.

Objects not specified to go into either the KEEP or RECYCLE caches will automatically be placed into the DEFAULT buffer cache.

Location	Oracle Performance Analyst statistics > Memory > Buffer Pools Tab
Metrics	When looking at overall database I/O activity, you should keep an eye out for objects that may lend themselves to being placed into a particular buffer cache. Consider using a KEEP cache for relatively small, frequently accessed tables that require fast response times. Large tables with random I/O activity and are scanned from beginning to end a lot are good candidates for a RECYCLE cache.
Correction	Objects can be placed into different buffer pools at object creation time (using the STORAGE...BUFFER_POOL option) or existing objects can be set into a different pool with the ALTER command. Note that tables, partitions, and indexes may be placed into the different caches. If you just want to use the DEFAULT buffer pool and not enable any special caches, you can still encourage Oracle to keep certain objects in the buffer cache as long as possible using the CACHE parameter. For example, issuing the command ALTER TABLE … CACHE specifies that the blocks retrieved for this table be placed at the most recently used end of the LRU list in the DEFAULT buffer cache when a full table scan is performed. The CACHE hint may also be used in SQL statements to cache a table, but used in this form, the blocks will only be cached until the next time the database is shut down. Once the database comes back up, the CACHE hint would have to be issued in a SQL statement again to cache the needed blocks of data.

Latch Activity Tab

The Latch Activity tab includes the following sections:

Immediate Latch Hit Ratio	Latch Detail	Latch Hit Ratio
Latch Sleep Ratio		

Latch Hit Ratio

Protecting the many memory structures in Oracle's SGA are latches. They ensure that one and only one process at a time can run or modify any memory structure at the same instant. Much more restrictive than locks (which at least allow for some collective user interaction), latches have no queuing mechanism so either you get it or you do not and are forced to continually retry. The latch hit ratio defines the number of times a process obtained a willing-to-wait.

Location	Oracle Performance Analyst statistics > Memory > Latch Activity Tab
Metrics	If the latch hit ratio falls below 99%, then action should be taken to resolve the amount of latch contention occurring.

>

Correction

The details regarding the latch contention should be examined. Increasing the shared_pool_size can assist in latch problems also. There are a few specialized cases of latch contention that can be rectified pretty easily. The table below describes them:

Contention	Description
Cache buffer chain latch	This latch is responsible for protecting paths to database block buffers in the buffer cache. Very high I/O loads tend to cause contention for this latch. You can alleviate contention somewhat by adding more buffers to the cache (through the db_block_buffers/db_cache_size parameter) or by adding more LRU latch chain latches with the db_block_lru_latches parameter.
Library cache latches	Protects cached SQL statements in the library cache area of the Oracle shared pool. Contention for this latch is the usual result of literals being used for SQL statements instead of bind variables.

Other routine latch contention problems used to include the redo allocation and redo copy latches, but these have pretty much been made obsolete in Oracle 8.1.5 and later.

Immediate Latch Hit Ratio

Protecting the many memory structures in Oracle's SGA are latches. They ensure that one and only one process at a time can run or modify any memory structure at the same instant. Much more restrictive than locks (which at least allow for some collective user interaction), latches have no queuing mechanism so either you get it or you do not and are forced to continually retry. The latch immediate miss ratio defines the number of times a process obtained a not-willing-to-wait latch.

Location	Oracle Performance Analyst statistics > Memory > Latch Activity Tab
Metrics	If the value falls below 99%, then action should be taken to resolve the amount of latch contention occurring.

Correction

The details regarding the latch contention should be examined. Increasing the shared_pool_size can assist in latch problems also. There are a few specialized cases of latch contention that can be rectified pretty easily. The table below describes them:

Contention	Description
Cache buffer chain latch	This latch is responsible for protecting paths to database block buffers in the buffer cache. Very high I/O loads tend to cause contention for this latch. You can alleviate contention somewhat by adding more buffers to the cache (through the db_block_buffers/db_cache_size parameter) or by adding more LRU latch chain latches with the db_block_lru_latches parameter.
Library cache latches	Protects cached SQL statements in the library cache area of the Oracle shared pool. Contention for this latch is the usual result of literals being used for SQL statements instead of bind variables.

Other routine latch contention problems used to include the redo allocation and redo copy latches, but these have pretty much been made obsolete in Oracle 8.1.5 and later.

Latch Sleep Ratio

Protecting the many memory structures in Oracle's SGA are latches. They ensure that one and only one process at a time can run or modify any memory structure at the same instant. Much more restrictive than locks (which at least allow for some collective user interaction), latches have no queuing mechanism so either you get it or you do not and are forced to continually retry.

A sleep indicates that a latch could not be obtained for a process, and that the attempt will be retried. A low ratio indicates many processes that had to sleep multiple times before obtaining a requested latch.

Location	Oracle Performance Analyst statistics > Memory > Latch Activity Tab
Metrics	Sleeps should be kept as low as possible. If the overall sleep ratio falls below 99%, then action should be taken to resolve the amount of latch contention that is occurring.

Correction

The details regarding the latch contention should be examined. Increasing the shared_pool_size can assist in latch problems also. There are a few specialized cases of latch contention that can be rectified pretty easily. The table below describes them:

Contention	Description
Cache buffer chain latch	This latch is responsible for protecting paths to database block buffers in the buffer cache. Very high I/O loads tend to cause contention for this latch. You can alleviate contention somewhat by adding more buffers to the cache (through the db_block_buffers/db_cache_size parameter) or by adding more LRU latch chain latches with the db_block_lru_latches parameter.
Library cache latches	Protects cached SQL statements in the library cache area of the Oracle shared pool. Contention for this latch is the usual result of literals being used for SQL statements instead of bind variables.

Other routine latch contention problems used to include the redo allocation and redo copy latches, but these have pretty much been made obsolete in Oracle 8.1.5 and later.

Latch Detail

Protecting the many memory structures in Oracle's SGA are latches. They ensure that one and only one process at a time can run or modify any memory structure at the same instant. Much more restrictive than locks (which at least allow for some collective user interaction), latches have no queuing mechanism so either you get it or you do not and are forced to continually retry.

The Latch Detail section presents a detailed view of latch activity. The table below lists the information available:

Column	Description
Latch Name	The name of the latch.
Gets	The number of times the latch was requested by a process.
Misses	The number of failed attempts to acquire the latch on the first attempt.
Immediate Gets	The number of nowait requests for a latch.
Immediate Misses	The number of failed nowait attempts to acquire the latch on the first attempt.
Sleeps	The number of requests that paused while waiting for a latch.

>

Location	Oracle Performance Analyst statistics > Memory > Latch Activity Tab

Metrics

Common indicators of latch contention are a latch miss ratio (which records willing-to-wait mode latch requests) and latch immediate miss ratio (which records no-wait mode latch requests). These statistics reflect the overall health of how often latch requests were made and satisfied without waiting. If either of these exceed 1%, then latch contention may be causing system slowdowns.

There are a few latch contention situations that you can recognize and get to the bottom of quickly. The table below describes them:

Contention Situation	Description
Cache buffer chain latch	This latch is responsible for protecting paths to database block buffers in the buffer cache. Very high I/O loads tend to cause contention for this latch. You can alleviate contention somewhat by adding more buffers to the cache (through the db_block_buffers parameter) or by adding more LRU latch chain latches with the db_block_lru_latches parameter.
Library cache latches	This latch protects cached SQL statements in the library cache area of the Oracle shared pool. Contention for this latch is the usual result of literals being used for SQL statements instead of bind variables.

Other routine latch contention problems used to include the redo allocation and redo copy latches, but these have pretty much been made obsolete in Oracle 8.1.5 and later.

Library Cache Tab

The Library Cache tab includes the following sections:

Library Cache Hit Ratio	Library Cache Performance	Pinned Objects

Library Cache Performance

The library cache holds commonly used SQL statements - basically database code objects. A method for improving performance in Oracle is to encourage the reuse of SQL statements so expensive parse operations may be avoided. The library cache assists this tuning effort.

The Library Performance sectioned tab provides insight into how efficiently the library cache is operating. The table below describes the information available in this section:

Column	Description
Namespace	The region of the library cache.
Gets	The number of times a lock was requested for objects in the particular namespace.
Get Hit Ratio	The percentage of times (with 100% being the maximum) that the object was found in the cache.
Pins	The number of times a pin was requested for objects of this namespace.
Pin Hit Ratio	The percentage of times (with 100% being the maximum) that pin requests were successful.
Reloads	The number of times a piece of the object had to be brought back in from disk to the cache, most likely because it was flushed from the shared pool.

Column	Description
Invalidations	The number of times objects in this namespace were marked invalid because a dependent object was modified.

Location	Oracle Performance Analyst statistics > Memory > Library Cache Tab
Metrics	**NOTE:** An overall high library cache hit ratio is a desirable thing. You should strive for a hit ratio between 95-100%, with 99% being a good performance benchmark for code reuse. When a database is first started, the library cache hit ratio, along with the individual region hit ratios, will not be at an optimal level because all code being used will be relatively new, and as such, must be parsed and placed into the shared pool. If, however, after a solid hour or two of steady database time, the library cache hit ratio has not increased to desirable levels, you should look into the possibility of increasing the shared_pool_size parameter. To keep important code objects from being aged out of the library cache, you can use the DBMS_SHARED_POOL package to pin frequently used code objects in memory so they will always be there when needed.

Pinned Objects

To keep important code objects from being aged out of the library cache, you can use the DBMS_SHARED_POOL package to pin frequently used code objects in memory so they will always be there when needed. The Pinned Objects section displays all objects in the library cache that have pinned. The table below lists the information available in this section:

Column	Description
Owner	The user account that owns the object.
Name	The name of the object.
Type	The type of the object: INDEX, TABLE, CLUSTER, VIEW, SET, SYNONYM, SEQUENCE, PROCEDURE, FUNCTION, PACKAGE, PACKAGE BODY, TRIGGER, CLASS, OBJECT, USER, DBLINK.
Shareable Memory	The amount of memory consumed by the object in the shared pool.
Loads	The number of times the object has been loaded into the cache. Note that this count also increases when an object has been invalidated.
Executions	The number of times the object has been executed by a session thread.
Locks	The number of users actively locking the object.
Pins	The number of user actively pinning the object.
Pinned?	Indicates (YES or NO) if this object has been pinned in memory with the DBMS_SHARED_POOL package.

Shared Pool Tab

The main components of the shared pool memory are the library and dictionary caches. Many of the caches in the shared pool increase or decrease in size automatically as necessary. The Shared Pool tab of the Memory Detail view presents the following sections:

Shared Space Pool Utilization - a graphical representation of the amount of used and free memory.

Shared Pool Fragmentation - displays the following statistics:

- **Bucket**: A work area defined by its optimal memory requirement.
- **Free Space:** The amount of free space in the bucket.

>

- **Free Chunks**: The number of free chunks that are available to be allocated from the shared pool.
- **Average Size:** The average size of a chunk.
- **Largest:** The size of the largest chunk.

Shared Pool Usage - The Shared Pool Usage section shows the following statistics:

- **Shared Pool Component Name**
- **Amount Used (bytes)**
- **Percent Used**

PGA Tab

The PGA tab displays statistical measures and techniques that let you manage a session's Program Global Area.

> **NOTE:** This information is available for Oracle9i and later.

The PGA tab includes the following sections:

PGA Overview	PGA Workarea Summary	PGA Session Detail

For more information, see Memory.

PGA Overview

Oracle's Program Global Area (PGA) is a private memory region containing data and control information for every connected process. Examples of PGA information are the runtime area of a cursor and memory sort space. If shared servers are used, then part of the runtime area can be located in the SGA.

Intensive queries cause a big portion of the PGA to be dedicated to work areas supporting memory intensive operators, such as the following:

- Sort-based operators (for example, ORDER BY, GROUP BY, ROLLUP, window functions)
- Hash-join
- Bitmap merge
- Bitmap create
- Write buffers used by bulk load operations

A sort operator uses a sort area to perform the in-memory sort of a set of rows. Similarly, a hash-join operator uses a hash area to build a hash table.

> **NOTE:** This information is available for Oracle9i and later.

Location	Oracle Performance Analyst statistics > Memory > PGA Tab

Metrics	The size of a PGA area can be controlled and tuned. Larger work areas can significantly improve performance of session work, with the trade off being higher memory consumption. The size of a session's PGA should be big enough so that it can accommodate the input data and auxiliary memory structures allocated by its requested SQL actions. This is known as the optimal size of a PGA work area. When the size of a PGA work area is smaller than optimal, the response time increases, because an extra pass is performed over part of the input data. This is known as the one-pass size of the PGA work area. Under the one-pass threshold, when the size of a work area is smaller than its input data size, multiple passes over the input data are needed. This could drastically increase the response time of the session. This is known as the multi-pass size of the work area. The goal is to have most work areas running with an optimal size (more than 90% or more for OLTP systems), while a smaller amount are running with a one-pass size (less than 10%). Multi-pass execution should be completely avoided. Prior to Oracle9i, the PGA was controlled using the SORT_AREA_SIZE, HASH_AREA_SIZE, BITMAP_MERGE_AREA_SIZE and CREATE_BITMAP_AREA_SIZE parameters. Setting these parameters is difficult, because the maximum PGA size is ideally selected based on data input size and total number of work areas active in the system. These two factors vary from one work area to another and from one point in time to another.

Correction

With Oracle9i, you can simplify and improve the way PGA memory is allocated, by enabling automatic PGA memory management. In this mode, Oracle dynamically adjusts the size of the portion of the PGA memory dedicated to work areas, based on an overall PGA memory target explicitly set by the DBA. To enable automatic PGA memory management, you have to set the initialization parameter PGA_AGGREGATE_TARGET.

The table below describes two metrics you should to watch:

Metric	Description
Over allocation count	Over-allocating PGA memory can happen if the value of PGA_AGGREGATE_TARGET is too small to accommodate the PGA other component in the previous equation plus the minimum memory required to execute the work area workload. When this happens, Oracle cannot honor the initialization parameter PGA_AGGREGATE_TARGET, and extra PGA memory needs to be allocated. If over-allocation occurs, you should increase the value of PGA_AGGREGATE_TARGET.
Cache hit percentage **ORACLE 9.2 AND LATER**	This metric is computed by Oracle to reflect the performance of the PGA memory component. A value of 100% means that all work areas executed by the system since instance start-up have used an optimal amount of PGA memory. When a work area cannot run optimally, one or more extra passes is performed over the input data. This reduces the cache hit percentage in proportion to the size of the input data and the number of extra passes performed.

PGA Workarea Summary

Oracle's Program Global Area (PGA) is a private memory region containing data and control information for every connected process. Examples of PGA information are the runtime area of a cursor and memory sort space. If shared servers are used, then part of the runtime area can be located in the SGA.

Intensive queries cause a big portion of the PGA to be dedicated to work areas supporting memory intensive operators, such as the following:

- Sort-based operators (for example, ORDER BY, GROUP BY, ROLLUP, window functions)
- Hash-join
- Bitmap merge
- Bitmap create
- Write buffers used by bulk load operations

>

A sort operator uses a sort area to perform the in-memory sort of a set of rows. Similarly, a hash-join operator uses a hash area to build a hash table.

NOTE: This information is available for Oracle9i and later.

Location	Oracle Performance Analyst statistics > Memory > PGA Tab
Metrics	The size of a PGA area can be controlled and tuned. Larger work areas can significantly improve performance of session work, with the trade off being higher memory consumption. The size of a session's PGA should be big enough so that it can accommodate the input data and auxiliary memory structures allocated by its requested SQL actions. This is known as the optimal size of a PGA work area. When the size of a PGA work area is smaller than optimal, the response time increases, because an extra pass is performed over part of the input data. This is known as the one-pass size of the PGA work area. Under the one-pass threshold, when the size of a work area is smaller than its input data size, multiple passes over the input data are needed. This could drastically increase the response time of the session. This is known as the multi-pass size of the work area. The goal is to have most work areas running with an optimal size (more than 90% or more for OLTP systems), while a smaller amount are running with a one-pass size (less than 10%). Multi-pass execution should be completely avoided. Prior to Oracle9i, the PGA was controlled using the SORT_AREA_SIZE, HASH_AREA_SIZE, BITMAP_MERGE_AREA_SIZE and CREATE_BITMAP_AREA_SIZE parameters. Setting these parameters is difficult, because the maximum PGA size is ideally selected based on data input size and total number of work areas active in the system. These two factors vary a lot from one work area to another and from one point in time to another.
Correction	With Oracle9i, you can simplify and improve the way PGA memory is allocated, by enabling automatic PGA memory management. In this mode, Oracle dynamically adjusts the size of the portion of the PGA memory dedicated to work areas, based on an overall PGA memory target explicitly set by the DBA. To enable automatic PGA memory management, you have to set the initialization parameter PGA_AGGREGATE_TARGET. If multi-pass percentages are consistently non-zero, you should increase the size of PGA_AGGREGATE_TARGET.

PGA Session Detail

Oracle's Program Global Area (PGA) is a private memory region containing data and control information for every connected process. Examples of PGA information are the runtime area of a cursor and memory sort space. If shared servers are used, then part of the runtime area can be located in the SGA.

Intensive queries cause a big portion of the PGA to be dedicated to work areas supporting memory intensive operators, such as the following:

- Sort-based operators (for example, ORDER BY, GROUP BY, ROLLUP, window functions)
- Hash-join
- Bitmap merge
- Bitmap create
- Write buffers used by bulk load operations

A sort operator uses a sort area to perform the in-memory sort of a set of rows. Similarly, a hash-join operator uses a hash area to build a hash table.

>

The PGA Session Detail section displays PGA usage for all connected sessions. The table below describes the information in this section:

Data	Description
SID	The unique identified given the session by Oracle.
Username	The user account a session is using.
O/S ID	The operating system ID of the process.
Logon Time	The time the session logged into Oracle.
Machine Name	The name of the machine running the process.
PGA Alloc Memory	The amount of allocated PGA memory for the process.
PGA Used Memory	The amount of freeable or usable memory for the process.
PGA Max Memory	The maximum amount of PGA memory used by the process.

NOTE: This information is available for Oracle9i and later.

Location	Oracle Performance Analyst statistics > Memory > PGA Tab

I/O

In addition to a Home page, the I/O category of Oracle Performance Analyst includes the following tabbed pages:

Access Pattern Detail Tab	Datafile I/O Tab	Job Queue Activity Tab
DBWR/LGWR Tab	Rollback I/O Tab	Sort Activity Tab
Table I/O Tab		

Home

The I/O Home page includes the following sections:

Key Systems Analysis pane	Bottleneck Analysis pane	I/O Access Patterns - I/O
SQL Analysis pane	Workload Analysis pane	

Key Systems Analysis pane

When complaints begin to surface about your database's performance, oftentimes the root cause can be traced to one or more issues with I/O. The thing to keep in mind when you begin to monitor the I/O of your database is that you are actually reviewing the success of your physical design model.

All the physical storage characteristics and placements, the table and index designs, and the speed with which it all works are on display when I/O is monitored. Because a database's main index of performance is measured by how fast I/O needs are satisfied, it is your responsibility to quickly interrogate Oracle to determine if a reported database slowdown is I/O related.

>

Physical I/O

Physical I/O consists of Oracle going to disk to gather or write data. The database writer (DBWR) and log writer (LGWR) processes typically perform all the I/O work in the database. You can also use other processes like the checkpoint and archive processes (CKPT and ARCH). Embarcadero Performance Analyst shows three key indicators of physical I/O. The table below describes these indicators:

Indicator	Description
Physical Reads	Physical Reads is the total number of physical reads performed on all datafiles since the last refresh.
Physical Writes	Physical Writes is the total number of times the DBWR process has performed writes to various database datafiles since the last refresh
Redo Writes	Oracle's redo logs are hotbeds of I/O activity in databases that store heavily modified data. Redo log files are used to perform database recovery in the event of a system crash. Redo logs are written to in a cyclical fashion - each log file is filled up before Oracle moves on to the next file. The redo writes statistic reflects the total number of redo writes by the LGWR process since the last refresh

Metrics

The table below describes metrics for Physical I/O statistics:

Statistic	Metrics
Physical Reads	Large numbers of physical reads could reflect a too small data/buffer cache. The Buffer Cache Hit Ratio is a better indicator of overall logical vs. physical I/O.
Physical Writes	Wait events related to I/O activity are good indicators of physical I/O problems. These events include `db file parallel write` and `db file single write`.

Correction

Doing the following can negate large numbers of continuous physical reads:

- Increasing the size of the data/buffer cache.
- Pinning often-referenced objects in memory by using the KEEP buffer pool (Oracle 8 and higher.)
- Placing heavily scanned objects in larger blocksize tablespaces (16-32KB). For Oracle9i or later.
- Tune SQL statements for better efficiency.

Logical I/O

Logical I/O refers to data access performed in memory. The database writer (DBWR) and log writer (LGWR) processes typically perform all the I/O work in the database. You can also use other processes like the checkpoint and archive processes (CKPT and ARCH). Embarcadero Performance Analyst shows three key indicators of logical I/O. The table below describes these indicators:

Indicator	Description
Logical Reads	Logical Reads is the total number of db block gets and consistent gets (data read from memory) since the last refresh.
Logical Changes	Logical Changes is the total number of changes that were made to all blocks in the SGA that were part of an update or delete operation. These changes generate redo log entries and are permanent if the transaction is committed. The number of logical changes is an approximate indication of total database work.

>

Indicator	Description
Consistent Reads	Consistent Reads is the total number of times a consistent read was requested for a database block. Such a read is performed from Oracle's rollback segments.

Metrics

Regarding raw logical I/O counts, no hard-core metrics exist. However, because physical I/O takes longer to complete than logical (memory) I/O, you should minimize physical read operations when possible. The Buffer Cache Hit Ratio is a better indicator of overall logical vs. physical I/O.

Correction

While logical I/O is still up to 1,400 times faster than physical disk access, it would be wise to investigate the top logical I/O process using Performance Analyst and see what SQL it is executing. If one process on the system is consuming between 25-50% of the overall amount, their SQL might require tuning.

Bottleneck Analysis pane

When an Oracle database is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting can mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. You can use Bottleneck Analysis to determine if perceived bottlenecks in a database are contributing to a performance problem.

Bottleneck analysis is a valid method of measuring performance because it helps you track where a database has been spending its time. If heavy table scan activity has been dragging a database's performance down, you can use bottleneck analysis to confirm the actual root cause. Once one or more wait events or other bottlenecks have been pinpointed as possible performance vampires, you can discover a fair amount of detail about which sessions and objects are causing the problem.

For wait event analysis to be properly carried out, it is imperative that the timed_statistics initialization parameter be set to TRUE. By default this parameter is set to FALSE, which disallows the collection of wait times for each wait event defined to the Oracle engine. For one to really understand the impact of wait events on database performance, you need to not only discover what the database is or has been waiting on, but the durations of the waits. Having both allows a complete picture to be formed regarding the magnitude of wait-initiated performance degradations. Almost all Oracle experts now agree that allowing the collection of timing statistics adds little if anything to database overhead, so setting timed_statistics to TRUE should not be a worry. The Performance Analyst I/O home page identifies the top database and log file waits that can cause decreased performance.

The Bottleneck Analysis section includes:

Hottest Database Files	Top Database File Waits	Top Log File Waits

Hottest Database Files

The Hottest Database Files section lists the most active database files as indicated by physical I/O activity.

Location	Oracle Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Metrics	Understanding the storage-level hot spots of a database is beneficial for a couple of reasons. First, you can get a feel for overworked physical disks by viewing I/O statistics at the tablespace and datafile levels. If a particular disk or set of disks is under too much strain, you can relocate the tablespaces to other less-used devices or create new tablespaces on different disks and move hot objects to them (assuming, of course, you have extra disks available). Second, if you have used standard DBA practice and placed indexes in their own tablespace, you can view the I/O statistics for that tablespace and see if the indexes are actually being used.

Correction	Here are some areas to consider when viewing Hottest Database Files: *Seeing much activity in the SYSTEM tablespace and datafiles can indicate recursive calls (space management, etc.). The use of locally-managed tablespaces (Oracle8i and later) can help with space management as it relates to data dictionary references.* *Temporary tablespaces (devoted to sort activity) showing higher volumes of physical I/O could indicate a problem with excessive disk sorts.* *You might want to quickly review all the physical I/O for each drive/file system and get a feel for the overworked disks on your server. If you have under utilized disk drives (with their own controllers), consider relocating some tablespaces that exhibit high I/O characteristics to those drives.*

Top Database File Waits

When viewing wait statistics, there are many levels of detail. The first level is the system view, which provides a global, cumulative snapshot of all the waits that have occurred on a system. The Top Database File Waits section filters all system wait events to present waits that are associated with database file activity. It ranks the top events by the count of total waits.

Location	Oracle Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Correction	Oracle documentation contains a listing and description of every current wait event defined in Oracle. DBAs unfamiliar with what each event represents should keep this listing close by as they examine wait-based event metrics. The most common database file wait is a db file scattered read, which indicates table scan activity. Small table scans are normally not a problem as Oracle can access data from small tables by quickly caching and scanning them. Scans of large tables can be confirmed by examining the I/O access pattern metrics on the Performance Analyst I/O home page. Information is also available on which sessions are causing the large table scans. Finally, one thing to watch is not only the actual waits, but also the wait time in seconds. If no wait time is observed with each wait type, then performance is likely not endangered by the event.

Top Log File Waits

When viewing wait statistics, there are many levels of detail. The first level is the system view, which provides a global, cumulative snapshot of all the waits that have occurred on a system. This section filters all system wait events to present waits that are associated with log file activity. It ranks the top events by the count of total waits.

Location	Oracle Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Correction	Oracle documentation contains a listing and description of every current wait event defined in Oracle. DBAs unfamiliar with what each event represents should keep this listing close by as they examine wait-based event metrics. The most common database file wait is a log file parallel write, which indicates how quickly Oracle can flush the log buffer. Excessive wait times associated with this event count indicate a bottleneck at the log buffer level. Finally, one thing to watch is not only the actual waits, but also the wait time in seconds. If no wait time is observed with each wait type, then performance is likely not endangered by the event.

SQL Analysis pane

Much of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured database down in terms of end user response times. SQL statements that use much memory can also cause a problem in a database.

>

Before you can identify problem SQL in your database, you have to ask the question, “What is bad SQL?” What criteria do you use when you begin the hunt for problem SQL in your critical systems? Even the seasoned experts disagree on what constitutes efficient and inefficient SQL; so there is no way to sufficiently answer this question to every Oracle professional's satisfaction. The SQL Analysis for I/O shows what SQL statements have consumed the largest percentages of physical and logical I/O, sort activity, and rows processed. The table below lists some general criteria you can use when evaluating the output from various database monitors or personal diagnostic scripts:

Criteria	Description
Overall Response (Elapsed) Time	The amount of time the query took to parse, execute, and fetch the data needed to satisfy the query. It should not include the network time needed to make the round trip from the requesting client workstation to the database server. This statistic is available in Oracle9i or later.
CPU Time	The amount of CPU time the query took to parse, execute, and fetch the data needed to satisfy the query.
Physical I/O	This is often used as the major statistic in terms of identifying good vs. bad SQL, this is a measure of how many disk reads the query caused to satisfy the user's request. While you certainly want to control disk I/O where possible, it is important that you not focus solely on physical I/O as the single benchmark of inefficient SQL. Make no mistake, disk access is slower than memory access and also consumes processing time making the physical to logical transition, but you need to look at the entire I/O picture of a SQL statement, which includes looking at a statements' logical I/O as well.
Logical I/O	This is a measure of how many memory reads the query took to satisfy the user's request. The goal of tuning I/O for a query should be to examine both logical and physical I/O, and use appropriate mechanisms to keep both to a minimum.
Repetition	This is a measure of how often the query has been executed. A problem in this area is not as easy to spot as the others unless you know your application well. A query that takes a fraction of a second to execute may still be causing a headache on your system if it has executed erroneously (for example, a query that executes in a runaway PL/SQL loop) over and over again.

There are other criteria that you can examine like sort activity or access plan statistics (that show items like Cartesian joins and the like), but more often than not, these measures are reflected in the criteria listed above. Fortunately, Oracle records all the above measures (some only in 9i), which makes tracking the SQL that has been submitted against an Oracle database much easier.

Metrics	When you begin to look for inefficient SQL in a database, there are two primary questions you want answered: *What HAS been the worst SQL that has historically been run in my database?* *What IS the worst SQL that is running right now in my database?* When troubleshooting a slow system, you should be on the lookout for any query that shows an execution count that is significantly larger than any other query on the system. It could be that the query is in an inefficient PL/SQL loop, or other problematic programming construct. Only by bringing the query to the attention of the application developers will you know if the query is being mishandled from a programming standpoint. There is the possibility that the SQL statement just is not tuned well. To determine that, you can go further into Performance Analyst's Top SQL view and, if you have Embarcadero SQL Tuner installed, you can port the SQL over to SQL Tuner to better optimize the statement.

I/O Access Patterns - I/O

There are many different routes that Oracle can take to access various objects contained in a database. The I/O Access Patterns section summarizes the methods Oracle is using to satisfy end-user requests for data.

Location	Oracle Performance Analyst statistics > I/O > Home > SQL Analysis pane

Metrics	Long table scans are typically an activity to avoid as they have the capability to cause needless physical and logical I/O as well as flood the buffer cache with seldom referenced blocks of data. You can discover which sessions have been causing the most large table scans. The table fetch continued row statistic is indicative of chained/migrated row I/O. Such activity is not desired because chained/migrated row access can cause twice the I/O needed to access a table. Oracle must pick do two or more I/Os to read a chained/migrated row in a table. You can discover which sessions have been accessing tables with chained/migrated rows.

Workload Analysis pane

It is not uncommon for one or two users to cause the majority of runtime problems that plague a database. The problem could be a runaway process, an untuned batch procedure, or other user-initiated operation. The Top I/O Hogs table shows the top sessions that have caused the most physical I/O on the system (both reads and writes).

Metrics	If you see one or two users who have caused more than 25-50% of the total memory usage, go further into the sessions to see the activities they are performing.

I/O Tablespace Tab

Physical I/O consists of Oracle going to disk to gather or write data. Logical I/O refers to data access performed in memory. The database writer (DBWR) and log writer (LGWR) processes typically perform all the I/O work in the database. Other processes like the checkpoint and archive processes (CKPT and ARCH) may also be used.

The Tablespace I/O tab displays details concerning the physical I/O activity at the tablespace level. The table below lists the information available on this tab:

Column	Description
Tablespace Name	The name of the tablespace.
Physical Reads	The cumulative number of physical reads.
Physical Writes	The cumulative number of physical writes.
Physical Block Reads	The cumulative number of physical block reads.
Physical Block Writes	The cumulative number of physical block writes.
Physical Read Time	The time spent reading from the tablespace (in hundredths of seconds).
Physical Write Time	The time spent writing to the tablespace (in hundredths of seconds).

>

Metrics	Generally, you want to see much more logical I/O activity than physical I/O, at least with respect to reads, although this in and of itself is a guarantee of good I/O performance. Seeing logical and physical reads keeping pace with one another is a sure sign that the Oracle SGA's buffer cache is sized too small or that needless, large table scans are occurring, which cause blocks of data to be continually read in and flushed out of the buffer cache. Other telltale signs of trouble brewing are large amounts of activity visible in user's TEMP tablespaces. The normal interpretation of such a thing is that a large number of disk sorts are taking place (perhaps because the Init.ora/spfile parameter SORT_AREA_SIZE may be set too small).

Datafile I/O Tab

Physical I/O consists of Oracle going to disk to gather or write data. Logical I/O refers to data access performed in memory. The database writer (DBWR) and log writer (LGWR) processes typically perform all the I/O work in the database against the physical datafile used to hold information The checkpoint and archive processes (CKPT and ARCH), also perform I/O work in the database.

The Datafile I/O tab displays details concerning the physical I/O activity at the datafile level. The table below lists the information available on this tab:

Column	Description
DataFile Name	The name of the datafile.
Tablespace Name	The name of the tablespace.
Physical Reads	The cumulative number of physical reads.
Physical Writes	The cumulative number of physical writes.
Block Reads	The cumulative number of physical block reads.
Block Writes	The cumulative number of physical block writes.
Read Time	The time spent reading from the tablespace (in hundredths of seconds).
Write Time	The time spent writing to the tablespace (in hundredths of seconds).

Metrics	Generally, you will want to see much more logical I/O activity than physical I/O, at least with respect to reads, although this in and of itself is a guarantee of good I/O performance. Seeing logical and physical reads keeping pace with one another is a sure sign that the Oracle SGA's buffer cache is sized too small or that needless, large table scans are occurring, which cause blocks of data to be continually read in and flushed out of the buffer cache. Other telltale signs of trouble brewing are large amounts of activity visible in user's TEMP tablespaces. The normal interpretation of such a thing is that a large number of disk sorts are taking place (perhaps because the Init.ora/spfile parameter SORT_AREA_SIZE may be set too small). This view also allows you to see the ‘hot’ files in terms of physical I/O activity. Seeing too much activity on one drive/file system may indicate a need for better file striping.

Rollback I/O Tab

The Rollback I/O tab includes the following sections:

Active Rollback Details	Rollback I/O	Session Rollback Activity

>

For more information, see Home.

Rollback I/O

To undo changes made to the Oracle database from within a transaction, Oracle writes data to individual rollback segments. Oracle also uses these segments to maintain read consistency for multiple users of data that is being modified. Because Oracle reads from and writes data to rollback segments, they can become very hot areas for I/O.

The Rollback I/O section presents everything necessary to view and troubleshoot rollback problems.The table below lists the information available in this section:

Column	Description
Name	The name of the rollback segment.
Size	The size of the rollback segment in KBs.
Shrinks	The number of times the rollback segment has decreased in size.
Extents	The number of times he rollback segment has increased in size.
Gets	The number of header gets (the segment has been used).
Waits	The number of header waits.
Writes	The number of bytes written to the rollback segment.
Active Trans	Indicates whether the rollback segment is active (non zero) or not (zero value).
Status	Indicates the status of the rollback segment, with the two main results being OFFLINE (a segment is offline and unavailable for transactions) and ONLINE (a rollback segment is online and available for transactional use).
High Water Mark	The largest size that the rollback segment has ever grown to.

Location	Oracle Performance Analyst statistics > I/O > Rollback I/O Tab
Metrics	To properly tune rollback I/O, you must first make sure that you have enough segments to accommodate the workload of the database. Constantly seeing a count of active rollback segments equal to or near the number of rollbacks defined for the database is an indicator that you should create more. An overall rollback contention ratio of 1% or greater is an indicator of too few rollbacks. Seeing wait counts greater than zero for each rollback segment is further evidence that you should create more rollback segments. Oracle9i provides the UNDO tablespace to automatically generate and eliminate the ‘correct’ number of rollback segments for a system given a certain workload. After ensuring that enough rollback segments exist in the database, you should then turn your attention to the question of sizing. Dynamic rollback extension can take a toll on performance when rollback segments are consistently enlarged to accommodate heavy transaction loads. Seeing rollback segments undergoing numerous extends and shrinks (as Oracle returns a segment back to its OPTIMAL setting), as well as rollback segments having current or high-water mark sizes greater than their OPTIMAL setting usually is a good indicator that they should be permanently enlarged. Again, Oracle9i’s automatic undo management can assist in this process.

>

Active Rollback Details

Rollback or undo segments are used to guarantee transactional integrity. When a transaction has not been committed or rolled back, a rollback segment will be in use with live transactional data. The Active Rollback Details section displays information concerning active rollback segments. The table below lists the information available in this section:

Column	Description
SID	The system ID of the session using the rollback segment.
Username	The account/schema name of the session.
Rollback	The rollback segment name containing the session's transaction.
Start Time	The time the transaction was initiated.
Blocks Used	The number of rollback blocks used by the transaction.

Location	Oracle Performance Analyst statistics > I/O > Rollback I/O Tab
Metrics	Seeing transactions with dated start times may indicate transactions that are the source of lock contention.

Session Rollback Activity

For all connected sessions, Oracle maintains statistics regarding each session's counts of rollbacks and commits. The Session Rollback Activity section displays a transactional summary/history that helps identify sessions with excessive rollback activity. The table below lists the information available in this section:

Column	Description
SID	The system ID of the session.
Username	The account/schema name of the session.
Machine Name	The machine name where the session logged on.
Logon Time	The time the session logged into Oracle.
User Rollbacks	The number of times the user manually issued the ROLLBACK statement or an error occurred during the user's transaction.
Transactions Rolled Back	The number of transactions successfully rolled back.
User Commits	The number of times a user transaction was committed.

Location	Oracle Performance Analyst statistics > I/O > Rollback I/O Tab
Metrics	Seeing high numbers of transactions rolled back might indicate an application error or other transactional problem.

DBWR/LGWR Tab

The DBWR/LGWR tab includes the following sections:

Daily Archive Log Summary (Last 7 Days)	Database Writer Detail	Log Writer Detail

Redo Wastage	Table I/O Tab	

Database Writer Detail

The database writer process (DBWR) handles the flow of information from Oracle's physical datafiles to and from the various memory structures in the system global area (SGA). On platforms that support it, you can configure and use multiple DBWR processes. The log writer (LGWR) process manages the information contained in Oracle's online redo log files and redo log buffer area. The table below lists the information available in this section:

Column	Description
Statistic	The specific metric for the database writer.
Amount	The value for the statistic.

Location	Oracle Performance Analyst statistics > I/O > DBWR/LGWR Tab
Metrics	Of all the statistics presented for the DBWR process, the summed dirty queue length statistic deserves attention. Non-zero values typically indicate buffers left in the write queue after a write request and may indicate that the DBWR process is falling behind.

Log Writer Detail

The database writer process (DBWR) handles the flow of information from Oracle's physical datafiles to and from the various memory structures in the system global area (SGA). On platforms that support it, you can configure and use multiple DBWR processes. The log writer (LGWR) process manages the information contained in Oracle's online redo log files and redo log buffer area. The table below lists the information available in this section:

Column	Description
Statistic	The specific metric for the log writer.
Amount	The value for the statistic.

Location	Oracle Performance Analyst statistics > I/O > DBWR/LGWR Tab
Metrics	For the LGWR process, non-zero values seen for the redo log space requests and redo log space wait time statistics could be a cause for concern. Redo log space requests reflect the number of times a user process waited for space in the redo log buffer, while the redo log space wait time presents the total time waited in milliseconds. Both could indicate the presence of contention in the redo log buffer. Possible remedies include increasing the log_buffer size in the SGA.

Daily Archive Log Summary (Last 7 Days)

To allow for point-in-time recovery, Oracle writes copies of redo log information to disk. When a database is running in archive log mode, a DBA can (with proper backup techniques in place) recovery nicely from a database error and roll forward to almost any point in time needed, as long as the proper archive logs are in place.

The I/O needed to write these archive logs is handled by Oracle's ARCH process. The Daily Archive Log Summary shows the number of archive logs written each day for the past seven days.

Location	Oracle Performance Analyst statistics > I/O > DBWR/LGWR Tab

Metrics	Heavy redo log activity can lead to many archive files being written out to disk. Batch jobs have the potential to move very fast - sometimes so fast that the online redo logs wrap back around before they have a chance to be archived. Messages of this nature will oftentimes show up in the Oracle alert log. If this is the case, you many want to think about increasing the size of the online redo log files, or increase the number of redo logs in general. Seeing archive files written more than one per hour or half-hour may indicate a too small redo size (or above average data modification load). If archive log protection is a concern (in other words, you do not want to lose an archive file that may be needed for recovery) and you are using Oracle8, you can now take advantage of the feature to write archive files to more than one destination on disk. It also allows multiple ARCH processes to now be invoked. Investigate the use of the Init.ora parameters log_archive_dest_n and log_archive_max_processes. Always remember one thing with respect to archive files and running Oracle in archive log mode: running out of archive file space on the server is a sure fire way to halt all activity in a database. Make sure you have plenty of free space available on your archive drives. Also, implement a purge procedure for older archives in conjunction with a good backup routine.

Redo Wastage

Oracle's redo logs are hotbeds of I/O activity in databases that store heavily modified data. Redo log files are used to perform database recovery in the event of a system crash. Redo logs are written to in a cyclical fashion - each log file is filled up before Oracle moves on to the next file. The Redo Wastage section shows how many bytes were wasted by redo log buffers being written before they were full.

Location	Oracle Performance Analyst statistics > I/O > DBWR/LGWR Tab
Metrics	Viewing a high percentage of wasted bytes to overall redo bytes can help you identify if redo wastage is a problem on your system.
Correction	Sometimes heavy redo wastage occurs when the log_checkpoint_interval parameter is set too high. If you think this is the case for your system, then *Edit the Init.ora file for the database.* *Change the amount of log_checkpoint_interval to a lower value.* *Cycle the Oracle server when possible to allow the new value to take effect.* *Monitor new value to see if performance improves.*

Table I/O Tab

I/O activity at the table level has historically been difficult to get with Oracle. Oracle9i and later contain better dictionary elements to help get a clearer picture of table I/O, but earlier versions of Oracle allow an estimation only. The table below lists the information available in the Estimated Physical I/O by table on this tab:

Column	Description
Owner Name	The owner of the table
Table Name	The name of the table or table partition.
Estimated Disk Reads	The estimated number of physical and physical direct reads for the table.

Metrics	Viewing high numbers of disk reads for a table could indicate excessive table scan activity. Such a table may benefit from better indexing or from being placed into a larger blocksize tablespace (Oracle9i and later).

Sort Activity Tab

The Sort Activity tab includes the following sections:

Current Disk Sorts	Session Sort Detail	

Session Sort Detail

Excessive sort activity can degrade a user's or overall database performance. When a sort operation occurs, Oracle attempts to perform the sort in a memory space, assigned by the DBA, which exists at the operating system level. If the sort is too large to be contained within this space, it will continue the sort on disk - specifically, in the user's assigned temporary tablespace.

The Session Sort Detail section displays historical sort statistics for every logged on session. The table below lists the information available in this section:

Column	Description
SID	The unique identified given the session by Oracle.
Username	The user account a session is using.
Machine Name	The machine name where the session logged on.
Logon Time	The time the session logged into Oracle.
Disk Sorts	The cumulative number of disk sorts for the session.
Memory Sorts	The cumulative number of memory sorts for the session.
Row Sorted	The cumulative number of rows sorted by the session.
Disk/Memory Sort Ratio	The percentage of times the session had to perform a disk sort.

Location	Oracle Performance Analyst statistics > I/O > Sort Activity Tab
Metrics	Techniques to include in your overall performance strategy are those that relate to minimizing the amount of sort activity overall and specifically sort activity that takes place on disk. A good place to start is by understanding things that cause sorts in the first place. A list of sort-related commands and SQL-related options include CREATE INDEX, ALTER INDEX … REBUILD, DISTINCT, ORDER BY, GROUP BY, UNION, INTERSECT, MINUS, IN, and NOT IN, as well as certain unindexed joins and certain correlated subqueries All of these SQL commands have the potential to create a sort. As a DBA, you probably will not know which queries will have their sorts performed entirely in memory and which ones will be forced to go to disk. There are times you simply cannot stop disk sort activity (such as in data warehousing environments). That being the case, you should ensure the following are true in your database: *Your users' TEMPORARY TABLESPACE assignment is not the SYSTEM tablespace, which is the default assignment.* *The TEMPORARY TABLESPACE assigned to your users is placed on a fast disk.* *The TEMPORARY TABLESPACE has the tablespace parameter TEMPORARY assigned to it, which allows sort activity to be performed in a more efficient manner.* If your overall memory sort ratio falls below 90%, you may want to increase the parameters devoted to memory sorts - sort_area_size and sort_area_retained_size. Keep in mind that Individual users might have the ability to alter their own sessions and increase their sort_area_size assignments! As a DBA, you may want to restrict users that have the ALTER SESSION privilege. In addition to increasing the amount of memory devoted to sorting, you should also hunt down inefficient SQL that cause needless sorts. For example, UNION ALL does not cause a sort whereas UNION does in an SQL query (to eliminate duplicate rows). DISTINCT oftentimes is coded inappropriately (especially by folks transferring from Microsoft Access, which used to use DISTINCT for nearly every SELECT query).

Current Disk Sorts

Excessive sort activity – especially disk sorts - can degrade a user's or overall database performance. When a sort operation occurs, Oracle attempts to perform the sort in a memory space, assigned by the DBA, which exists at the operating system level. If the sort is too large to be contained within this space, it will continue the sort on disk - specifically, in the user's assigned TEMPORARY TABLESPACE.

The Current Disk Sorts section displays sort statistics for every session currently performing a disk sort. The table below lists the information available in this section:

Column	Description
SQL Text	The SQL statement causing the disk sort.
SID	The unique identified given the session by Oracle.
Username	The user account a session is using.
Machine Name	The machine name where the session logged on.
Tablespace	The tablespace containing the sort segments used by the sort.
Extents	The number of extents being used for the sort.
Blocks	The number of blocks being used for the sort.

Location	Oracle Performance Analyst statistics > I/O > Sort Activity Tab
Metrics	Techniques to include in your overall performance strategy are those that relate to minimizing the amount of sort activity overall and specifically sort activity that takes place on disk. A good place to start is by understanding things that cause sorts in the first place. A list of sort-related commands and SQL-related options include CREATE INDEX, ALTER INDEX … REBUILD, DISTINCT, ORDER BY, GROUP BY, UNION, INTERSECT, MINUS, IN, and NOT IN, as well as certain unindexed joins and certain correlated subqueries All of these SQL commands have the potential to create a sort. As a DBA, you probably will not know which queries will have their sorts performed entirely in memory and which ones will be forced to go to disk. There are times you simply cannot stop disk sort activity (such as in data warehousing environments). That being the case, you should ensure the following are true in your database: *Your users' TEMPORARY TABLESPACE assignment is not the SYSTEM tablespace, which is the default assignment.* *The TEMPORARY TABLESPACE assigned to your users is placed on a fast disk.* *The TEMPORARY TABLESPACE has the tablespace parameter TEMPORARY assigned to it, which allows sort activity to be performed in a more efficient manner.* If your overall memory sort ratio falls below 90%, you may want to increase the parameters devoted to memory sorts - sort_area_size and sort_area_retained_size. Keep in mind that Individual users might have the ability to alter their own sessions and increase their sort_area_size assignments! As a DBA, you may want to restrict users that have the ALTER SESSION privilege. In addition to increasing the amount of memory devoted to sorting, you should also hunt down inefficient SQL that cause needless sorts. For example, UNION ALL does not cause a sort whereas UNION does in an SQL query (to eliminate duplicate rows). DISTINCT oftentimes is coded inappropriately (especially by folks transferring from Microsoft Access, which used to use DISTINCT for nearly every SELECT query).

Job Queue Activity Tab

The Job Queue Activity tab includes the following sections:

Job Queue Summary	Problem Jobs	Running Jobs Detail

>

Job Queue Summary

Oracle provides a way for Oracle developers to submit and run jobs through the job queue feature of the Oracle database. The table below lists the information available on this section:

Column	Description
Total Jobs	The number of jobs defined to the database.
Broken	The number of jobs broken (disabled).
Failures	The number of job failures.
Running	The number of jobs currently running.

Problem Jobs

Oracle provides a way for Oracle developers to submit and run jobs through the job queue feature of the Oracle database. Sometimes a job may fail or become disabled (broken). The table below lists the information available in this section:

Column	Description
Job ID	The job number for the job.
Login User	The user account used to run the job.
Last Date Successfully Run	The date/time of when the job last successfully ran.
Broken	Indicates if the job is disabled or not.
Failures	The number of failures the job has experienced.

Location	Oracle Performance Analyst statistics > I/O > Job Queue Activity Tab

Running Jobs Detail

Oracle provides a way for Oracle developers to submit and run jobs through the job queue feature of the Oracle database. The Running Jobs Detail section provides information about currently running jobs. The table below lists the information available in this section:

Column	Description
SID	The unique identified given the session by Oracle.
Job ID	The job number for the job.
Start Date/Time	Indicates when the job started.
Submitted by	Indicates the session that submitted the job.
Run As	Indicates which account is used to run the job.
Parse As	Indicates which account is used to parse the job.
Next Run Date/Time	Indicates when the job will next run.
Job Contents	Indicates the actual command used by the job.

>

Location	Oracle Performance Analyst statistics > I/O > Job Queue Activity Tab

Access Pattern Detail Tab

When a session submits a set of SQL or a transaction, there are many different combinations of methods the Oracle kernel will use to satisfy the query or transaction. The Access Pattern Detail tab displays all currently connected sessions, their demographic information, and a count of the various different methods that Oracle uses to handle incoming requests. The table below lists the information available on this tab:

Column	Description
SID	The unique identified given the session by Oracle.
Username	The user account a session is using.
Machine Name	The machine name where the session logged on.
Logon Time	The time the session logged into Oracle.
Large Table Scans	The cumulative count of ‘large’ table scans with ‘large’ normally being defined as any table over five blocks.
Small Table Scans	The cumulative count of ‘small’ table scans with ‘small’ normally being defined as any table under five blocks.
Table Scan ROWID Ranges	The cumulative count of table scans where the table in question has the CACHE property enabled.
Table Scan Cache Partitions	The cumulative count for partitions that have the CACHE property enabled.
Table Scan Direct Reads	The cumulative count of table scans performed with Oracle bypassing the buffer cache (data read only once - straight from disk and not from both disk and cache).
Chained Row Accesses	The cumulative count of the number of times a chained/migrated row was accessed.
Index by ROWID	The cumulative count of the number of rows fetched from a table by ROWID (normally index access).
Index Scans	The cumulative count of fast full index scans.
Index Scan ROWID Ranges	The cumulative count of fast full index scans with ROWID endpoints specified.
Index Scan Direct Read	The cumulative count of fast full index scans using direct read (data read only once - straight from disk and not from both disk and cache).

Space

In addition to a Home page, the Space category of Oracle Performance Analyst includes the following tabbed pages:

Extent Analysis Tab	Object Space Tab	Fragmentation Tab
Tablespace Detail Tab	Tablespace Growth Tab	User Space Tab

>

Home

The Space Home page includes the following sections:

Bottleneck Analysis pane	Redo Logs	Space Analysis pane
SQL Analysis pane	User Space Analysis pane	

Space Analysis pane

The Oracle tablespace is the logical container for the physical storage needs of a database. The Space Analysis section displays details for all the tablespaces for a particular database, including their total, used, free space, and percent free, as well as if each can automatically extend to support incoming space requests.

Metrics

If any tablespace's free space percent amount goes below 10%, and at least one tablespace's datafiles does not have AUTOEXTEND enabled (or the datafile has reached its extend limit), you should take action to ensure that the tablespace does not run out of available free space.

Correction

There are two things you can do to ensure that a tablespace does not run out of available free space:

1. First, you should look into the use of Oracle's AUTOEXTEND feature. AUTOEXTEND lets you give an Oracle tablespace the ability to auto-grow when it has exhausted the free space contained within. You can let a tablespace grow in an unlimited fashion or put constraints on it to stop at a certain point. You can also dictate how much more free space the tablespace gets each time it needs more space than is available. However, AUTOEXTEND enabled for a tablespace does not mean that you cannot run out of space. Remember you still have the physical server limitations to contend with. Make sure you (or your sysadmin) keep a careful eye on the server drives that house your Oracle database files for available free space.
2. If the free space on a server drive nears its limit, disable AUTOEXTEND for the datafile(s) that are on the target drive, and use the ALTER TABLESPACE … ADD DATAFILE command to place a new datafile for the tablespace on another drive that has more free space to offer.

    TIP: AUTOEXTEND is not a replacement for proactive space planning. When extra space is needed by the database, and AUTOEXTEND is activated by Oracle, performance slows as Oracle allocates more space for the tablespace. Avoiding AUTOEXTEND aids performance, albeit in a small way.

Bottleneck Analysis pane

When using bottleneck analysis to troubleshoot a database, you cannot rely only on the information contained in the wait event views that Oracle provides. For example, an object may attempt to extend into another extent of space in a tablespace and yet be denied if no such free space exists. Such a failure is not reflected in any wait event, but still represents a very real bottleneck to the database. In the same way that you cannot depend on only a few ratios to properly carry out ratio-based performance analysis, an administrator must include other statistical metrics in their overall bottleneck analysis framework to obtain an accurate performance risk assessment. Performance Analyst works to identify bottlenecks in your database that fall outside of pure wait events so you can get a total picture of all stoppages in your system. The Performance Analyst Space home page displays space-related bottlenecks that can decrease overall performance of your system.

The following bottleneck indicators are used on the Performance Analyst Space home page to succinctly communicate the general overall space situation of the monitored database:

Autoextend Datafiles at Extend Limit	Objects at Maximum Extents	Objects with No Room to Expand
Offline Tablespaces	Top Fragmented Datafiles	

Autoextend Datafiles at Extend Limit

You can specify that one or more tablespace datafiles can automatically grow by enabling the datafile's AUTOEXTEND property. you can also specify that the datafile not be permitted to grow beyond a certain limit, and once that limit is reached, the datafile and/or the tablespace may be permanently out of space.

Location	Oracle Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	You should not see any positive indicator that any datafile has reached its extend limit. While it might be okay from the standpoint that a datafile cannot simply grow anymore (if the tablespace has other datafiles that can grow or that no more space is added to the tablespace), disable AUTOEXTEND on the datafile.
Correction	If a datafile has reached its extend limit, you can take a number of actions: *Eliminate the extend limit for the datafile either by increasing the limit size or allowing to grow in an unlimited fashion.* *If the datafile is on a disk or file system that has no free space left, you can add another datafile to the tablespace to allow for space expansion and alter the problem datafile to have AUTOEXTEND disabled.* *If the datafile is on a disk or file system that has no free space left, the datafile can be physically relocated to another drive or file system with more free space.*

Objects at Maximum Extents

When an object in Oracle (table, index, table partition, etc.) needs to expand, Oracle automatically allocates another extent of space to accommodate the incoming data. However, a table or index can reach its maximum extent limit, which is the maximum number of extents for an object.

Location	Oracle Performance Analyst statistics > Space > Home > Bottleneck Analysis pane Oracle Performance Analyst statistics > Objects > Home > Bottleneck Analysis pane
Metrics	Unless an object is static in nature, and is not expected to grow in size, no object should ever get to its maximum extent limit.
Correction	If a maximum extent limit is reached for an object, you can take a number of actions: *The object can be reorganized into a locally-managed tablespace where no maximum extent limit is applied to any object.* *The object can be altered to have unlimited extents.* *The object can be reorganized into one or a few extents.*

The good news concerning locally-managed tablespaces is that unlimited extents are the rule, so if you choose to use these storage structures in your database, you will not have to worry about an object ever reaching a maximum extent limit. Some DBAs have expressed concerns over whether an object having hundreds or thousands of extents experiences performance problems when full table scans or similar operations are performed against them. While most Oracle experts agree that such I/O degradation might have been experienced in earlier Oracle versions, most feel that such is not the case any longer, especially when locally-managed tablespaces are used.

Objects with No Room to Expand

When an object in Oracle (table, index, table partition, etc.) needs to expand, Oracle automatically allocates another extent of space to accommodate the incoming data. However, an object may be prohibited from expanding for a number of reasons:

- The tablespace does not have enough free space to accommodate the extent request, and cannot automatically extend.

- The tablespace has enough total free space, but the space is not contiguous in nature (and the tablespace cannot automatically extend) and therefore the object cannot extend.

Location	Oracle Performance Analyst statistics > Space > Home > Bottleneck Analysis pane Oracle Performance Analyst statistics > Objects > Home > Bottleneck Analysis pane
Metrics	Unless an object is static in nature, and is not expected to grow in size, no object should ever be prohibited from acquiring its next extent of space.
Correction	If an object cannot extend, you can take a number of actions including: *One or more of the tablespace's datafiles can have its AUTOEXTEND property enabled, which lets the object to grow even though there is no space currently available in the tablespace.* *The object can be reorganized into a locally-managed tablespace all extents are uniform in nature.* *The object can be altered to have a reduced next extent size.*

Offline Tablespaces

A tablespace ONLINE status indicates that the tablespace is in normal read/write mode. Tablespaces marked read only displays a status of READ ONLY, while tablespaces engaged in backup mode indicates such in their status. An OFFLINE tablespace during normal processing hours is a cause for concern.

Location	Oracle Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	No tablespace should have a status of OFFLINE during normal business hours, unless you are performing maintenance on it.
Correction	If a tablespace has a status of OFFLINE, you can make it ONLINE using the ALTER TABLESPACE command.

Top Fragmented Datafiles

This section shows the tablespaces datafiles suffering from the highest levels of free space fragmentation in the database. Tablespaces are made up of object segments and space extents. Extents are either allocated to object segments or are free. When a tablespace is initially populated, all objects are neatly packed together in the front of the tablespace and all remaining free space is in one free chunk at the end. As objects grow (or extend) they are given new extents of space in the tablespace/datafile. If they are dropped, pockets of free space begin to appear throughout the tablespace. These pockets of space take one of two forms. The table below describes these forms:

Free Space	Description
Honeycombs	Pockets of free space that are adjacent to one another.
Bubbles	Pockets of free space that are trapped between object extents in the tablespace.

Location	Oracle Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	If you see a datafile that has many chunks of free space, determine if the tablespace is experiencing honeycomb or bubble fragmentation. You can handle honeycomb fragmentation quite easily, whereas bubble fragmentation is more difficult to solve.

Correction

You can eliminate honeycomb fragmentation with the ALTER TABLESPACE…COALESCE command. Issuing this command combines all pockets of adjacent free space into one extent. Each database maintenance plan should include a job that coalesces all the free honeycombs in a tablespace into one free chunk. Although Oracle is supposed to perform this operation automatically through the SMON process, it requires you to have the PCTINCREASE parameter of the tablespace set to a nonzero value. Having PCTINCREASE set to a value greater than zero encourages tablespace fragmentation through disparately sized extents. Plus, using SMON in this way is not efficient or entirely reliable.

You can temporarily solve bubble fragmentation by performing a total tablespace reorganization. A better long-term solution for Oracle databases 8.1.5 or later is to convert tablespaces over to locally managed tablespaces. 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. Problems caused by fragmentation then become a thing of the past.

To help stave off fragmentation problems:

- Set PCTINCREASE to zero for all tablespaces and objects to promote same-sized extents.
- Specify equal-sized allotments for your INITIAL and NEXT object storage parameters.
- Group objects with like growth and storage needs together in their own tablespaces.
- Considering converting any dictionary-managed tablespace to locally-managed.

You should also avoid fragmentation in the SYSTEM tablespaces. The best ways to do this include:

- 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.

SQL Analysis pane

Much of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured database down in terms of end user response times.

The SQL Analysis section of the Space home page shows a count of all SQL statements that have executed INSERT, UPDATE, DELETE, and TRUNCATE commands. Seeing these statistics should help you determine how much space-related SQL activity is occurring on your system.

Metrics	The statistics for this section can be used in many different ways. For example, large numbers of inserts could indicate a growing database and that you should keep a watchful eye on free space. Heavy update numbers could mean trouble for objects containing many indexes (indexes typically undergo splits and fragmentation when large volumes of updates are thrown their way). Many deletes might signal potential table high water mark problems.

Redo Logs

Oracle's redo logs are hotbeds of I/O activity in databases that store heavily modified data. Redo log files are used to perform database recovery in the event of a system crash. Redo logs are written to in a cyclical fashion - each log file is filled up before Oracle moves on to the next file. The redo bytes statistic reflects the total amount of redo generated in bytes since the last refresh. The redo wastage bytes indicate how many bytes were wasted by redo log buffers being written before they were full.

Location	Oracle Performance Analyst statistics > Space > Home > SQL Analysis pane
Metrics	Viewing a high percentage of wasted bytes to overall redo bytes can help you identify if redo wastage is a problem on your system.

Correction	Sometimes heavy redo wastage occurs when the log_checkpoint_interval parameter is set too high. To fix this: *Edit the Init.ora file for the database.* *Change the amount of log_checkpoint_interval to a lower value.* *Cycle the Oracle server when possible to allow the new value to take effect.* *Monitor new value to see if performance improves.* For Oracle9i and later, you can dynamically alter the log_checkpoint_interval parameter without stopping/starting the database.

User Space Analysis pane

The Top Space Hogs section details the schema in the database that own the most space. In databases where object owners can be many, it is oftentimes a good idea to take a quick look at which schema are leading the way in terms of space usage. This is not restricted to production environments only, but can be extended to dynamic development systems as well, especially if there are many developers in a database that have their own sets of objects. Developers sometimes have a bad habit of creating objects (or copies of objects) that they leave in a database even though they are not being used.

Correction	If a user owns many objects, or copies of objects, you can reclaim space in your database by contacting the user to see if those objects can be dropped.

Tablespace Detail Tab

The Tablespace Detail tab includes the following sections:

Datafile Detail	Tablespace Object Summary	

Tablespace Object Summary

Tablespaces are the logical containers for database objects inside Oracle. The Tablespace Object Summary section provides an overview of how much space is consumed in each tablespace by object type. The table below lists the information available in this section:

Column	Description
Tablespace Name	The name of the tablespace.
Status	Indicates the status of the tablespace.
Autoextend	Indicates whether one or more of the tablespace's datafiles has the ability to automatically grow.
Total Space (MB)	The total allocated space for the tablespace.
Table (MB)	The amount of used table and table partition space, in MB.
Index (MB)	The amount of used index and index partition space, in MB.
Rollback (MB)	The amount of used rollback and undo segment space, in MB.
Cluster (MB)	The amount of used cluster space, in MB.
LOB (MB)	The amount of used large object space, in MB.

Location	Oracle Performance Analyst statistics > Space > Tablespace Detail Tab

Datafile Detail

Tablespaces are the logical containers for database objects inside Oracle, and datafiles are the actual physical files that make up a tablespace. The Datafile Detail section provides an overview of the size of each datafile and other demographic information. The table below lists the information available in this section:

Column	Description
File Name	The name of the datafile.
Status	Indicates the status of the tablespace.
Autoextend	Indicates whether the datafile has the ability to automatically grow.
Tablespace Name	The name of the tablespace.
Size (MB)	The total allocated space for the datafile.
Max Size (MB)	The maximum size the datafile can automatically grow to.

Location	Oracle Performance Analyst statistics > Space > Tablespace Detail Tab

Object Space Tab

The Object Space tab displays a summary of space and extent information for all or selected objects in the database. The table below lists the information available on this tab:

Column	Description
Owner	The owner of the object.
Object Name	The name of the object.
Object Type	The type of object (table, index, etc.).
Space (KB)	The size of the object in KB.
Extents	The number of extents used by the object.
Max Extents	The maximum number of extents allowed for the object.
Initial Extent (KB)	The initial extent size for the object.
Next Extent (KB)	The next extent size for the object.
Tablespace	The tablespace where the object resides.

Metrics

One thing to keep an eye on are objects that are close to their maximum extent limit. Such objects can be altered to have more or unlimited extents if need be. The use of locally-managed tablespace eliminates the threat of an object ever encountering a maximum extent limit. For more information, see the Home page under Space.

Tablespace Growth Tab

Growing tablespaces do not have to spell problems for a DBA. Of course, up front planning is the key to sizing tablespaces correctly. Unfortunately, DBAs may not have all the information they need at tablespace creation time (or they may have the wrong information), so a tablespace can approach the end of its free space from time to time. The DBA can allow a tablespace to grow automatically (AUTOEXTEND) to prevent an out-of-space condition. Enabling AUTOEXTEND for a tablespace is quite reassuring for a DBA, but it introduces a new concept for the tablespace: monitoring data file growth. You should monitor your tablespaces that have AUTOEXTEND set so you can get an idea of the growth that is occurring in your database. Monitoring them lets you perform some mini-capacity planning and helps you get a jump-start on server upgrades.

The Tablespace Growth tab displays growth statistics for the database currently being monitored. The table below lists the information available on this tab:

Column	Description
Tablespace Name	The name of the tablespace.
Datafile Name	Thee name of the tablespace's datafile.
Autoextend	Indicates whether the tablespace has one or more datafiles that have the AUTOEXTEND feature set (allowing the datafile to automatically grow).
Created Size	The starting size for the physical datafile.
Current Size	The current size of the physical datafile.
Growth	The percentage of growth experienced for the datafile (if any).

Metrics

Using AUTOEXTEND can be very important for tablespaces devoted to disk sorting. Large data warehouses often must endure large disk sort operations. Having AUTOEXTEND enabled for tablespaces used for temporary segments (sort activity) helps large sort queries to complete when they might otherwise fail due to running out of sort space.

Just having AUTOEXTEND enabled for a tablespace does not mean you cannot run out of space. Remember, you still have to contend with the physical server limitations. Make sure you (or your sysadmin) keep a careful eye on the server drives that house your Oracle database files for available free space. If the free space on a server drive nears its limit, disable AUTOEXTEND for the datafile(s) that are on that drive, and use the traditional ALTER TABLESPACE…ADD DATAFILE command to place a new datafile for the tablespace on another drive that has more free space to offer.

Fragmentation Tab

Tablespaces are made up of object segments and space extents. Extents are either allocated to object segments or are free. When a tablespace is initially populated, all objects are neatly packed together in the front of the tablespace and all remaining free space are in one free chunk at the end. Unfortunately, this is not how things continue to be in a tablespace. As objects grow (or extend) they are given new extents of space in the tablespace. If you drop objects, pockets of free space begins to appear throughout the tablespace. These pockets of space are either honeycombs or bubbles. Honeycombs are not so difficult to deal with, but bubbles are another story.

The Fragmentation tab displays two different views of tablespace fragmentation. The first grid displays fragmentation at the tablespace level and the second grid displays fragmentation at the datafile level. The table below describes the information available in the Tablespace Fragmentation grid:

Column	Description
Tablespace Name	The tablespace name.
Free Space (MB)	The total amount of free space in MB for the tablespace.
Free Chunks	The total number of free chunks in the tablespace.

Column	Description
Largest Chunk (MB)	The largest free chunk (in MB) for the tablespace.

The table below describes the information available in the Datafile Fragmentation grid on the Fragmentation tab:

Column	Description
Datafile	The name of the datafile.
Autoextend	Indicates whether the datafile can automatically grow in size.
Tablespace	The tablespace name.
Free Chunks	The number of free chunks in the datafile.
Largest Chunk	The largest free chunk (in MB) for the datafile.

Metrics

To spot and correct fragmentation in your tablespaces, you should periodically monitor the fragmentation levels of your tablespaces at a global level. Doing so helps you quickly spot tablespaces that are experiencing fragmentation issues. Seeing a tablespace with only one chunk of free space is a sign that a tablespace is not having fragmentation problems. Seeing a tablespace with a couple of free chunks may not be a big deal either, because the tablespace could be made up of more than one datafile. Each datafile has its own chunk or chunks of free space.

If you see a tablespace that has many chunks of free space, the next thing to do is drill down into it and find out if the tablespace is experiencing honeycomb or bubble fragmentation. Honeycomb fragmentation occurs when pockets of free space exist that are adjacent to one another. Bubbles are pockets of free space that are trapped between object segment extents.

You can eliminate honeycomb fragmentation with the ALTER TABLESPACE...COALESCE command. This command combines all pockets of adjacent free space into one extent. It is important to do this because when an object is newly introduced to a tablespace (or an existing object needs to extend), and a contiguous chunk of free space does not exist to accommodate an object's INITIAL or NEXT size allocation, Oracle must manually coalesce all available honeycombs to try and form a large enough free chunk. This is a performance hit. If possible, you should try to minimize performance hits.

Bubble fragmentation is a more serious matter and is normally only corrected through tablespace or database reorganization. The standard technique is to perform an export of all objects in the tablespace, drop the tablespace, and then import all the objects back into the tablespace. However, this technique may just treat the symptom and not the cause of bubble fragmentation. The real issue is to address the reuse of space extents within a tablespace so that bubble fragmentation does not occur in the first place.

Oracle8 and later offers the concept of locally-managed tablespaces, which can all but eliminate tablespace fragmentation. It totally does away with the storage parameters of MINEXTENTS, MAXEXENTS, 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. Problems caused by fragmentation then become a thing of the past.

Correction

What can you do manually to help stave off fragmentation problems? A few suggestions include:

- Set PCTINCREASE to zero for all tablespaces and objects to promote same-sized extents.
- Specify equal-sized allotments for your INITIAL and NEXT object storage parameters.
- Group objects with like growth and storage needs together in their own tablespaces.

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:

1 Ensure no user has a DEFAULT or TEMPORARY tablespace assignment of SYSTEM.

2 Ensure no user has a quota set for SYSTEM.

3 Ensure no user has been granted the UNLIMITED TABLESPACE privilege.

Extent Analysis Tab

Object fragmentation results when objects consume multiple extents of space in a tablespace rather than a single block of space. Although performance problems with respect to object fragmentation are not as severe as they were in previous versions of Oracle, some DBAs believe that response-time penalties can still be chalked up to this situation. When multiple extents exist for an object, the amount of time it takes Oracle to scan it can be longer than if the object was made up of only one extent. This typically holds true when extents are scattered on different parts of the physical server disk. In addition, a performance hit is taken each time an object must extend into another extent of space.

The Extent Analysis tab displays objects whose extent count has exceeded a user-suggested numerical limit. The table below lists the information available on this tab if any objects are found:

Column	Description
Owner	The user account that owns the object.
Object Name	The name of the object.
Object Type	The type of object (TABLE, INDEX, etc.).
Next Extent	The next extent amount of space the object will consume.
Total Extents	The number of extents the object is currently consuming.
Max Extents	The MAXEXTENTS limit imposed on the object by the DBA.
Largest Free Chunk (KB)	The size of the largest chunk.
Tablespace Name	The name of the tablespace.

Metrics

When is object fragmentation a problem for an object? Opinions on this vary widely. As a general rule of thumb, if an object is exhibiting response time degradations, and no other factor can be attributed to the slowdown, examine the object to see how badly fragmented it is. If extent numbers are seen exceeding 100 or more, than you might want to think about reorganizing the object or using locally managed tablespaces. This rule of thumb does not apply to objects assigned to locally-managed tablespaces as extent fragmentation is expected and not a performance issue.

Correction

The best preventative for this problem is specifying the correct allocation of space for the object in the first place, but what can you do if you have objects with high numbers of extents? For tables and indexes you have three options:

1 Use Oracle's export/import utility to export, drop, and import the fragmented objects back into the database with the export parameter COMPRESS=Y. This will bring the objects back into the database with one large extent. Make sure that large enough chunks of available free space exist to accept the object back in, or you may experience space allocation errors.

2 With Oracle8, you can use the ALTER TABLE … MOVE command to reorganize a table back into one extent in a very safe and efficient manner.

3 Use ALTER INDEX … REBUILD to reorganize indexes that have moved into multiple extents.

Another thing to keep an eye on is the next extent size of an object vs. the largest free chunk of space in the tablespace. Seeing objects with larger next extents than free space may signal a problem unless the tablespace can automatically extend to accommodate the object's next extent.

User Space Tab

The User Space tab includes the following sections:

Top Space User Accounts	User Space Detail	

Top Space User Accounts

It is vital that a DBA know how the various accounts assigned to a database consume space. This is true in either production or development databases, because even in development databases, space can become an issue that causes major headaches. If development accounts are creating and cloning many objects without removing them, the result can be a lot of wasted space.

The Top Space User Accounts statistic shows how much total allocated space the top twenty accounts possess in the database.

Location	Oracle Performance Analyst statistics > Space > User Space Tab
Metrics	Seeing owner account with larger than expected data volumes could warrant further investigation. In addition, seeing accounts that have objects in more tablespaces than they have quota's for should also cause you to examine their account more closely to ensure the correct space privileges are in place.

User Space Detail

The User Space Detail section allows you to drill-down and see the specifics of the Top Space User accounts.

Location	Oracle Performance Analyst statistics > Space > User Space Tab

Objects

In addition to a Home page, the Objects category of Oracle Performance Analyst includes the following tabbed pages:

Hot Objects Tab	Invalid Objects Tab	Objects Summary Tab
Tables/Indexes in Same Tablespace Tab		

Home

The Objects Home page includes the following sections:

Bottleneck Analysis pane	Key Ratio Analysis pane	SQL Analysis pane
Object/Buffer Pool Placement pane	User Object Analysis pane	

Key Ratio Analysis pane

Object-related database activity can be examined using both ratio-based and wait/bottleneck-based analysis. Ratio-based analysis involves examining a number of key database ratios and statistical readings that can be used to indicate how active certain object types are. Performance ratios serve as roll-up mechanisms for busy DBAs to use for at-a-glance performance analysis.

When using ratio-based analysis, there are some standards to adhere to. To start with, many of the formulas that make up ratio-based analysis must be derived from delta measurements instead of cumulative statistics. Many of the global ratios that you examine come from the v$sysstat performance view. The performance view maintains a count of all the occurrences (in the VALUE column) of a particular database incident (in the NAME column) since the database was brought up. For databases that are kept up for long periods of time, these values can grow quite large and impacts how a particular ratio is interpreted. However, if delta statistics are used (taking, for a specified sampling period, the before and after counts of each statistic that make up a ratio), then an accurate and current portrayal of the various ratios can be had.

A final thing to remember about using ratio-based analysis is that, while there are a number of rules of thumb that can be used as starting points to begin the evaluation of database performance, DBAs must determine each database's individual 'personality' with respect to the key performance ratios. Some hard and fast rules simply do not apply to every database. The danger in using blanket ratio standards is that they can lead you to haphazardly take action, which can at times contribute nothing to the situation, and sometimes even degrade performance.

The following ratios are used on the Performance Analyst Objects home page to succinctly communicate the general session-based performance levels of the monitored database:

Active Rollback Ratio	Chained Row Fetch Ratio	

The following satistic, available on this pane, duplicates a statistic available on the Key Ratio Analysis pane of the Home page:

- Rollback Contention Ratio

Active Rollback Ratio

To undo changes made to the Oracle database from within a transaction, Oracle writes data to individual rollback segments. You can also use these to maintain read consistency for multiple users of modified data. Because Rollback Segments read and write data, they can become very hot areas for I/O. This statistic is a ratio that displays the percentage of rollback segments currently involved in work.

Seeing consistent ratios of 80% or higher could indicate the need for more rollback segments on the database.

Location	Oracle Performance Analyst statistics > Objects > Home > Key Ratio Analysis pane
Correction	If you are using Oracle8i or earlier, begin by creating new rollback segments and altering them to be online for use. Then monitor the overall active rollback ratio to see if it begins to drop. For Oracle9i or later, consider the use of UNDO management, which is where Oracle itself automatically manages rollback segments in special UNDO tablespaces.

Chained Row Fetch Ratio

In normal circumstances, a row of data should fit completely inside one Oracle block. Sometimes, however, this is not the case and the table suddenly contains chained or migrated rows (rows that span more than one data block).

Chaining occurs when a row is initially too large to fit inside one block. Two or more blocks are used by Oracle to hold the row. Migration deals with rows that have grown so much that they can no longer be contained within their original block. When this occurs, Oracle relocates the row out of its original block into another block, but leaves a pointer behind to indicate the relocation.

>

Both chaining and migration force Oracle to perform more than one I/O to retrieve data that could normally be obtained with a single I/O operation, with the end result being degraded performance.

The chain row fetch ratio statistic represents the percentage of rows fetched that were of either the chained or migrated row type.

Location	Oracle Performance Analyst statistics > Objects > Home > Key Ratio Analysis pane
Metrics	If the amount of chained rows in your tables exceeds 25-30%, you should take steps to eliminate the problem. Further, if the amount of chained rows accessed in your system vs. total rows accessed approaches 20-30%, you can start the process of eliminating the chained and migrated rows.

Correction

You can locate tables that contain chained rows. Once found, there are a couple of ways to reorganize tables to remove the chaining/migration problem. However, the best way to deal with chained and migrated rows is to prevent them from occurring. The table below describes two methods:

Method	Description
Use a large block size	Because chaining and migrations occur when the Oracle block is too small to hold the rows in question, make sure you are using a large enough block size for each database you create. An 8KB block size or higher is normally recommended to help stave off chained and migrated rows. If you are using Oracle9i or later, you can create special tablespaces that have larger block sizes (16-32KB) than the overall database block size and place any table that is a candidate for chained/migrated rows into them.
Use proper values of PCTFREE for tables	The necessary amount of percent free in a table helps prevent row migrations from occurring. If you have a database that houses rows with the potential to grow substantially over their initially inserted size, provide a liberal amount of PCTFREE for each of your tables.

Bottleneck Analysis pane

When using bottleneck analysis to troubleshoot a database, you cannot rely only on the information contained in the wait event views that Oracle provides. For example, an object attempts to extend into another extent of space in a tablespace and yet be denied if no such free space exists. Such a failure is not reflected in any wait event, but still represents a very real bottleneck to the database. In the same way that you cannot depend on only a few ratios to properly carry out ratio-based performance analysis, an administrator must include other statistical metrics in their overall bottleneck analysis framework to obtain an accurate performance risk assessment. Performance Analyst works to identify bottlenecks in your database that fall outside of pure wait events so you can get a total picture of all stoppages in your system. The Performance Analyst Objects home page displays object-related bottlenecks decreases the overall performance of your system.

The following bottleneck indicators are used on the Performance Analyst Objects home page to communicate the general overall conditions of objects in the monitored database:

Chained Tables	Session Blocks	Invalid/ Objects
Locked Objects	Table High Water Mark Problems	Tables/Indexes in Same Tablespace

The following statistics, available on this pane, duplicate statistics available on the .Top Session Bottlenecks of the Home page:

Enqueue Waits	Free List Waits	Objects at Maximum Extents
Objects with No Room to Expand		

Chained Tables

A row of data should fit completely inside one Oracle block. Sometimes, a table suddenly contains chained or migrated rows, which are rows that span more than one data block.

Chaining occurs when a row is initially too large to fit inside one block. Two or more blocks are used by Oracle to hold the row. Migration deals with rows that have grown so much that they can no longer be contained within their original block. Oracle relocates the row out of its original block into another block, but leaves a pointer behind to indicate the relocation.

Both chaining and migration force Oracle to perform more than one I/O to retrieve data that could normally be obtained with a single I/O operation, with the end result being degraded performance.

The Chained Tables statistic identifies the number of tables in the database that contain chained or migrated rows.

Location	Oracle Performance Analyst statistics > Objects > Home > Bottleneck Analysis pane
Metrics	Seeing many chained/migrated tables could indicate a database whose block size is undersized. However, even one chained/migrated row table can be a problem if that table is an often-referenced or scanned table.

Correction

You can locate the actual tables that contain chained/migrated rows. Once found, there are a couple of ways to reorganize tables to remove the chaining/migration problem. However, the best way to deal with chained and migrated rows is to prevent them from occurring. The table below describes two methods:

Method	Description
Use a large block size	Because chaining and migrations occur when the Oracle block is too small to hold the rows in question, make sure you are using a large enough block size for each database you create. An 8KB block size or higher is normally recommended to help stave off chained and migrated rows. If you are using Oracle9i or later, you can create special tablespaces that have larger block sizes (16-32KB) than the overall database block size and place any table that is a candidate for chained/migrated rows into them.
Use proper values of PCTFREE for tables	The necessary amount of percent free in a table helps prevent row migrations from occurring. If you have a database that houses rows with the potential to grow substantially over their initially inserted size, provide a liberal amount of PCTFREE for each of your tables.

For more information, see Bottleneck Analysis pane.

Session Blocks

A single blocking user has the potential to stop work for nearly all other processes on a small system, and can cause major headaches even on large systems. Although Oracle supports unlimited row-level locking, blocking lock situations do crop up. Blocks are most often caused by user processes holding exclusive locks and not releasing them via a proper COMMIT frequency.

NOTE: This statistic is also called Current Object Blocks on the Home page, and Sessions Blocked on the Users Home page.

Location	Oracle Performance Analyst statistics > Objects > Home > Bottleneck Analysis pane
Metrics	Investigate any indicator above zero immediately before the situation has a chance to mushroom.

Correction	Once discovered, a blocking lock situation can normally be quickly remedied. You can issue a KILL against the offending process, which eliminates the user's hold on the objects they were accessing. Other user processes then almost always complete in an instant. Discovering the blocked lock situation is eased by using tools like Performance Analyst, but preventing the blocking lock situation in the first place is more difficult. You can look at the Users detail and view all current blocking locks to see exactly which sessions are holding the currently restrictive locks. The culprit of blocking lock scenarios is usually the application design, or the SQL being used within the application itself. Properly coding an application to reference database objects in an efficient order, and then using the right SQL to get the job done, is an art. Most DBAs who have had to face Oracle Forms applications have suffered through the dreaded SELECT … FOR UPDATE statements that place unnecessary restrictive locks on nearly every read operation, and know all too well that good coding practice is important. The key to avoiding lock contention is to process user transactions in the quickest and most efficient manner possible - something not always easy to do.

Invalid/ Objects

Objects like procedures, packages, functions, triggers, and views can become invalidated for a variety of reasons. The main cause is generally a dependent object that has been altered or removed from the system. However, other objects, like indexes, can become invalid also due to scenarios like SQL*Loader problems. If an object that has become invalid is still referenced (through an application or SQL query tool), a variety of problems can result. Sometimes Oracle reports a clear error stating the problem, while other times seemingly odd behavior is exhibited by the database.

Location	Oracle Performance Analyst statistics > Objects > Home > Bottleneck Analysis pane
Metrics	There is no reason to have invalid objects in a production database. If your production databases have invalid objects that are no longer needed, promptly remove them from each system. Any needed objects that are indicating an invalid status should quickly be fixed before access problems develop. It is very normal for development databases to have invalid objects because developers create, modify, and compile objects all the time. The only invalid object that really should not be present in either a development or production database is an invalid index.
Correction	If code objects have become invalidated, you can issue an ALTER … COMPILE command to see if they compile properly and become valid once again. If they do not, then check the USER_ERRORS view for any resulting error messages. Indexes can be validated once more by using the ALTER INDEX … REBUILD command.

Locked Objects

Locked Objects is a count of all objects on the system that currently have some form of lock against them.

NOTE: Drilling down into the count of locked objects displays detail on each object that is locked, along with the user process holding the lock and the type of lock held.

Location	Oracle Performance Analyst statistics > Objects > Home > Bottleneck Analysis pane

Table High Water Mark Problems

One form of object fragmentation revolves around what is called a table's high water mark. Tables that are the victim of much insert-delete activity can develop performance problems due to wasted space being present. A table's high-water mark equals the last block in the table that was used to hold data. The problem in Oracle is that this high-water mark is not reset by DELETE activity, so it is possible for a table to have absolutely no data in it but contain a high-water mark that is many blocks high. When such a table is scanned, Oracle will read up to the high-water mark even if no rows exist in the table at all. This can make for some unnecessarily large scan times. For a real world example, see Bottleneck Analysis pane.

Location	Oracle Performance Analyst statistics > Objects > Home > Bottleneck Analysis pane
Metrics	Seeing non-zero numbers for this metric could indicate that table scans are taking longer than they should.
Correction	Performance Analyst does not currently provide diagnostics that identify the actual tables that contain artificially inflated high water marks. Embarcadero Space Analyst provides all the levels of detail regarding table high water marks plus a reorganization wizard that can reset the high water mark and thereby improve performance of table scans against tables suffering from this problem.

Tables/Indexes in Same Tablespace

One common Oracle practice is to physically separate tables and their child indexes onto distinct physical drives or file systems. Doing so reduces the potential for I/O contention when each object is accessed. The Tables/Indexes in Same Tablespace statistic displays a count of all tables whose index(es) reside in the same tablespace as they do.

Location	Oracle Performance Analyst statistics > Objects > Home > Bottleneck Analysis pane
Metrics	For development databases where performance is not an issue, seeing high counts of this metric is not an issue. Even on production systems, high counts of this metric may not be a problem, or cannot be avoided, if the server is configured with one large RAID5 disk configuration, each distinct drive or file system does not have its own controller, or there simply are not enough disks on the server to perform adequate object striping.
Correction	If the server is configured to allow for intelligent object striping, and you see a large count of tables and indexes that are housed in the same tablespace, you can: *Create a separate tablespace for indexes on a drive different than that used for tables, and perform an ALTER INDEX … REBUILD operation on all indexes to move them to the new tablespace. Certain versions of Oracle even allow the rebuild operation to occur online (using the ONLINE keyword of the REBUILD command).* *Use the export/import utility to move indexes to a new tablespace.* *Use the ALTER TABLE … MOVE command to move tables to a new tablespace in Oracle8i and later. Index-organized tables can be moved online.*

SQL Analysis pane

Certain objects in an Oracle database are accessed more than others. These objects can become a source of contention given certain conditions. The SQL Analysis section of the Objects home page identifies the most accessed tables on the database.

Metrics	DML activity against tables can cause contention for space management objects like free lists. Oracle9i or later provides automatic segment management, which can remove problems with free lists and the like.

Object/Buffer Pool Placement pane

Because data objects can be referenced with different types of usage patterns, Oracle8 offers the option to intelligently place objects into one of three buffer caches. The table below describes the buffer caches:

Buffer Cache	Description
KEEP	Designed to minimize misses in the buffer cache. Small objects that are frequently referenced are candidates for the KEEP buffer pool.
RECYCLE	Designed to avoid having large numbers of infrequently accessed blocks of data crowd out objects that need to be referenced in RAM, the RECYCLE pool is good for large objects that are scanned from beginning to end without the need for keeping all their data in RAM.
DEFAULT	The traditional cache for all other data objects.

NOTE: Unless you specify the KEEP or RECYCLE buffer cache, Oracle automatically places objects into the DEFAULT buffer cache.

The Object/Buffer Pool Placement section identifies how many objects are currently assigned to the various caches.

Location	Oracle Performance Analyst statistics > Objects > Home > SQL Analysis pane
Metrics	When looking at overall database I/O activity, you should keep an eye out for objects that you can place into a particular buffer cache. Consider using a KEEP cache for relatively small, frequently accessed tables that require fast response times. Large tables with random I/O activity and are scanned from beginning to end are good candidates for a RECYCLE cache.
Correction	Objects can be placed into different buffer pools at object creation time (using the STORAGE... BUFFER_POOL option in Oracle8 or later) or existing objects can be set into a different pool with the ALTER command. **NOTE:** Tables, partitions, and indexes can be placed into the different caches. If you just want to use the DEFAULT buffer pool and not enable any special caches, you can still encourage Oracle to keep certain objects in the cache as long as possible using the CACHE parameter. For example, issuing the command ALTER TABLE…CACHE specifies that the blocks retrieved for a table be placed at the most recently used end of the LRU list in the DEFAULT buffer cache when a full table scan is performed. You can also use the CACHE hint in SQL statements to cache a table, but used in this form, the blocks are only cached until the next time the database is shut down.

User Object Analysis pane

All objects are created inside of user accounts called schema. The User Object Analysis section of the Performance Analyst Objects home page shows schema that are ranked the worst in the several categories. The table below lists the categories:

Worst Schema for	Description
Chained Row Tables	The schema containing the most tables with chained rows is identified.
Invalid Objects	The schema containing the most invalid objects is identified.
Extent Problems	The schema containing the most objects with maximum extent limit or extent space deficit problems is identified.

Correction	To determine the degree and severity of each type of problem, go into detailed views that present specific information for each type of object problem.

Hot Objects Tab

The Hot Objects tab includes the following sections:

- Hot Code
- Hot Tables

Hot Tables

Certain objects in an Oracle database are accessed more than others. These objects can become a source of contention given certain conditions. The Hot Tables section identifies tables that are being frequently accessed through various SQL statements. The table below describes the information available in this section:

Column	Description
Table Owner	The owner of the table.
Table Name	The name of the table.
Issued Command	The SQL statement command issued against the table.
Executions	The number of SQL executions the object has experienced.
Disk Reads	The number of estimated disk reads from the object.
Buffer Gets	The number of estimated logical I/Os from the object.
Rows Processed	The number of estimated rows processed from the object.

Location	Oracle Performance Analyst statistics > Objects > Hot Objects Tab
Metrics	DML activity against tables can cause contention for space management objects like free lists. Oracle9i and later provides automatic segment management, which can remove problems with free lists and the like.

Hot Code

Certain objects in an Oracle database are accessed more than others. Data objects can become a source of contention given certain conditions, while code objects rarely cause contention issues. The Hot Code section identifies code objects (procedure, functions, etc.) that are being frequently accessed through various SQL statements. The table below describes the information available in this section:

Column	Description
Object Owner	The owner of the object.
Object Name	The name of the objects.
Object Type	The type of object (package, etc.).
Executions	The number of estimated executions for the object.
Loads	The number of times the object was loaded into the shared pool.
Locks	The number of locks the object has experienced.
Pins	The number of times the object was pinned in the shared pool.

Location	Oracle Performance Analyst statistics > Objects > Hot Objects Tab

Metrics	Often referenced code objects should be pinned in the shared pool using the Oracle DBMS_SHARED_POOL package. This can also be done through DBArtisan's pinned code utility. Objects with many executions and loads should be considered candidates for pinning.

Invalid Objects Tab

The Invalid Objects tab includes the following sections:

Invalid Objects Detail	Invalid Objects Summary	

Invalid Objects Summary

Objects like procedures, packages, functions, triggers, and views can become invalidated for a variety of reasons, with the main cause being a dependent object that has been altered or removed from the system. However, other objects, like indexes, can become invalid also due to scenarios like SQL*Loader problems. If an object that has become invalid is still referenced (through an application or SQL query tool), a variety of problems can result. Sometimes Oracle will report a clear error stating the problem, while other times seemingly quirky behavior will be exhibited by the database. In any event, as a DBA you should be on the lookout for objects in your database that have suddenly become invalid.

The Invalid Objects Summary section displays a count of the various objects in the database that are invalid.

Location	Oracle Performance Analyst statistics > Objects > Invalid Objects Tab
Metrics	There is no reason to have invalid objects in a production database. If your production databases have invalid/unusable objects that are no longer needed, you should promptly remove them from each system. Any needed objects that are indicating an invalid status should quickly be fixed before access problems develop. Correcting invalid objects like procedures and views often involves performing an ALTER…COMPILE operation. If the object status does not return to VALID, then further examination is warranted. It is very normal for development databases to have invalid objects because developers will no doubt be creating, modifying, and compiling objects all the time. The only invalid object that really should not be present in either a development or production database is an invalid index.

Invalid Objects Detail

Objects like procedures, packages, functions, triggers, and views can become invalidated for a variety of reasons, with the main cause being a dependent object that has been altered or removed from the system. However, other objects, like indexes, can become invalid also due to scenarios like SQL*Loader problems. If an object that has become invalid is still referenced (through an application or SQL query tool), a variety of problems can result. Sometimes Oracle will report a clear error stating the problem, while other times seemingly quirky behavior will be exhibited by the database. In any event, as a DBA you should be on the lookout for objects in your database that have suddenly become invalid.

The Invalid Objects Detail section displays information concerning invalid objects. The table below describes the information available in this section:

Column	Description
User Account	The user account that owns the objects.
Object Name	The name of the invalid object.
Object Type	The type of object (procedure, view, etc.).
Created On	The timestamp of when the objects was created.
Last DDL Date	The last structural modification date for the object.

Location	Oracle Performance Analyst statistics > Objects > Invalid Objects Tab
Metrics	There is no reason to have invalid objects in a production database. If your production databases have invalid objects that are no longer needed, you should promptly remove them from each system. Any needed objects that are indicating an invalid status should quickly be fixed before access problems develop. Correcting invalid objects like procedures and views often involves performing an ALTER…COMPILE operation. If the object status does not return to VALID, then further examination is warranted. It is very normal for development databases to have invalid objects because developers will no doubt be creating, modifying, and compiling objects all the time. The only invalid object that really should not be present in either a development or production database is an invalid index.

Objects Summary Tab

The Objects Summary tab includes the following sections:

Section	Description
Global Summary	This section displays a count of all user objects in the database (minus SYS and SYSTEM objects).
User Summary	This section displays a count across user accounts of all the major object types in the database. **NOTE** SYS and SYSTEM accounts are not included.

Tables/Indexes in Same Tablespace Tab

To help avoid physical I/O contention, it is wise to separate a table and its indexes onto separate physical drives. This is normally accomplished by putting tables into one tablespace (whose datafiles exist on one drive) and their indexes into a different tablespace (whose datafiles exist on another drive).

The Tables/Indexes in Same Tablespace tab identifies indexes whose tables reside in the same tablespace as they do. The table below describes the information available in this section:

Column	Description
Table Owner	The owner of the table.
Table Name	The name of the table.
Index Name	The name of the index that resides in the same tablespace as the table.
Tablespace Name	The name of the tablespace.

Metrics	If tables and indexes are found in the same tablespace, then either the table or index should be moved to a different tablespace if the server is set up in such a way that different drives and controllers are available that make the relocation worthwhile. Indexes are traditionally easier to move than tables. They can be relocated through many reorganization tools (like Embarcadero Space Analyst) or through the manual use of the ALTER INDEX … REBUILD command.

OS

In many scenarios, an optimally tuned database may not perform well because there are constraints imposed by the system where the database is running. These constraints may include processes competing with the database sever for resources (CPU, I/O, or Memory), a slow CPU, insufficient or slow I/O devices, and insufficient memory. The OS Statistics page of Performance Analyst lets you examine operating system metrics for the following platforms:

- AIX
- HP-UX

  NOTE: To view processor info and swap disk info on an HP-UX box, you need to login as ROOT in the OS login.

- Linux
- Solaris
- Unix
- Windows XP and 2000

  NOTE: The statistics available on the OS page depend on the operating system platform.

  TIP: If you magnify a graph to change back to the original size or minimize the size, close the OS Page and then reopen it again. Performance Analyst displays the graph in its original form and size.

In addition to a Home page, the OS category of Oracle Performance Analyst includes the following tabbed pages:

CPU Tab	I/O Tab	Memory Tab
Network Tab	Processes Tab	Location

Home

The OS Home page includes the following sections:

Bottleneck Analysis pane	Disk Analysis pane	Key Resource Usage pane
Memory Analysis pane	Workload Analysis pane	

Key Resource Usage pane

The following ratios are used on the Performance Analyst OS home page to communicate the general overall performance levels of the operating system:

Disk Time	Load Average	Processor Time
Paged Memory Used (Windows)	Swap Memory Used (AIX, HP-UX, Linux, Solaris, Unix)	

TIP: To open the CPU tab, right-click Processor Time. To open the I/O tab, right-click Disk Time. To open the Memory tab, right-click Paged Memory Used.

Disk Time

The Disk Time statistic is the percentage of elapsed time that the selected disk drive/device was busy servicing read or write requests.

Location	Oracle Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	You should avoid consistently seeing values for this statistic greater then 90%.

Correction	Add more disk drives and partition the files among all of the drives.

Load Average

The Load Average statistic represents the system load averages over the last 1, 5, and 15 minutes.

Location	Oracle Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	High load averages usually mean that the system is being used heavily and response time is correspondingly slow.

Processor Time

The Processor Time statistic indicates the percentage of time the processor is working. This counter is a primary indicator of processor activity.

Location	Oracle Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	If your computer seems to be running sluggishly, this statistic could be displaying a high percentage.
Correction	Upgrade to a processor with a larger L2 cache, a faster processor, or install an additional processor.

Paged Memory Used

The Paged Memory Used statistic is the ratio of Commit Memory Bytes to the Commit Limit. Committed memory is where memory space has been reserved in the paging file if it needs to be written to disk. The commit limit is determined by the size of the paging file. As the paging file increases, so does the commit limit.

NOTE: This statistic is available for the Windows platform.

Location	Oracle Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	This value displays the current percentage value only and not an average. If the percentage of paged memory used is above 90%, you may be running out of memory.
Correction	Increase the size of page file.

Swap Memory Used

The Swap Memory Used statistic is the percentage of swap space currently in use.

Location	Oracle Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	If the percentage of swap memory used is above 90%, you may be running out of memory.
Correction	Increase the size of your swap files.

Bottleneck Analysis pane

The following ratios are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Average Disk Queue Length	Interrupts/Sec	Network Output Queue Length/Network Queue Length
Processor Queue Length	Processor Queue Length	Processor Speed
Available Physical Memory		

TIP: To open the I/O tab, right-click any Details menu item. To open the CPU tab, right-click the Item Processor Queues Length Details menu item. To open the Network tab, right-click Network Output Queues Length.

NOTE: The statistics available in this section depend on the platform of operating system.

Average Disk Queue Length

The Average Disk Queue Length statistic is the average number of both read and write requests that were queued for the selected disk during the sample interval.

Location	Oracle Performance Analyst statistics > OS > Home > Bottleneck Analysis pane
Metrics	This metric is useful in identifying I/O related bottlenecks. If the disk queue lengths for certain disks are consistently much higher than others, you may need to redistribute the load among available disks. If the disk queues lengths for all disks are consistently large, and you see a high amount of I/O activity, your disks may be inefficient.
Correction	Some things you can do if you have problems with this statistic include: *Redistribute the data on the disk with the large average disk queue to other disks.* *Upgrade to faster disk(s).*

Interrupts/Sec

Interrupts/Sec is the average rate, in incidents per second, at which the processor received and serviced hardware interrupts. This value is an indirect indicator of the activity of devices that generate interrupts, such as the system clock, the mouse, disk drivers, data communication lines, network interface cards, and other peripheral devices. These devices normally interrupt the processor when they have completed a task or require attention. Normal thread execution is suspended. The system clock typically interrupts the processor every 10 milliseconds, creating a background of interrupt activity. This statistic shows the difference between the values observed in the last two samples, divided by the duration of the sample interval.

Location	Oracle Performance Analyst statistics > OS > Home > Bottleneck Analysis pane Oracle Performance Analyst statistics > OS > CPU Tab
Metrics	A high value indicates possible excessive hardware interrupts; justification is dependent on device activity.

Network Output Queue Length/Network Queue Length

The Network Output Queue Length statistic is the number of threads in the processor queue.

NOTE: The name of this statistic depends on the platform of the operating system.

Location	Oracle Performance Analyst statistics > OS > Home > Bottleneck Analysis pane

Metrics	Unlike the disk counters, this counter shows ready threads only, not threads that are running. There is a single queue for processor time even on computers with multiple processors. Therefore, if a computer has multiple processors, you need to divide this value by the number of processors servicing the workload. A sustained processor queue of less than 10 threads per processor is normally acceptable, dependent of the workload.
Correction	A sustained high value in the Processor Queue Length could indicate that a processor bottleneck has developed due to threads of a process requiring more process cycles than are available. If this is the case, you should look at installing a faster (or an additional) processor.

Page Faults/Sec

The Page Faults/Sec statistic is the overall rate faulted pages are handled by the processor. It is measured in numbers of pages faulted per second. A page fault occurs when a process requires code or data that is not in its working set. This counter includes both hard faults and soft faults.

Location	Oracle Performance Analyst statistics > OS > Home > Bottleneck Analysis pane Oracle Performance Analyst statistics > OS > Memory Tab
Metrics	This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
Correction	If the number of page faults remains consistently high, you can check with your Windows System Administrator for further investigation. Often, large numbers of page faults are not a problem so long as they are soft faults. However, hard faults, that require disk access, can cause delays.

Processor Queue Length

The Processor Queue Length statistic is the number of threads in the processor queue.

Location	Oracle Performance Analyst statistics > OS > Home > Bottleneck Analysis pane Oracle Performance Analyst statistics > OS > CPU Tab
Metrics	Unlike the disk counters, this counter shows ready threads only, not threads that are running. There is a single queue for processor time even on computers with multiple processors. Therefore, if a computer has multiple processors, you need to divide this value by the number of processors servicing the workload. A sustained processor queue of less than 10 threads per processor is normally acceptable, dependent of the workload.
Correction	A sustained high value in the Processor Queue Length could indicate that a processor bottleneck has developed due to threads of a process requiring more process cycles than are available. If this is the case, you should look at installing a faster (or an additional) processor.

Processor Speed

The Processor Speed statistic displays the speed of the active processor in MHz. The speed is approximate.

Location	Oracle Performance Analyst statistics > OS > Home > Bottleneck Analysis pane

Processor

The Processor Statistic displays the type of processor currently in use, for example, GenuineIntel.

Location	Oracle Performance Analyst statistics > OS > Home > Bottleneck Analysis pane

Memory Analysis pane

The following metrics are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Available Paged Memory (Windows)	Available Physical Memory	Available Swap Memory (AIX, HP-UX, Linux, Solaris, Unix)
Total Paged Memory/Total Swap Memory	Total Physical Memory	

TIP: To open the Memory tab, right-click any Details menu item.

NOTE: The statistics available in this section depend on the platform of operating system.

Available Paged Memory

The Available Paged Memory statistic shows the amount of virtual memory available for the processes.

NOTE: This statistic is available for the Windows platform.

Location	Oracle Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	If the available virtual memory is less than 10% of the total virtual memory, your system may run out of memory.
Correction	Increase the size of page file.

Available Physical Memory

The Available Physical Memory statistic represents the amount of RAM available to all processes.

Location	Oracle Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	This counter displays the last observed value only and not an average. Use this value with the Total physical memory and paging metrics (Memory details page). If the available physical memory is very small compared to this value, and the paging activity is high, your system may be running low on memory.
Correction	Some things you can do if you have problems with this statistic include: *Check the running processes to see if there are any memory leaks.* *Stop any services that are not required.* *Install additional RAM.*

Available Swap Memory

The Available Swap Memory statistic represents the amount of virtual memory available for the processes.

Location	Oracle Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	If the available Available Swap Memory is less than 10% of the total Swap Memory, your system may run out of memory.
Correction	Increase the size of swap files.

Total Paged Memory/Total Swap Memory

The Total Paged Memory statistic shows the maximum amount of virtual memory available to all processes.

NOTE: The name of this statistic depends on the platform of the operating system.

Location	Oracle Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	It is recommended that this be 1½ - 3 times the amount of RAM on the system.

Total Physical Memory

The Total Physical Memory statistic shows the amount of physical memory installed on your computer.

Location	Oracle Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	This is an informational metric and displays the total amount installed on the machine. Use this value with the available physical memory and paging metrics (Memory details page). If the available physical memory is very small compared to this value, and the paging activity is high, your system may be running low on memory.

Disk Analysis pane

The following ratios are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Free Disk Space	Total Disk Space	Used Disk Space

TIP: To open the Space tab, right-click any Details menu item.

NOTE: The statistics available in this section depend on the platform of operating system.

Free Disk Space

The Free Disk Space statistic shows the unallocated space, in megabytes on all logical disk drives.

Location	Oracle Performance Analyst statistics > OS > Home > Disk Analysis pane
Correction	There are many ways to ensure a database does not encounter a space problem due to physical space limitations: *If a database currently resides on a disk that has little free space, you can add more files to the database. Of course, you should add the new files to other physical hard disks that can accommodate a growing database.* *You should examine hard disks with shrinking disk space to see if you can relocate or delete files to allow more free space.*

Used Disk Space

The Used Disk Space statistic shows the amount of allocated space, in megabytes on all logical disk drives.

Location	Oracle Performance Analyst statistics > OS > Home > Disk Analysis pane

Correction	There are many ways to ensure a database does not encounter a space problem due to physical space limitations: *If a database currently resides on a disk that has little free space, you can add more files to the database. Of course, you should add the new files to other physical hard disks that can accommodate a growing database.* *You should examine hard disks with shrinking disk space to see if you can relocate or delete files to allow more free space.*

Total Disk Space

Total Disk Space displays the total allocated and unallocated space, in megabytes on all logical disk drives.

Location	Oracle Performance Analyst statistics > OS > Home > Disk Analysis pane
Correction	There are many things a DBA can do to ensure that a database does not encounter a space problem due to physical space limitations, here are two: *If a database currently resides on a disk that has little free space, you can add more files to the database. Of course, you should add the new files to other physical hard disks that can accommodate a growing database.* *You should examine hard disks with shrinking disk space to see if you can relocate or delete files to allow more free space.*

Workload Analysis pane

The following ratios are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Number of Logins	Number of Processes	Top CPU Process
Top I/O Process	Top Memory Process	

TIP: To open the Processes tab, right-click any Details menu item.

NOTE: The statistics available in this section depend on the platform of operating system.

Number of Logins

This statistic displays the total number of logins on the server.

Location	Oracle Performance Analyst statistics > OS > Home > Workload Analysis pane

Number of Processes

This statistic displays the total number of processes on the server.

Location	Oracle Performance Analyst statistics > OS > Home > Workload Analysis pane

>

Top Memory Process

Top Memory Process shows the current process that is consuming the most amount of memory. The information displayed is dependent on the platform of the operating system. Information displayed includes the name of the process, process ID, amount of memory consumed expressed in KB, amount of CPU expressed as a percentage, the amount of Major Page Faults, and the amount of I/O expressed in KB/sec.

Location	Oracle Performance Analyst statistics > OS > Home > Workload Analysis pane
Metrics	If you are running out of memory on the system, this is a quick way to identify the top memory user. If the displayed process is using a significant portion of the total memory, it could be causing the memory issues.

Top CPU Process

Top CPU Process shows the current process that is consuming the most amount of CPU. The information displayed is dependent on the platform of the operating system. Information displayed includes the name of the process, process ID, amount of memory consumed expressed in KB, amount of CPU expressed as a percentage, the amount of Major Page Faults, and the amount of I/O expressed in KB/sec.

Location	Oracle Performance Analyst statistics > OS > Home > Workload Analysis pane
Metrics	If the amount of CPU time used by this process is close to 100% and the CPU usage is very high, this process may be the bottleneck on the server.
Correction	Investigate the process further to see if it is in an inconsistent state. Also, look at minimum requirements for CPU speed for the process. You may need to upgrade your CPU.

Top I/O Process

The Top I/O Process statistic shows the current process that is consuming the most amount of CPU. The information displayed is dependent on the platform of the operating system. Information displayed includes the name of the process, process ID, amount of memory consumed expressed in KB, amount of CPU expressed as a percentage, the amount of Major Page Faults, and the amount of I/O expressed in KB/sec.

Location	Oracle Performance Analyst statistics > OS > Home > Workload Analysis pane

CPU Tab

The CPU tab of the OS Detail includes the following sections:

Context Switches/Sec	CPU Events	CPU Utilization
Load Averages		

The following statistics available on this tab, are duplicates of statistics available on the Bottleneck Analysis pane of the OS Home page:

Interrupts/Sec	Processor Queue Length	

NOTE: The sections available on this tab depend on the platform of operating system.

CPU Utilization

The CPU Utilization section includes the following information:

% Privileged Time	% User Time	

% Privileged Time

The % Privileged Time statistic is the percentage of elapsed time that the process threads spent executing code in privileged mode.

NOTE: For Windows systems, when a Windows system service is called, the service will often run in privileged mode to gain access to system-private data. Such data is protected from access by threads executing in user mode. Calls to the system can be explicit or implicit, such as page faults or interrupts. These kernel commands, are considered privileged to keep the low-level commands executing and prevent a system freeze. Unlike some early operating systems, Windows uses process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. Some work done by Windows on behalf of the application might appear in other subsystem processes in addition to the privileged time in the process.

Location	Oracle Performance Analyst statistics > OS > CPU Tab
Metrics	The ideal range should be 0-40% (less than 40% indicates excessive system activity).
Correction	If your CPU consistently runs at less than 40% you may need to upgrade your system to include a faster processor(s).

% User Time

The % User Time statistic is the percentage of elapsed time the processor spends in the user mode. User mode is a restricted processing mode designed for applications, environment subsystems, and integral subsystems. The alternative, privileged mode, is designed for operating system components and allows direct access to hardware and all memory. The operating system switches application threads to privileged mode to access operating system services. This counter displays the average busy time as a percentage of the sample time.

Location	Oracle Performance Analyst statistics > OS > CPU Tab
Metrics	If the Privileged Time is high in conjunction with Physical Disk Reads, consider upgrading the disk I/O subsystem.

CPU Events

The CPU Events section includes the following information:

Context Switches/Sec	Load Averages	System Calls/Sec

NOTE: The statistics available in this section depend on the platform of operating system.

Context Switches/Sec

The Context Switches/Sec section shows the combined rate at which all processors on the computer are switched from one thread to another. Context switches occur when a running thread voluntarily relinquishes the processor, is preempted by a higher priority ready thread, or switches between user-mode and privileged (kernel) mode to use an Executive or subsystem service.

Location	Oracle Performance Analyst statistics > OS > CPU Tab
Metrics	The ideal range should be between 0-10,000. GA number greater then 10,000 may indicate too many threads contending for resources.

System Calls/Sec

System calls are functions which require the services of the operating system. Usually they involve access to some data that users must not be permitted to corrupt or even change.

Location	Oracle Performance Analyst statistics > OS > CPU Tab
Correction	Try to find a good ratio between the amount of time the operating system runs (fulfilling system calls and doing process scheduling), and the amount of time the application runs. Your goal should be running 60% to 75% of the time in application mode, and 25% to 40% of the time in operating system mode. If you find that the system is spending 50% of its time in each mode, then you should investigate to determine what is wrong.

Load Averages

The Load Average statistic represents the system load averages over the last 1, 5, and 15 minutes.

Location	Oracle Performance Analyst statistics > OS > CPU Tab
Metrics	High load averages usually mean that the system is being used heavily and the response time is correspondingly slow.

Processes Tab

The Processes tab of the OS Detail page succinctly communicates the general overall performance levels of processes. The columns available in this table depend on the platform of operating system. The table below describes the information available in the table on this tab:

Column	Description
Process	The name of the process.
User	The user of the process.
ID	The ID Process is the unique identifier of this process. ID Process numbers are reused, so they only identify a process for the lifetime of that process.
CPU	The CPU is the percentage of elapsed time that all of process threads used the processor to execution instructions.
User Mode	The User Mode is the percentage of elapsed time that the process threads spent executing code in user mode.
Memory **WINDOWS ONLY**	Memory is the current size, in bytes, of the virtual address space the process is using. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is finite, and the process can limit its ability to load libraries.

Column	Description
Memory (MB)	Memory is the current size, in bytes, of the virtual address space the process is using. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is finite, and the process can limit its ability to load libraries.
Memory	Memory is the percentage of the memory used of the total memory.
Active Memory	Active Memory is the amount of committed virtual memory, in bytes for this process. Active memory is the physical memory which has space reserved on the disk paging file(s). There can be one or more paging files on each physical drive. This counter displays the last observed value only; it is not an average.
I/O Data	The rate at which the process is reading and writing bytes in I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/Os.
Elapsed Time	The total elapsed time, in seconds, that this process has been running.
Thread Count	The number of threads currently active in this process. An instruction is the basic unit of execution in a processor, and a thread is the object that executes instructions. Every running process has at least one thread.
Handle Count	The total number of handles currently open by this process. This number is equal to the sum of the handles currently open by each thread in this process.
Priority	The current base priority of this process. Threads within a process can raise and lower their own base priority relative to the process' base priority.
Creating Proc ID	The Creating Process ID value is the Process ID of the process that created the process. The creating process may have terminated, so this value may no longer identify a running process.
Page Faults/Sec	Page Faults/Sec is the rate at which page faults by the threads executing in this process are occurring. A page fault occurs when a thread refers to a virtual memory page that is not in its working set in main memory. This may not cause the page to be fetched from disk if it is on the standby list and hence already in main memory, or if it is in use by another process with whom the page is shared.
Page File	Page File is the current number of kilobytes that this process has used in the paging file(s). Paging files are used to store pages of memory used by the process that are not contained in other files. Paging files are shared by all processes, and the lack of space in paging files can prevent other processes from allocating memory.
Private	Private is the current size, in kilobytes, of memory that this process has allocated that cannot be shared with other processes.

I/O Tab

The sections available on this tab depend on the platform of operating system. The I/O tab includes the following sections:

Average Disk Queue/Total Disk Queue	Bytes per I/O Operation	Disk I/O Time
Disk Transfers/Sec	I/O Details	Transfer Rate

>

I/O Details

The table below describes the information available in this section:

Column	Description
Disk	The disk number assignment.
Reading (KB/s)	The amount of bytes read from the device.
Writing (KB/s)	The amount of bytes written to the device.
Disk Read Time	Disk Read Time is the percentage of elapsed time that the selected disk drive was busy servicing read requests.
Disk Write Time	Disk Write Time is the percentage of elapsed time that the selected disk drive was busy servicing write requests.
Disk Time	Disk Time is the percentage of elapsed time that the selected disk was busy servicing requests.
Avg. Read Queue	Avg. Disk Read Queue Length is the average number of read requests that were queued for the selected disk during the sample interval.
Avg. Write Queue	Avg. Disk Write Queue Length is the average number of write requests that were queued for the selected disk during the sample interval.
Disk Reads/Sec	Disk Reads/Sec is the rate of read operations on the disk.
Disk Writes/Sec	Disk Writes/Sec is the rate of write operations on the disk.

NOTE: The columns available in this table depend on the platform of operating system.

Location	Oracle Performance Analyst statistics > OS > I/O Tab

Transfer Rate

The Transfer Rate section succinctly communicates the general overall performance levels of the server's I/O. The table below describes the statistics in this section:

Statistic	Description	Metrics
Reading	The average number of bytes transferred from the disk during read operations.	It is measured in number of KBs per second.
Writing	The average number of bytes transferred to the disk during write operations.	It is measured in number of KBs per second.
Transferred	The amount of time required to transfer the data between the device's cache and the actual device media.	N/A

NOTE: The statistics available in this section depend on the platform of operating system.

Location	Oracle Performance Analyst statistics > OS > I/O Tab

>

Bytes per I/O Operation

The Bytes per I/O Operation section of the OS Detail succinctly communicates the general overall performance levels of the server's I/O. The table below describes the statistics in this section:

Statistic	Description	Metrics
Reading	The average number of bytes transferred from the disk during read operations.	It is measured in number of bytes per second.
Writing	The average number of bytes transferred to the disk during write operations.	It is measured in number of bytes per second.

Location	Oracle Performance Analyst statistics > OS > I/O Tab

Average Disk Queue/Total Disk Queue

The Average Disk Queue section of the OS Detail page succinctly communicates the general overall performance levels of the server's I/O. The table below describes the statistics in this section:

Statistic	Description	Information
Average Read Queue	Average Read Queue Length is the average number of read requests that were queued for a particular disk.	Sustained loads on any of the disk queue lengths could represent problems within your system.
Average Write Queue	Average Write Queue is the average number of write requests that were queued for a particular disk.	Sustained loads on any of the disk queue lengths could represent problems within your system. Contact your Windows System Administrator for assistance in troubleshooting these problems.
Average Queue	Average time in ms. to service I/O requests which for disks includes seek, rotational latency and data transfer time).	Sustained loads on any of the disk queue lengths could represent problems within your system. Contact your Unix System Administrator for assistance in troubleshooting these problems.

NOTE: The statistics available in this section depend on the platform of operating system.

Location	Oracle Performance Analyst statistics > OS > I/O Tab
Correction	Contact your Windows System Administrator for assistance in troubleshooting these problems.

>

Disk I/O Time

The Disk I/O Time section of the OS Detail page succinctly communicates the general overall performance levels of the disk's I/O. The table below describes the statistics in this section:

Statistic	Description	Metrics
% Disk Read Time	% Disk Read Time is the percentage of time during the sample interval that the disk is busy servicing read requests.	Use this in conjunction with the % Disk Write Time to identify the disks that may be running at capacity and may require upgrades. This metric may also be used to balance the load between under/over utilized disks.
% Disk Write Time	% Disk Write Time is the percentage of elapsed time that the selected disk drive was busy servicing write requests.	Use this in conjunction with the % Disk Read Time to identify the disks that may be running at capacity and may require upgrades. This metric may also be used to balance the load between under/over utilized disks.
% Disk Time	% Disk Time is the percentage of time during the sample interval that the disk is busy servicing requests.	Use this statistic to identify the disks that may be running at capacity and may require upgrades. This metric may also be used to balance the load between under/over utilized disks.

NOTE: The statistics available in this section depend on the platform of operating system.

Location	Oracle Performance Analyst statistics > OS > I/O Tab

Disk Transfers/Sec

The Disk Transfers/Sec section of the OS Detail page succinctly communicates the general overall performance levels of the disk transfers. The table below describes the statistics in this section

Statistic	Description
Reads/Sec	The the rate of read operations on the disk.
Writes/Sec	The rate of write operations on the disk.
Transfers/Sec	The rate of transfer operations per second.

NOTE: The statistics available in this section depend on the platform of operating system.

Location	Oracle Performance Analyst statistics > OS > I/O Tab

Memory Tab

The Memory tab of the OS Detail page includes the following sections:

Buffer Size	Cache Efficiency	Cache Size
Page Replacements	Paging Activity	Page Replacements

The following statistic, available on this tab, duplicates astatistic available on the Bottleneck Analysis pane of the OS Home page:

>

- Page Faults/Sec

NOTE: The sections available on this tab depend on the platform of operating system.

Buffer Size

The Buffer Size statistic represents the amount of cache used by the kernel in kilobytes.

Location	Oracle Performance Analyst statistics > OS > Memory Tab

Cache Efficiency

The Cache Efficiency section of the Memory tab succinctly communicates the general overall performance levels of the server's memory. The following statistics are available in this section:

Copy Read Hits %	Data Map Hits %	MDL Read Hits %
Pin Read Hits %		

Copy Read Hits %

The Copy Read Hits % statistic is the percentage of cache copy read requests that hit the cache and does not require a disk read to provide access to the page in the cache.

Location	Oracle Performance Analyst statistics > OS > Memory Tab
Metrics	When the page is pinned in the memory, the page's physical address in the file system cache will not be altered. A copy read is a file read operation where a page in the cache is copied to the application's buffer. Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.
Correction	Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate

Data Map Hits %

The Data Map Hits % statistic is the percentage of data maps in the file system cache that could be resolved without having to retrieve a page from the disk.

Location	Oracle Performance Analyst statistics > OS > Memory Tab
Metrics	Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.
Correction	Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate.

MDL Read Hits %

The MDL Read Hits % statistic is the percentage of Memory Descriptor List Read requests to the file system cache that hit the cache and does not require disk access to provide memory access to the pages in the cache.

Location	Oracle Performance Analyst statistics > OS > Memory Tab

Metrics	Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.
Correction	Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate.

Pin Read Hits %

The Pin Read Hits % statistic is the percentage of pin read requests that hit the file system cache and does not require a disk read in order to provide access to the page in the file system cache.

Location	Oracle Performance Analyst statistics > OS > Memory Tab
Metrics	Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.
Correction	Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate.

Cache Size

The Cache Size statistic represents the amount of cache used by the kernel in kilobytes.

Memory Available

The Memory Available section of the OS Detail page succinctly communicates the general overall performance levels of the server's memory. The statistics available in this section depend on the platform of operating system. The table below describes these statistics:

Statistic	Description	Metrics	Troubleshooting
Physical Memory	The Physical Memory statistic is the amount of committed virtual memory, in megabytes. Committed memory is where memory space has been reserved in the paging file if it needs to be written to disk.	This counter displays the last observed value only and not an average.	N/A
Physical	The Physical statistic is the total physical memory available.	N/A	N/A
Paged Memory	The Paged Memory statistic is the amount of committed virtual memory, in bytes. Used Memory is the physical memory which has space reserved on the disk paging file(s). There can be one or more paging files on each physical drive.	This counter displays the last observed value only; it is not an average.	N/A
Swap Memory	The Swap Memory statistic is the free, unreserved swap space.	N/A	Increase the size of your swap file or add additional files if you consistently see a shortage of free, unreserved swap space.
Swap	The Swap statistic is the total swap memory being used.	N/A	N/A
Total Physical	The Total Physical statistic is the total physical memory available.	N/A	N/A

Statistic	Description	Metrics	Troubleshooting
Total Swap	The Total Swap statistic is the total swap memory available.	N/A	N/A
Swap-ins	The number of memory pages paged in from swapspace.	N/A	If the system is consistently low on free memory (over a 30 second average), the memory scheduler will start to swap out processes. Increase the size of your swap file or add additional files.
Swap-outs	The number of memory pages paged in from swapspace.	N/A	If the system is consistently low on free memory (over a 30 second average), the memory scheduler will start to swap out processes. Increase the size of your swap file or add additional files.
Buffers	The Buffers statistic represents the amount of buffers used by the kernel in kilobytes.	N/A	N/A
Page Faults/Sec	The Page Faults/Sec statistic is the overall rate faulted pages are handled by the processor. It is measured in numbers of pages faulted per second. A page fault occurs when a process requires code or data that is not in its working set. This counter includes both hard faults and soft faults.	This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.	If the number of page faults remains consistently high, you can check with your Windows System Administrator for further investigation. Often, large numbers of page faults are not a problem so long as they are soft faults. However, hard faults, that require disk access, can cause delays.

Location	Oracle Performance Analyst statistics > OS > Memory Tab

Paging Activity

The Paging Activity section includes the following statistics:

Blocks Input/Sec	Blocks Output/Sec	Paged In
Paged Out	Pages Input/Sec	Pages Output/Sec

NOTE: The statistics available in this section depend on the platform of operating system.

Blocks Input/Sec

The Blocks Input/sec statistic is the number of blocks sent to a block device per second.

Location	Oracle Performance Analyst statistics > OS > Memory Tab

Correction	Although it never hurts to have as much physical memory as your system can handle, there are some things you can check within your system to alleviate the memory bottleneck. *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your page file. Normally, the bigger the initial size of your page file, the better, in performance terms.*

Blocks Output/Sec

The Blocks Output/sec statistic is the number of blocks received from a block device per second).

Location	Oracle Performance Analyst statistics > OS > Memory Tab
Correction	Although it never hurts to have as much physical memory as your system can handle, there are some things you can check within your system to alleviate the memory bottleneck. *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your page file. Normally, the bigger the initial size of your page file, the better, in performance terms.*

Paged In

The number of disk block paged into core from disk which occurs when the server is experiencing a shortage of RAM. Any nonzero value indicates excessive swap activity. In any virtual memory environment, it is not uncommon to see RAM pages moved out to a swap disk. This is a special disk area in UNIX that's reserved for holding memory pages so that the processor is capable of addressing RAM in excess of its full capability. While page-out of operations are a normal part of any server's operations, page-in operations indicate that the real amount of available RAM has been exceeded and that additional RAM pages are required on the server.

Location	Oracle Performance Analyst statistics > OS > Memory Tab
Correction	High values indicate the possible need to increase the amount of RAM on the system.

Paged Out

The number of disk block paged out of core from disk. In any virtual memory environment, it is not uncommon to see RAM pages moved out to a swap disk. This is a special disk area in UNIX that's reserved for holding memory pages so that the processor is capable of addressing RAM in excess of its full capability. While page-out of operations are a normal part of any server's operations, page-in operations indicate that the real amount of available RAM has been exceeded and that additional RAM pages are required on the server.

Location	Oracle Performance Analyst statistics > OS > Memory Tab

Pages Input/Sec

The Pages Input/Sec statistic is the number of pages read from disk to resolve hard page faults. Hard page faults occur when a process requires code or data that is not in its working set or elsewhere in physical memory, and must be retrieved from disk.

Location	Oracle Performance Analyst statistics > OS > Memory Tab

Metrics	This value was designed as a primary indicator of the kinds of faults that cause system-wide delays. It includes pages retrieved to satisfy faults in the file system cache (usually requested by applications) and in non-cached mapped memory files. This counter counts numbers of pages, and can be compared to other counts of pages, such as Memory: Page Faults/sec, without conversion. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
Correction	Although it never hurts to have as much physical memory as your system can handle, there are some things you can check within your system to alleviate the memory bottleneck. *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your page file. Normally, the bigger the initial size of your page file, the better, in performance terms.*

Pages Output/Sec

The Pages Output/Sec statistic is the number of pages written to disk to free up space in physical memory. Pages are written back to disk only if they are changed in physical memory. A high rate of pages output might indicate a memory shortage.

Location	Oracle Performance Analyst statistics > OS > Memory Tab
Metrics	Windows NT writes more pages back to disk to free up space when low in physical memory. This counter counts numbers of pages, and can be compared to other counts of pages, without conversion. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
Correction	Although it never hurts to have as much physical memory as your system can handle, there are some things you can check within your system to alleviate the memory bottleneck. *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your page file. Normally, the bigger the initial size of your page file, the better, in performance terms.*

Page Replacements

The following statistics are available in this section:

Memory Freed (Pages/sec)	Clock Algorithm Scans (Pages/sec)	

Memory Freed (Pages/sec)

Freed or destroyed (Kb/s).

Location	Oracle Performance Analyst statistics > OS > Memory Tab

Clock Algorithm Scans (Pages/sec)

The VMM uses a technique known as the clock algorithm to select pages to be replaced. This technique takes advantage of a referenced bit for each page as an indication of what pages have been recently used (referenced). When the page-stealer routine is called, it cycles through the PFT, examining each page's referenced bit.

Location	Oracle Performance Analyst statistics > OS > Memory Tab

>

Space Tab

The Space tab of the OS Detail page includes the following sections:

- Device Summary
- Device Detail

Disk Space Free

The Disk Space Free metric displays the amount of free disk space in megabytes.

Location	Oracle Performance Analyst statistics > OS > Space Tab

Disk Space Detail

The Disk Space Detail section of the Space tab succinctly communicates the general overall performance levels of the server's disks and space allotment. The table below describes the statistics in this section:

Statistic	Description
Partition	The drive letter of the disk.
Local Filesystem	The name of the file system.
Type	The type of file system.
Total Space	Total size of the disk/device's capacity expressed in MBs.
Used Space	Amount of MBs currently allocated on the particular disk/device.
Free Space	Amount of MBs currently unallocated and free on the particular disk/device.
Capacity	The percentage of space used on the device.
Mounted On	The mount point of the device.

Location	Oracle Performance Analyst statistics > OS > Space Tab

Device Summary

The Device Summary section of the Space tab displays a bar chart for all devices. The Device Summary section succinctly communicates the general overall performance levels of the server's disks and space allotment. The table below describes the statistics in this section:

Statistic	Description
Used	Amount of MBs currently allocated on the particular disk/device.
Free	Amount of MBs currently unallocated and free on the particular disk/device.

Location	Oracle Performance Analyst statistics > OS > Space Tab

>

Device Detail

The Device Detail section of the Space tab succinctly communicates the general overall performance levels of the server's disks and space allotment. The table below describes the statistics in this section:

Statistic	Description
Partition	The drive letter of the disk.
Local Filesystem	The name of the file system.
Type	The type of file system.
Total Space	Total size of the disk/device's capacity expressed in MBs.
Used Space	Amount of MBs currently allocated on the particular disk/device.
Free Space	Amount of MBs currently unallocated and free on the particular disk/device.
Capacity	The percentage of space used on the device.
Mounted On	The mount point of the device.

NOTE: The statistics available in this section depend on the platform of operating system.

Location	Oracle Performance Analyst statistics > OS > Space Tab

Network Tab

The Network tab of the OS Detail page succinctly communicates the general overall performance levels of the server's networking. The Network tab of the OS Detail page includes the following sections:

Network Details	Packet Collisions	Packet Discards
Packet Errors	Transmission Rate	Transmission Rate (Bytes)
Transmission Queue Length		

NOTE: The sections available on this tab depend on the platform of operating system.

Network Details

The statistics available in this section depend on the platform of operating system. The table below describes the information available in this section:

Column	Description
Network Interface	The name of network interface.
INET Address/Address	The IP address assigned to the network interface.
Pkts Sent/Sec	The number of packets sent per second.
Pkts Received/Sec	The number of packets received per second.
Sent (KB/Sec)	The number of bytes sent per second.
Received (KB/Sec)	The number of bytes received per second.
Out Pkts Discarded	The number of outbound packets discarded.
In Pkts Discarded	The number of inbound packets discarded.
Out Pkt Errors	The number of outbound packet errors.

>

Column	Description
In Pkt Errors	The number of inbound packet errors.
Queue Length	The queue length.
Collisions	The number of collisions.
Packets Discarded	The number of packets discarded.

Location	Oracle Performance Analyst statistics > OS > Network Tab

Transmission Rate

The Transmission Rate section of the Network tab succinctly communicates the packet transmission rate. The following statistics are available in this section:

Packets Received/Sec	Packets Sent/Sec	

Packets Sent/Sec

The Packets Sent/Sec statistic is the number of packets sent over each network adapter per second.

Location	Oracle Performance Analyst statistics > OS > Network Tab

Packets Received/Sec

The Packets Received/Sec statistic is the number of packets received over each network adapter per second.

Location	Oracle Performance Analyst statistics > OS > Network Tab

Transmission Rate (Bytes)

The Transmission Rate (Bytes) section of the Network tab succinctly communicates the packet transmission rate. The following statistics are available in this section:

Received (KB/Sec)	Sent (KB/Sec)	

Sent (KB/Sec)

The Sent (KB/Sec) statistic is the rate at which bytes are sent over each network adapter including framing characters.

Location	Oracle Performance Analyst statistics > OS > Network Tab

Received (KB/Sec)

The Received (KB/Sec) statistic is the rate at which bytes are received over each network adapter, including framing characters.

Location	Oracle Performance Analyst statistics > OS > Network Tab

>

Transmission Queue Length

The number of pending outgoing packets on either the software transmit queue or the hardware transmit queue

Location	Oracle Performance Analyst statistics > OS > Network Tab

Packet Collisions

A collision occurs when two devices on the network sense that the network is idle and end up trying to send data at the same time. Some collisions are normal on a repeated network, but excess collisions can cause serious performance problems.

Location	Oracle Performance Analyst statistics > OS > Network Tab
Correction	Contact your network administrator if there is high rate of persistent Packet Collisions.

Packet Discards

The Packet Discards section of the Network tab succinctly communicates the packet discard rate. The following statistics are available in this section:

Inbound Packets Discarded	Outbound Packets Discarded	Packet Discards

Outbound Packets Discarded

The Outbound Packets Discarded statistic is the number of outbound packets that were discarded even though no errors had been detected to prevent from being transmitted.

Location	Oracle Performance Analyst statistics > OS > Network Tab
Correction	A possible reason for discarding such a packet could be to free up buffer space.

Inbound Packets Discarded

Received Packets Discarded is the number of received packets that were discarded even though no errors had been detected to prevent from being transmitted.

Location	Oracle Performance Analyst statistics > OS > Network Tab

Correction

A possible reason for discarding such a packet could be to free up buffer space.

Packet Discards

The Packet Discards statistic represents the number of network packets discarded.

Location	Oracle Performance Analyst statistics > OS > Network Tab
Correction	Contact your network administrator if there is high rate of persistent packet discards.

>

Packet Errors

The Packet Errors section of the Network tab succinctly communicates the packet error rate. The following statistics are available in this section:

Inbound Packet Errors	Outbound Packet Errorss	

Outbound Packet Errors

The Outbound Packet Errors statistic is the outbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.

Location	Oracle Performance Analyst statistics > OS > Network Tab

Inbound Packet Errors

The Packets Received Errors statistic is the received packets that contained errors preventing them from being deliverable to a higher-layer protocol.

Location	Oracle Performance Analyst statistics > OS > Network Tab

Top SQL

One or two bad queries can cause a lot of trouble for the remaining sessions in a database, and so can one or two bad queries. It is important to find them before they get into a production system, but sometimes a few sneak through.

By applying custom filters and performance-related thresholds, the Top SQL view locates inefficient SQL. By applying filters to certain I/O and statistical counters, you hope to isolate queries that far exceed their nearest competitors in the same area (like disk reads). When you find them, you should reduce the number of sorts a query performs. Or, for a query that returns only a few records, you should try to minimize the number of rows a query processes.

When you begin to look for inefficient SQL in a database, there are two primary questions you need to answer:

1 What has been the worst SQL that's historically been run in my database?

2 What is the worst SQL that's running right now in my database?

When troubleshooting a slow system, you should be on the lookout for any query that shows an execution count that is significantly larger than any other query on the system. It may be that the query is in an inefficient Transact SQL loop, or other problematic programming construct. Only by bringing the query to the attention of the application developers will you know if the query is being mishandled from a programming standpoint.

The Top SQL view displays requested SQLOracle datasources. The following tabs are available on the Top SQL page:

Summary	Criteria	SQL Details
Recent History (Oracle 10g only)		

Summary

The Summary tab of the Top SQL page includes the following sections:

Top SQL Statements	Top SQL Details	

Top SQL Statements

A lot of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured database down in terms of end user response times.

Before you can identify problem SQL in your database, you have to ask the question of what 'bad SQL' is. What criteria do you use when you begin the hunt for problem SQL in your critical systems? Understand that even the seasoned experts disagree on what constitutes efficient and inefficient SQL; so there is no way to sufficiently answer this question to every Oracle professional's satisfaction. The table below describes some general criteria you can use when evaluating the output from various database monitors or personal diagnostic scripts:

Criteria	Description
Overall Response (Elapsed) Time	This is how much time the query took to parse, execute, and fetch the data needed to satisfy the query. It should not include the network time needed to make the round trip from the requesting client workstation to the database server. This statistic is available in Oracle9i and later.
CPU Time	This is how much CPU time the query took to parse, execute, and fetch the data needed to satisfy the query.
Physical I/O	Often used as the major statistic in terms of identifying good vs. bad SQL, this is a measure of how many disk reads the query caused to satisfy the user's request. While you certainly want to control disk I/O where possible, it is important that you not focus solely on physical I/O as the single benchmark of inefficient SQL. Make no mistake, disk access is slower than memory access and also consumes processing time making the physical to logical transition, but you need to look at the entire I/O picture of a SQL statement, which includes looking at a statements' logical I/O as well.
Logical I/O	This number of memory reads the query took to satisfy the user's request. The goal of tuning I/O for a query should be to examine both logical and physical I/O, and use appropriate mechanisms to keep both to a minimum.
Repetition	This is a measure of how often the query has been executed. A problem in this area is not as easy to spot as the others unless you know your application well. A query that takes a fraction of a second to execute may still be causing a headache on your system if it is executed erroneously (for example, a query that executes in a runaway PL/SQL loop) over and over.

Top SQL Details

A lot of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured database down in terms of end user response times.

Before you can identify problem SQL in your database, you have to ask the question of what 'bad SQL' is. What criteria do you use when you begin the hunt for problem SQL in your critical systems? Even the seasoned experts disagree on what constitutes efficient and inefficient SQL; so there is no way to sufficiently answer this question to every Oracle professional's satisfaction. The Top SQL Details section displays the top SQL statements across various performance categories. The table below describes the information available in this section:

Information	Description
Top For	Indicates the performance category for the SQL statement.
Usage	Displays the percent usage for this SQL statement in this performance category.
Hash value	A unique identifier for the SQL statement assigned by Oracle.
SQL text	A snippet of the actual SQL statement.
User name	The username of the account that originally parsed the SQL.
Disk reads	The number of disk reads for the statement.
Buffer gets	The number of logical reads for the statement.

>

Information	Description
Parse calls	The number of parse calls for the statement.
Sorts	The number of sorts caused by the statement.
Executions	The number of times that statement has been executed.
Rows processed	The number of rows processed by the statement.
First load time	The date/time of when the statement was first loaded into the shared pool.
Sharable memory	The amount of sharable memory used by the statement.
Persistent memory	The amount of persistent memory used by the statement.
Runtime memory	The amount of runtime memory used by the statement.
Address	The address of the SQL statement.

Location	Oracle Performance Analyst statistics > Top SQL > Summary

Criteria

The Top SQL view displays all SQL running or executed in the current datasource. You can use the lists and boxes to reorder the display according to your specific needs. The Criteria tab of the Top SQL page includes the following sections that you can mix and match to give you the information you want to find.

- Demographic Criteria
- Statistical Criteria
- Sort Criteria

SQL Details

The SQL Details tab displays the following sections:

Top SQL Results	Selected SQL Text	

Top SQL Results

A lot of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured database down in terms of end user response times.

Before you can identify problem SQL in your database, you have to ask the question of what 'bad SQL' is. What criteria do you use when you begin the hunt for problem SQL in your critical systems? Understand that even the seasoned experts disagree on what constitutes efficient and inefficient SQL; so there is no way to sufficiently answer this question to every Oracle professional's satisfaction. The Top SQL results grid displays the top SQL statements across various performance categories. The table below describes the information available in this section

Information	Description
SQL Text	A snippet of the full SQL statement. Clicking the statement will cause the full statement to be presented in the Selected SQL Text grid.
User name	The username of the account that originally parsed the SQL.

>

Information	Description
Disk reads per execution	The average number of disk reads caused by the statement each time it is executed.
Buffer gets per execution	The average number of logical reads caused by the statement each time it is executed.
Buffer gets	The number of logical reads for the statement.
Disk reads	The number of disk reads for the statement.
Parse calls	The number of parse calls for the statement.
Sorts	The number of sorts caused by the statement.
Executions	The number of times that statement has been executed.
Rows processed	The number of rows processed by the statement.
Hit Ratio	The cache hit ratio for the statement.
First load time	The date/time of when the statement was first loaded into the shared pool.
Sharable memory	The amount of sharable memory used by the statement.
Persistent memory	The amount of persistent memory used by the statement.
Runtime memory	The amount of runtime memory used by the statement.
Elapsed Time per Exec	For Oracle9i and later, this is the average elapsed response time for the statement.
CPU Time per Exe	For Oracle9i and later, this is the average CPU time for the statement.
CPU Time	For Oracle9i and later, this is the total CPU time for the statement.
Elapsed Time	For Oracle9i and later, this is the total elapsed time for the statement over all executions.
Hash value	A unique identifier for the SQL statement assigned by Oracle.
Address	The address of the SQL statement.

Location	Oracle Performance Analyst statistics > Top SQL > SQL Details

Selected SQL Text

The Selected SQL Text window allows you to select any of the Top SQL result rows and get the complete SQL text. From here you can open an Explain Plan or Tune the SQL.

Location	Oracle Performance Analyst statistics > Top SQL > SQL Details

Recent History (Oracle 10g only)

The Recent History tab of the Top SQL view displays:

SQL Activity (last 24 hours)	SQL Details	

SQL Activity (last 24 hours)

The SQL Activity graph shows you the resource usage of SQL statements over the last 24-hour period. You can see cumulative CPU, elapsed time, and wait time for all SQL statements that were executed during the last 24 hours.

>

You can also select a specific time period in the graph by drawing a box around part of the graph with your mouse. The graph will change to match your selection.

Location	Oracle Performance Analyst statistics > Top SQL > Recent History (Oracle 10g only)

SQL Details

The SQL Details section of the Recent History view displays the SQL statements that correspond to the time period shown in the SQL Activity graph. The following columns are displayed:

SQL ID: The unique identifier for the SQL statement

CPU Time: The total number of CPU seconds used by the statement during the specified time period

CPU %: The percentage of CPU used by the statement during the specified time period

Elapsed Time: The total elapsed time for the statement (in seconds) during the specified time period

Elapsed %: The percentage of elapsed time for the statement during the specified time period

Wait Time: The total wait time for the statement (in seconds) during the specified time period

Wait %: The percentage of wait time for the statement during the specified time period

SQL Time: The total SQL execution time for the statement (in seconds) during the specified time period

PL/SQL Time: The total PL/SQL execution for the statement (in seconds) during the specified time period

Java Time: The total java execution time for the statement (in seconds) during the specified time period

Executions: The total number of executions for the statement (in seconds) during the specified time period

Disk Reads: The total number of disk reads for the statement (in seconds) during the specified time period

Buffer Gets: The total number of buffer gets for the statement (in seconds) during the specified time period

Sorts: The total number of sorts for the statement (in seconds) during the specified time period

Parse Calls: The total number of parse calls for the statement (in seconds) during the specified time period

Location	Oracle Performance Analyst statistics > Top SQL > Recent History (Oracle 10g only)
Metrics	Potential SQL tuning targets will be those statements that have unusually large amounts of CPU time, elapsed time, wait time, or number of executions.

Users

In addition to a Home page, the Users category of Oracle Performance Analyst includes the following tabbed pages:

Locks/Blocking Locks Tab	Session Waits Tab	System Waits Tab
Top Sessions Tab		

Home

The Users home page includes the following sections:

Bottleneck Analysis pane	Key User Analysis pane	Bottleneck Analysis pane

>

Transaction Analysis pane	Workload Analysis pane	

Key User Analysis pane

User database activity can be examined using both ratio-based and wait/bottleneck-based analysis. Ratio-based analysis involves examining a number of key database ratios and statistical readings that can be used to indicate how efficiently users are accessing the database. Performance ratios serve as roll-up mechanisms for busy DBAs to use for at-a-glance performance analysis.

When using ratio-based analysis, there are some standards to adhere to. To start with, many of the formulas that make up ratio-based analysis must be derived from delta measurements instead of cumulative statistics. Many of the global ratios that you examine come from the v$sysstat performance view. The performance view maintains a count of all the occurrences (in the VALUE column) of a particular database incident (in the NAME column) since the database was brought up. For databases that are kept up for long periods of time, these values can grow quite large and impacts how a particular ratio is interpreted. However, if delta statistics are used (taking, for a specified sampling period, the before and after counts of each statistic that make up a ratio), then an accurate and current portrayal of the various ratios can be had.

A final thing to remember about using ratio-based analysis is that, while there are a number of rules of thumb that can be used as starting points to begin the evaluation of database performance, DBAs must determine each database's individual 'personality' with respect to the key performance ratios. Some hard and fast rules simply do not apply to every database. The danger in using blanket ratio standards is that they can lead you to haphazardly take action, which can at times contribute nothing to the situation, and sometimes even degrade performance.

The following ratios are used on the Performance Analyst Users home page to succinctly communicate the general session-based performance levels of the monitored database:

Active Sessions	CPU/Parse Ratio	Inactive Sessions
Open Cursors		

The following statistics, available on this pane are duplicates of statistics available on the Key Ratio Analysis pane of the Home page:

Memory Sort Ratio	Parse/Execute Ratio	

Active Sessions

The Active Sessions statistic is the total number of active and open threads currently reported in the database as well as the number of processes actively performing work.

NOTE: For related information, see Active User Processes.

For more information, see Inactive Sessions.

Location	Oracle Performance Analyst statistics > Users > Home > Key User Analysis pane

CPU/Parse Ratio

CPU resources are obviously required to accomplish parse activities. The CPU/Parse Ratio determines the percentage of CPU spent on user-initiated parse tasks.

Location	Oracle Performance Analyst statistics > Users > Home > Key User Analysis pane

Metrics	Low numbers (0-5%) are desired with the CPU/parse ratio. Larger values can indicate a problem with excessive parse activity.
Correction	If you suspect an excessive parsing problem on your database, you can look into increasing the shared pool size (in hopes that SQL statements are not aged out and therefore reused) or investigate the usage of bind variable in SQL statements. If an undersized shared pool is suspected on Oracle8i or earlier: *Edit the Init.ora file for the database.* *Increase the amount of shared_pool_size to a higher value (take caution to not over-allocate; ensure enough free memory exists on server before increasing value).* *Cycle the Oracle server when possible to allow the new value to take effect.* *Monitor the new value to see if performance improves.* If an undersized shared pool is suspected on Oracle9i or later: *Increase the size of the shared_pool_size parameter through use of the ALTER SYSTEM SET shared_pool_size command value (take caution to not over-allocate; ensure enough free memory exists on server before increasing value).* *Monitor the new value to see if performance improves.* *If using an SPFILE, save the new configuration values so Oracle reuses them each time the database is stopped and re-started.* If you determine that SQL literals are causing SQL to not be reused, you can (in Oracle 8.1.6 and later): *Change the cursor_sharing parameter to FORCE by using the ALTER SYSTEM SET cursor_sharing=FORCE command.* *Monitor database to see if parse activity is reduced and library cache reloads shrink* If using an SPFILE, save the new configuration values so Oracle reuses them each time the database is stopped and re-started. If using an Init.ora file, add the cursor_sharing=FORCE parameter to the file.

Inactive Sessions

The Inactive Sessions statistic is the total number of threads logged on to the database that are currently idle.

NOTE: For related information, see Inactive User Processes.

Location	Oracle Performance Analyst statistics > Users > Home > Key User Analysis pane
Metrics	A large number of inactive users could indicate user sessions that have mistakenly been left logged on. Because each user thread consumes a portion of memory on the Oracle server, to reduce resource usage, you should sever any session that does not need a connection.
Correction	Double-click this statistic to open the Top Sessions Tab of the Users Detail page. On this tab you can check sessions that have many seconds idle and/or that have been logged on for very long periods of time as indicated by the logon time column. After verifying that the session is no longer necessary, you can KILL the session. For more information, see Active Sessions.

Open Cursors

Open Cursors is the total number of all SQL open cursors that exist on the system. In some cases, Oracle cached cursors that have been open by PL/SQL procedures can be kept open for certain lengths of time, even though the actual activity has ceased.

Location	Oracle Performance Analyst statistics > Users > Home > Key User Analysis pane Oracle Performance Analyst statistics > Session Details > Current Work Tab

Metrics	You should monitor sessions to make sure that they do not approach the Open Cursor limit (specified in the Init.ora file). The parameter, open_cursors, limits how many open cursors (context areas) a session can have open at one time.
Correction	If the total number of open cursors approaches the open_cursors limit on Oracle8i or earlier, then: *Ensure that user processes are efficiently using cursors before editing the Init.ora file.* *Edit the Init.ora file for the database.* *Increase the amount of open_cursors to a higher value.* *Cycle the Oracle server when possible to allow the new value to take effect.* If the total number of open cursors approaches the open_cursors limit on Oracle9i or later then: *Change the open_cursors parameter to FORCE by using the ALTER SYSTEM SET open_cursors=< new value > command.*

Bottleneck Analysis pane

User database activity can be examined using both ratio-based and wait/bottleneck-based analysis. When connected to an Oracle database, every process is either busy doing work or waiting to perform work. A process that is waiting may mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. You can use the Bottleneck Analysis section to determine if perceived bottlenecks in a database are contributing to a performance problem.

Bottleneck analysis is a valid method of measuring performance because it helps you track where a database has been spending its time. If user latch contention or heavy table scan activity has been dragging a database's performance down, you can use bottleneck analysis to confirm the actual root cause. Once one or more wait events or other bottlenecks have been pinpointed as possible performance vampires, you can oftentimes discover a fair amount of detail about which sessions and objects are causing the problem.

For wait event analysis to be properly carried out, it is imperative that the timed_statistics initialization parameter be set to TRUE. By default this parameter is set to FALSE, which disallows the collection of wait times for each wait event defined to the Oracle engine. For one to really understand the impact of wait events on database performance, you need to not only discover what the database is or has been waiting on, but the durations of the waits. Having both allows a complete picture to be formed regarding the magnitude of wait-initiated performance degradations. Almost all Oracle experts now agree that allowing the collection of timing statistics adds little if anything to database overhead, so setting timed_statistics to TRUE should not be a worry. The Performance Analyst Users home page identifies the current top session waits as well as the top session-related bottlenecks that have seen the most waits on the system.

When using bottleneck analysis, you cannot rely only on the information contained in the wait event views that Oracle provides. Other user-related bottlenecks include current disk sort activity. Performance Analyst works to identify bottlenecks in your database that fall outside of pure wait events so you can get a total picture of all stoppages in your system.

The following bottleneck indicators are used on the Performance Analyst Users home page to succinctly communicate the general overall performance level of the monitored database:

Sessions Blocked	Sessions in Disk Sorts	Sessions with Enqueue Waits
Sessions Waiting	Top Bottleneck Events	.Top Session Bottlenecks

Sessions Blocked

A single blocking user has the potential to stop work for nearly all other processes on a small system, and can cause major headaches even on large systems. Although Oracle supports unlimited row-level locking, blocking lock situations do crop up. User processes holding exclusive locks and not releasing them via a proper COMMIT generally cause most blocks.

NOTE: This statistic is also called Current Object Blocks on the Home page and Session Blocks on the Objects Bottleneck Analysis pane.

Location	Oracle Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	You should immediately investigate any indicator above zero for a blocking lock statistic before the situation has a chance to grow out of control.
Correction	Once discovered, a blocking lock situation can normally be quickly remedied. You can issue a KILL against the offending process, which eliminates the user's stranglehold on the objects they were accessing. Other user processes then nearly almost always complete in an instant. Embarcadero Performance Analyst makes it easier to discover the blocked lock situation, but the trick is to prevent the blocking lock situation in the first place. You can look at the Users detail and view all current blocking locks to see exactly which sessions are holding the currently restrictive locks. The culprit of blocking lock scenarios is usually the application design, or the SQL being used within the application itself. Properly coding an application to reference database objects in an efficient order, and then using the right SQL to get the job done, is an art. Most DBAs who have had to face Oracle Forms applications have suffered through the dreaded SELECT … FOR UPDATE statements that place unnecessary restrictive locks on nearly every read operation, and know all too well that good coding practice is important. The key to avoiding lock contention is to process user transactions in the quickest and most efficient manner possible - something not always easy to do.

Sessions in Disk Sorts

Oracle's SGA is not the only memory structure used by Oracle for database work. One of the other memory areas used by Oracle8i and earlier for normal activity is an area set aside for sort actions. When a sort operation occurs, Oracle attempts to perform the sort in a memory space that exists at the operating system level. If the sort is too large to be contained within this space, it continues the sort on disk - specifically, in the user's assigned TEMPORARY TABLESPACE. Oracle records the overall number of sorts that are satisfied in memory as well as those that end up being finalized on disk. Using these numbers, you can calculate the percentage of memory sorts vs. disk sorts and get a feel for how fast your sort activity is being resolved. Obviously, memory sorts completes many times faster than sorts forced to use physical I/O to accomplish the task at hand.

Oracle9i or later now has the option of running automatic PGA memory management. Oracle has introduced a new Oracle parameter called pga_aggregate_target. When this parameter is set and you are using dedicated Oracle connections, Oracle ignores all of the PGA parameters in the Oracle file, including sort_area_size, hash_area_size and sort_area_retained_size. Oracle recommends that the value of pga_aggregate_target be set to the amount of remaining memory (less a 10% overhead for other server tasks) on a server after the instance has been started.

The sessions in disk sorts statistic displays a count of all sessions currently performing a disk sort.

Location	Oracle Performance Analyst statistics > Users > Home > Bottleneck Analysis pane

Metrics	Consistently seeing non-zero numbers for this statistic (as well as low values observed for the memory sort ratio) can indicate excessive disk sort activity. If you are on Oracle8i or earlier, increase the parameters devoted to memory sorts - sort_area_size and sort_area_retained_size. If you are using Oracle9i or later, investigate the use of pga_aggregate_target. Once the pga_aggregate_target has been set, Oracle automatically manages PGA memory allocation, based upon the individual needs of each Oracle connection. Oracle9i or later allows the pga_aggregate_target parameter to be modified at the instance level with the alter system command, thereby lets you dynamically adjust the total RAM region available to Oracle9i. Oracle9i also introduces a new parameter called workarea_size_policy. When this parameter is set to automatic, all Oracle connections benefit from the shared PGA memory. When workarea_size_policy is set to manual, connections allocate memory according to the values for the sort_area_size parameter. Under the automatic mode, Oracle tries to maximize the number of work areas that are using optimal memory and uses one-pass memory for the others.
Correction	If you find a problem, do the following: *Edit the Init.ora or SPFILE file for the database.* *Increase the amount of sort_area_size to a higher value (take care not to not over-allocate; ensure enough free memory exists on server before increasing value). Realize that EVERY user receives this amount for sorting).* *Cycle the Oracle server when possible to allow the new value to take effect.* *Monitor new value to see if performance improves.* In addition to increasing the amount of memory devoted to sorting, you should also locate inefficient SQL that causes needless sorts. For example, in an SQL query (to eliminate duplicate rows) UNION ALL does not cause a sort whereas UNION does. People frequently code DISTINCT inappropriately (especially people transferring from Microsoft Access, which uses DISTINCT for most SELECT queries). There are times you simply cannot stop sort activity. When this happens, you should try to keep it in memory whenever possible. However, large data warehousing systems frequently exhaust RAM sort allotments, so if disk sorts must occur, ensure three things: *Your user's TEMPORARY TABLESPACE assignment is not the SYSTEM tablespace, which is the default assignment. In Oracle9i or later, you can specify a default tablespace other than SYSTEM for every user account that is created.* *The TEMPORARY TABLESPACE assigned to your users is placed on a fast disk.* The TEMPORARY TABLESPACE has the tablespace parameter TEMPORARY assigned to it, which allows sort activity to be performed in a more efficient manner.

Sessions with Enqueue Waits

An enqueue is an advanced locking device that allows multiple database processes to share certain resources. Enqueue waits typically occur when sessions wait to be granted a requested lock. Sometimes these locks are internal Oracle locks while other times they could be locks for rows of data. Note that enqueues are issued implicitly by Oracle.

Location	Oracle Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	Any enqueue waits that read consistently above 1 or more (delta statistics) should be investigated.
Correction	Removing contention for enqueues is almost always an application design issue. If many enqueue waits are seen, this normally indicates either contention for certain rows in the database, or the result of database-initiated lock escalation. You should examine the use of indexes to make sure all referencing foreign keys are indexes and that SQL is tuned to not tarry over rows in the database during modification operations. Enqueue waits can also be the result of space management tasks (such as objects extending) and disk sorts (mainly in tablespaces that do not make use of the TEMPORARY tablespace parameter).

Sessions Waiting

User connections that are waiting on a system generally occur for two reasons:

1 A process waits because a requested resource is not available.

2 A process waits for Oracle to perform a prerequisite task for its given operation.

Idle waits (processes waiting because they have no work) are not usually a concern. However the two wait causes mentioned above are the ones worth your time and investigation. The sessions waiting statistic is a count of all sessions that are currently waiting for one reason or another.

Location	Oracle Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	To determine the actual wait causes currently experienced by user connections, you should drill down from the global count of users waiting, into the actual system and user wait details of a database.
Correction	If you find a problem, drill down into wait details to determine whether the waits are resource-related.

Top Bottleneck Events

When viewing wait statistics, there are many levels of detail. The first level is the system view, which provides a global, cumulative snapshot of all the waits that have occurred on a system. The second level is the session view, which shows details on which events connected sessions have experienced. Viewing these numbers can help you determine which session-related wait events have caused the most commotion in a database thus far. The Top Bottleneck Events section identifies the top waits that have occurred on the Oracle database based on the amount of time waited per event.

Location	Oracle Performance Analyst statistics > Users > Home > Bottleneck Analysis pane

Correction

Appendix A in the Oracle Reference manual contains a listing and description of every current wait event defined in Oracle. DBAs unfamiliar with what each event represents should keep this listing close by as they examine wait-based event metrics. For example, a ‘db file scattered read’ event is typically indicative of table scan operations. If many of these events are observed, you should begin to see if large table scans are occurring in the database. Like the ‘db file scattered read’ event, each wait event has its own meaning and individual end-resolution diagSQL Analysis - Users

Much of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured database down in terms of end user response times. SQL statements that use much memory can also cause a problem in a database.

Before you can identify problem SQL in your database, you have to ask the question, “What is bad SQL?” What criteria do you use when you begin the hunt for problem SQL in your critical systems? Understand that even the seasoned experts disagree on what constitutes efficient and inefficient SQL; so there is no way to sufficiently answer this question to every Oracle professional's satisfaction. The SQL Analysis for Users shows what SQL statements have consumed the largest percentages of executions, parse calls, CPU time, and elapsed time. The table below lists some general criteria you can use when evaluating the output from various database monitors or personal diagnostic scripts:

Criteria	Description
Overall Response (Elapsed) Time	The amount of time the query took to parse, execute, and fetch the data needed to satisfy the query. It should not include the network time needed to make the round trip from the requesting client workstation to the database server. This statistic is available in Oracle9i or later.
CPU Time	The amount of CPU time the query took to parse, execute, and fetch the data needed to satisfy the query.

Criteria	Description
Physical I/O	This is often used as the major statistic in terms of identifying good vs. bad SQL, this is a measure of how many disk reads the query caused to satisfy the user's request. While you certainly want to control disk I/O where possible, it is important that you not focus solely on physical I/O as the single benchmark of inefficient SQL. Make no mistake, disk access is slower than memory access and also consumes processing time making the physical to logical transition, but you need to look at the entire I/O picture of a SQL statement, which includes looking at a statements' logical I/O as well.
Logical I/O	The number of memory reads the query took to satisfy the user's request. The goal of tuning I/O for a query should be to examine both logical and physical I/O, and use appropriate mechanisms to keep both to a minimum.
Repetition	The number of times the query has been executed. A problem in this area is not as easy to spot as the others unless you know your application well. A query that takes a fraction of a second to execute may still be causing a headache on your system if it has executed erroneously (for example, a query that executes in a runaway PL/SQL loop) over and over again.

There are other criteria that you can examine like sort activity or access plan statistics (that show items like Cartesian joins and the like), but more often than not, these measures are reflected in the criteria listed above.

Fortunately, Oracle records all the above measures (some only in 9i), which makes tracking the SQL that has been submitted against an Oracle database much easier.

Location	Oracle Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	When you begin to look for inefficient SQL in a database, there are two primary questions you want answered: *What HAS been the worst SQL that has historically run in my database?* *What IS the worst SQL that is running right now in my database?*
Correction	When troubleshooting a slow system, you should be on the lookout for any query that shows an execution count that is significantly larger than any other query on the system. It may be that the query is in an inefficient PL/SQL loop, or other problematic programming construct. Only by bringing the query to the attention of the application developers will you know if the query is being mishandled from a programming standpoint. There is the possibility that the SQL statement just is not tuned well. To determine that, you can go further into Performance Analyst's Top SQL view and, if you have Embarcadero SQL Tuner installed, you can port the SQL over to SQL Tuner to better optimize the statement.

Transaction Analysis pane

On a transaction-oriented system, user transactions perform COMMITs to make permanent any data additions, changes, or deletions to the database. Or they ROLLBACK any modifications to undo any work that has encountered an error or other transaction-halting situation. The transaction analysis breakdown shows current commits and rollbacks and the percentage that each activity accounts for on the database.

Metrics	Large volumes of rollbacks can indicate a transaction breakdown on the system, that may need investigation at the application level.

Workload Analysis pane

When the database population as a whole experiences a system slowdown, it is not uncommon to find one or two users who are responsible for bringing the system to its knees. In the best of worlds, users have a pretty evenly divided amount of memory usage, disk I/O, CPU utilization, and parse activity. However, users submit large batch jobs during peak OLTP activity, or when sessions that are fire off untuned queries on a critical system.

>

If you are seeing a slowdown in your database, and cannot seem to find a root cause, examine the resource consumption of the leading sessions on a system. A handful of users can overwhelm the I/O capabilities of Oracle (through untuned queries or runaway batch jobs) or hammer the CPU or memory structures of the database and server.

Performance Analyst makes it easy to pinpoint the top sessions by showing the leading processes at the physical I/O, logical I/O, memory, disk sort, table scan, and CPU usage levels.

Metrics	If any one session uses more than 50% of a total resource (CPU, memory, etc.) go into the session to find out what they are currently executing.

Top Sessions Tab

The Top Sessions tab includes the following sections:

Top CPU Process	Top I/O Process	Top Memory Process
Top Sessions		

Top I/O Process

This section identifies the Oracle process that currently has caused the most I/O usage on the database.

Location	Oracle Performance Analyst statistics > Users > Top Sessions Tab
Metrics	To obtain more details on the top I/O process, locate the SID in the Top Sessions grid and drill down to obtain more granular information.

Top Memory Process

The Top Memory Process section identifies the Oracle process that currently is using the highest percentage of memory in the database.

Location	Oracle Performance Analyst statistics > Users > Top Sessions Tab
Metrics	To obtain more details on the top memory process, locate the SID in the Top Sessions grid and drill down to obtain more granular information.

Top CPU Process

The Top CPU Process section identifies the Oracle process that currently has the highest percentage of CPU usage in the database.

Location	Oracle Performance Analyst statistics > Users > Top Sessions Tab
Metrics	To obtain more details on the top CPU process, locate the SID in the Top Sessions grid and drill down to obtain more granular information.

Top Sessions

When a system experiences heavy activity, sometimes you will find that all the user connections are contributing somewhat equally to the overall load. More often than not, however, one or two user connections are causing most of the activity. It may be that a large batch load or other typical process is running that is perfectly okay for your system. Or it may be a runaway process or other rogue connection that needs to be tracked down and possibly eliminated.

>

The Top Sessions section displays information regarding all key metrics for all current sessions on a database. The table below describes the information available in this section:

Column	Description
SID	The system ID of the process.
User Name	The username of the process.
O/S ID	The operating system ID of the process.
Machine Name	The workstation where the process originated.
Logon Time	The date/time the process logged on to Oracle.
Tot Physical I/O	A summation of all physical reads and physical writes for the process (includes direct).
Tot Logical I/O	A summation of all logical I/O (buffer gets, consistent gets, etc.) for the process.
Hit Ratio	Shows how often object definitions are found in memory vs. having to read them in from disk.
Total Memory Usage	The total amount of memory used.
Parses	The total number of parses the process has caused.
Total CPU	The total CPU usage for the process.
Parse CPU	The amount of process CPU caused by parse activity.
Recursive CPU	The amount of process CPU caused by recursive calls.
Other CPU	Total CPU usage minus parse and recursive activity. This can be a negative number if the Oracle kernel is not properly reporting total CPU usage.
Disk Sorts	The total number of disk sorts caused by the process.
Memory Sorts	The total number of memory sorts caused by the process.
Rows Sorted	The total number of rows sorted by the process.
Commits	The total number of commits executed by the process.
Rollbacks	The total number of rollbacks executed by the process.
Executions	The total number of executions.
Physical Reads	The total number of physical reads produced by the process.
DB Block Gets	The total number of block gets produced by the process.
Consistent Gets	The total number of consistent gets produced by the process.
Consistent Changes	The total number of consistent changes produced by the process.
Last Wait	Event Defines the last wait event experienced by the session
Wait Class	The wait class of the last wait event
Wait time	The sessions last wait time
P1text	Text that describes the first wait parameter
P2text	Text that describes the second wait parameter
P2 r	The second wait parameter
P2Raw	The data for the second parameter
P3text	Text that describes the third wait parameter
P3	The third wait parameter
P3Raw	The data for the third parameter

Location	Oracle Performance Analyst statistics > Users > Top Sessions Tab
Metrics	To obtain more details on any process, drill down to obtain more granular information.

System Waits Tab

The System Waits tab includes the following sections:

System Waits	Wait Percent by Time Waited	Wait Percent by Total Waits

System Waits

Waits on a system generally occur for three reasons:

1. A process waits because it has no work to do.
2. A process waits because a requested resource is not available.
3. A process waits for Oracle to perform a prerequisite task for its given operation.

Idle waits (processes waiting because they have no work) are not normally a problem, however the other two wait causes are the ones worth your time and investigation. From a global database level, there are many different types of waits and sources of contention. The system Waits section presents all the various system waits that have occurred on the system since startup. The table below lists the information available in this section:

Column	Description
Event	The name of the wait event.
Class	Wait classes contain a collection of related wait events (like User I/O, System I/O, etc.)
Total Waits	The total number of waits for the event.
Percent of Total	The percentage of all waits that this event makes up.
Time Waited (Secs)	The total amount of time waited for the event, in seconds.
Percent Time Waited	The percentage of time that this event makes up.
Total Timeouts	The total number of timeouts for the event.
Percent Total Timeouts	The percentage of timeouts that this event makes up.
Average Waits (Secs)	The average amount of time waited for the event, in seconds.

Location	Oracle Performance Analyst statistics > Users > System Waits Tab
Metrics	Which waits should be a cause for concern and which waits can be ignored (for the most part)? Common wait events that can generally be ignored include **SQL*Net more data from client**, **SQL*Net message from client**, **client message**, **rdbms ipc message**, **pipe get**, **pmon timer**, **smon time**r, and **Null event**.

>

Correction

The table below describes wait events that deserve attention:

Wait Event	Description
Enqueue waits	If many enqueue waits are seen, this normally indicates either contention for certain rows in the database, or the result of database-initiated lock escalation. You should examine the use of indexes to make sure all referencing foreign keys are indexes and that SQL is tuned to not tarry over rows in the database during modification operations. Enqueue waits can also be the result of space management tasks (such as objects extending) and disk sorts (mainly in tablespaces that do not make use of the TEMPORARY tablespace parameter).
Buffer busy waits	Buffer busy waits normally center around contention for rollback segments, too small an INITRANS setting for tables, or insufficient free lists for tables. The remedy for each situation would be increasing the number of rollback segments, or altering tables to have larger settings for INITRANS to allow for more transactions per data block, and more free lists.
db file scattered read waits	An interesting wait event is the db file scattered read event. This event is indicative of table scans occurring on the system. Large numbers of them may indicate heavy scan activity and the need to revisit your indexing/physical design.
Latch free waits	Latch free waits indicate contention for latches.

The SQL*Net message to client can help a DBA diagnose wait activity outside the database.

The key statistics for waits is not the number of waits for each event, but rather the wait time and average wait time. These indicate if the event is truly causing a significant bottleneck in the database. To view the wait times for sessions and overall system wait events, you must set the timed_statistics parameter to TRUE for your Oracle databases. You can either set this parameter in your Init.ora file or alter the system dynamically with the ALTER SYSTEM SET TIMED_STATISTICS=TRUE command.

Wait Percent by Total Waits

Waits on a system generally occur for three reasons:

- A process waits because it has no work to do.
- A process waits because a requested resource is not available.
- A process waits for Oracle to perform a prerequisite task for its given operation.

Idle waits (processes waiting because they have no work) are not normally a problem, however the other two wait causes are the ones worth your time and investigation. From a global database level, there are many different types of waits and sources of contention. The Wait Percent by Total Waits section displays the percentage of waits for all top non-idle wait events.

Location	Oracle Performance Analyst statistics > Users > System Waits Tab
Metrics	Which waits should be a cause for concern and which waits can be ignored (for the most part)? Common wait events that can generally be ignored include **SQL*Net more data from client**, **SQL*Net message from client**, **client message**, **rdbms ipc message**, **pipe get**, **pmon timer**, **smon time**r, and **Null event**.

>

Correction

The table below describes wait events that deserve attention:

Wait Event	Description
Enqueue waits	If many enqueue waits are seen, this normally indicates either contention for certain rows in the database, or the result of database-initiated lock escalation. You should examine the use of indexes to make sure all referencing foreign keys are indexes and that SQL is tuned to not tarry over rows in the database during modification operations. Enqueue waits can also be the result of space management tasks (such as objects extending) and disk sorts (mainly in tablespaces that do not make use of the TEMPORARY tablespace parameter).
Buffer busy waits	Buffer busy waits normally center around contention for rollback segments, too small an INITRANS setting for tables, or insufficient free lists for tables. The remedy for each situation would be increasing the number of rollback segments, or altering tables to have larger settings for INITRANS to allow for more transactions per data block, and more free lists.
db file scattered read waits	An interesting wait event is the db file scattered read event. This event is indicative of table scans occurring on the system. Large numbers of them may indicate heavy scan activity and the need to revisit your indexing/physical design.
Latch free waits	Latch free waits indicate contention for latches.

The SQL*Net message to client can help a DBA diagnose wait activity outside the database.

Wait Percent by Time Waited

Waits on a system generally occur for three reasons:

1. A process waits because it has no work to do.
2. A process waits because a requested resource is not available.
3. A process waits for Oracle to perform a prerequisite task for its given operation.

Idle waits (processes waiting because they have no work) are not normally a problem, however the other two wait causes are the ones worth your time and investigation. From a global database level, there are many different types of waits and sources of contention. The Wait Percent by Time Waited section displays the percentage of waits for all top non-idle wait events.

Location	Oracle Performance Analyst statistics > Users > System Waits Tab
Metrics	The key statistics for waits is not the number of waits for each event, but rather the wait time and average wait time. These indicate if the event is truly causing a significant bottleneck in the database. Which waits should be a cause for concern and which waits can be ignored (for the most part)? Common wait events that can generally be ignored include **SQL*Net more data from client**, **SQL*Net message from client**, **client message**, **rdbms ipc message**, **pipe get**, **pmon timer**, **smon time**r, and **Null event**.

Correction

The table below describes wait events that deserve attention:

Wait Event	Description
Enqueue waits	If many enqueue waits are seen, this normally indicates either contention for certain rows in the database, or the result of database-initiated lock escalation. You should examine the use of indexes to make sure all referencing foreign keys are indexes and that SQL is tuned to not tarry over rows in the database during modification operations. Enqueue waits can also be the result of space management tasks (such as objects extending) and disk sorts (mainly in tablespaces that do not make use of the TEMPORARY tablespace parameter).

Wait Event	Description
Buffer busy waits	Buffer busy waits normally center around contention for rollback segments, too small an INITRANS setting for tables, or insufficient free lists for tables. The remedy for each situation would be increasing the number of rollback segments, or altering tables to have larger settings for INITRANS to allow for more transactions per data block, and more free lists.
db file scattered read waits	An interesting wait event is the db file scattered read event. This event is indicative of table scans occurring on the system. Large numbers of them may indicate heavy scan activity and the need to revisit your indexing/physical design.
Latch free waits	Latch free waits indicate contention for latches.

The SQL*Net message to client can help a DBA diagnose wait activity outside the database.

To view the wait times for sessions and overall system wait events, you must set the timed_statistics parameter to TRUE for your Oracle databases. You can either set this parameter in your Init.ora file or alter the system dynamically with the ALTER SYSTEM SET TIMED_STATISTICS=TRUE command.

Session Waits Tab

The Session Waits tab includes the following sections:

Current Waits	Historical Waits	

Historical Waits

Waits on a system generally occur for three reasons:

1. A process waits because it has no work to do.
2. A process waits because a requested resource is not available.
3. A process waits for Oracle to perform a prerequisite task for its given operation.

Idle waits (processes waiting because they have no work) are not normally a problem, however the other two wait causes are the ones worth your time and investigation. From a global database level, there are many different types of waits and sources of contention. The Historical Waits section presents all the various waits that have occurred for all currently connected oracle sessions. The table below lists the information available in this section:

Column	Description
Event	The name of the wait event.
Wait Class	Wait classes contain a collection of related wait events (like User I/O, System I/O, etc.)
Total Waits	The number of waits for the event.
Percent of Total Waits	The percentage of all waits that this event makes up.
Time Waited (Secs)	The total amount of time waited for the event, in seconds.
Average Wait (Secs)	The average amount of time waited for the event, in seconds.
Max Wait (Secs)	The maximum amount of time waited for the event, in seconds.
Total Timeouts	The number of timeouts for the event.

Location	Oracle Performance Analyst statistics > Users > Session Waits Tab Oracle Performance Analyst statistics > Session Details > Waits Tab

Metrics	Which waits should be a cause for concern and which waits can be ignored (for the most part)? Common wait events that can generally be ignored include **SQL*Net more data from client**, **SQL*Net message from client**, **client message**, **rdbms ipc message**, **pipe get**, **pmon timer**, **smon time**r, and **Null event**.

Correction

The table below describes wait events that deserve attention:

Wait Event	Description
Enqueue waits	If many enqueue waits are seen, this normally indicates either contention for certain rows in the database, or the result of database-initiated lock escalation. You should examine the use of indexes to make sure all referencing foreign keys are indexes and that SQL is tuned to not tarry over rows in the database during modification operations. Enqueue waits can also be the result of space management tasks (such as objects extending) and disk sorts (mainly in tablespaces that do not make use of the TEMPORARY tablespace parameter).
Buffer busy waits	Buffer busy waits normally center around contention for rollback segments, too small an INITRANS setting for tables, or insufficient free lists for tables. The remedy for each situation would be increasing the number of rollback segments, or altering tables to have larger settings for INITRANS to allow for more transactions per data block, and more free lists.
db file scattered read waits	An interesting wait event is the db file scattered read event. This event is indicative of table scans occurring on the system. Large numbers of them may indicate heavy scan activity and the need to revisit your indexing/physical design.
Latch free waits	Latch free waits indicate contention for latches.

The SQL*Net message to client can help a DBA diagnose wait activity outside the database.

The key statistics for waits is not the number of waits for each event, but rather the wait time and average wait time. These indicate if the event is truly causing a significant bottleneck in the database. To view the wait times for sessions and overall system wait events, you must set the timed_statistics parameter to TRUE for your Oracle databases. You can either set this parameter in your Init.ora file or alter the system dynamically with the ALTER SYSTEM SET TIMED_STATISTICS=TRUE command.

Current Waits

Waits on a system generally occur for three reasons:

- A process waits because it has no work to do.
- A process waits because a requested resource is not available.
- A process waits for Oracle to perform a prerequisite task for its given operation.

Session contention is merely a subset of contention that is viewable at the global database level. Oftentimes, it takes analysis at the session level to pinpoint the exact source of contention that is occurring globally. So you need to become accustomed to viewing contention statistics and waits at the user process level.

When monitoring waits with respect to user sessions, there are two areas of interest:

1 What HAS the user session been waiting on?

2 What IS the user session waiting on?

>

Oracle records both sets of wait statistics for you. In reviewing previous waits for a session, you can see which types of things have caused the session to be bottlenecked. The Current Waits section displays information to help you troubleshoot session contention. The table below lists the information available in this section:

Column	Description
Event	The name of the wait event.
Wait (Secs)	The number of seconds in the current wait.
Last Wait Time	A non-zero value is the session's last wait time. A zero value means the session is currently waiting. A -2 value indicates the platform cannot support a fast timing mechanism or TIMED_STATISTICS is not set for the system.
State	Indicates the possible values. They includes the following: WAITING - the session is waiting. WAITED UNKNOWN TIME - duration of last wait is not known. WAITED SHORT TIME - last wait was less than 1/100th of a second. WAITED KNOWN TIME - the wait is equal to the time of the last wait.
P1 - P3	Any additional wait parameters.
P1TEXT - P3TEXT	Any descriptions of additional parameter.
P1RAW - P3RAW	Any additional wait parameters.

Location	Oracle Performance Analyst statistics > Users > Session Waits Tab Oracle Performance Analyst statistics > Session Details > Waits Tab
Metrics	Which waits should be a cause for concern and which waits can be ignored (for the most part)? Common wait events that can generally be ignored include **SQL*Net more data from client**, **SQL*Net message from client**, **client message**, **rdbms ipc message**, **pipe get**, **pmon timer**, **smon time**r, and **Null event**.

Correction

The table below describes wait events that deserve attention:

Wait Event	Description
Enqueue waits	If many enqueue waits are seen, this normally indicates either contention for certain rows in the database, or the result of database-initiated lock escalation. You should examine the use of indexes to make sure all referencing foreign keys are indexes and that SQL is tuned to not tarry over rows in the database during modification operations. Enqueue waits can also be the result of space management tasks (such as objects extending) and disk sorts (mainly in tablespaces that do not make use of the TEMPORARY tablespace parameter).
Buffer busy waits	Buffer busy waits normally center around contention for rollback segments, too small an INITRANS setting for tables, or insufficient free lists for tables. The remedy for each situation would be increasing the number of rollback segments, or altering tables to have larger settings for INITRANS to allow for more transactions per data block, and more free lists.
db file scattered read waits	An interesting wait event is the db file scattered read event. This event is indicative of table scans occurring on the system. Large numbers of them may indicate heavy scan activity and the need to revisit your indexing/physical design.
Latch free waits	Latch free waits indicate contention for latches.

The SQL*Net message to client can help a DBA diagnose wait activity outside the database.

>

The key statistics for waits is not the number of waits for each event, but rather the wait time and average wait time. These indicate if the event is truly causing a significant bottleneck in the database. To view the wait times for sessions and overall system wait events, you must set the timed_statistics parameter to TRUE for your Oracle databases. You can either set this parameter in your Init.ora file or alter the system dynamically with the ALTER SYSTEM SET TIMED_STATISTICS=TRUE command.

Locks/Blocking Locks Tab

The Locks/Blocking Locks tab includes the following sections:

All Locks	Blocking Locks	

All Locks

To modify database information or structures, a user session must obtain a lock on the object to perform its task. Ibesides user locks, Oracle itself will issue lock requests to carry out its internal duties. The All Locks section displays information about all locks currently on a system. The table below lists the information available in this section:

Column	Description
Machine	The workstation of the Oracle process.
Serial #	The serial number of the session holding the lock.
Request	The lock requests that Oracle issues to carry out its internal duties.
Table	The table name of the locked table.
Mode	The lock mode (EXCLUSIVE, SHARE, etc.).
Blocking?	Indicates if the lock is a blocking lock.

Location	Oracle Performance Analyst statistics > Users > Locks/Blocking Locks Tab Oracle Performance Analyst statistics > Session Details > Locks Tab
Metrics	Locks that are held for unusually long periods of time may be candidates for further investigation. The application logic may be inefficient or perhaps the program is not issuing frequent enough COMMITs.

Blocking Locks

Without a doubt, blocking lock situations can give the appearance of a frozen database almost more than anything else (except, perhaps, for a stuck archiver process). A single blocking user has the potential to stop work for nearly all other processes on a small system, and can cause major headaches even on large systems. Although Oracle supports unlimited row-level locking, blocking lock situations do crop up - sometimes frequently.

The Blocking Locks section contains information relating to user accounts that are currently blocked and the sessions that are blocking them. The table below describes the information available in this section:

Column	Description
Blocked User	The user account of the session waiting for the lock.
Blocking User	The user account of the session holding the offending lock.
Waiting SID	The session identifier of the session waiting for the lock.
Holding SID	The session identifier of the session holding the offending lock.
Type	The type of lock (TRANSACTION, DML, etc.).

>

Column	Description
Mode	The lock mode (EXCLUSIVE, SHARE, etc.).
Request	The type of lock being requested by the session.
Row	The rowid of the row being held.
ID 1	Lock identifier #1 (depends on type).
ID 2	Lock identifier #2 (depends on type).
Time Waited (Secs)	The amount of time, in seconds, the blocked user has been waiting.
Locked Object	The name of the object being locked.

Location	Oracle Performance Analyst statistics > Users > Locks/Blocking Locks Tab Oracle Performance Analyst statistics > Session Details > Locks Tab
Metrics	Once discovered, a blocking lock situation can normally be quickly remedied - the DBA issues a KILL against the offending process, which eliminates the user's stranglehold on the objects they were accessing. Other user processes then nearly almost always complete in an instant. Discovering the blocked lock situation is made easier by using tools like Performance Analyst, but preventing the blocking lock situation in the first place is where it gets tricky. The culprit of blocking lock scenarios is usually the application design, or the SQL being used within the application itself. Properly coding an application to reference database objects in an efficient order, and then using the right SQL to get the job done, is an art. Most DBAs who have had to face Oracle Forms applications have suffered through the dreaded SELECT … FOR UPDATE statements that place unnecessary restrictive locks on nearly every read operation, and know all too well that good coding practice is important. The key to avoiding lock contention is to process user transactions in the quickest and most efficient manner possible - something not always easy to do. Data warehouses (whose data is mostly read) can benefit from tablespaces set in READ ONLY mode. READ ONLY signals to the other database that exclusive locks need not be used for the data contained within the tablespace. This is especially helpful in Oracle Parallel Server environments and drastically reduces ping activity.

RAC Detail View

The Performance Analyst for Oracle has been upgraded to support for the monitoring of Oracle Real Application Clustering (RAC) environments (Oracle 9i and higher).

Oracle RAC allows a single physical Oracle database to be accessed by simultaneous instances of Oracle running across several CPUs. All statistics you see are listed by instance ID.

The RAC Detail view is divided into the following tabs:

Overview Tab	Memory	Ping Activity
Locks	Sessions	

Overview

The Overview tab is of the RAC Detail view divided into the following sections:

Key Ratio Analysis pane	Instance Analysis pane	Bottleneck Analysis pane

Workload Analysis pane		

Key Ratio Analysis pane

Database performance analysts typically use one of two methods for examining the performance levels of a database: ratio-based or wait/bottleneck-based analysis. Ratio-based analysis involves examining a number of key database ratios that can be used to indicate how well a database is running. Performance ratios serve as very good roll-up mechanisms for busy DBAs to use for at-a-glance performance analysis. Many DBAs have large database farms to contend with and can't spend time checking detailed wait-based analysis outputs for each and every database they oversee. Succinctly presented performance ratios can assist in such situations by giving DBAs a few solid indicators that can be quickly scanned to see if any database needs immediate attention.

While there are many opinions as to what rules to follow, there are some standards that should always be adhered to. To start with, many of the formulas that make up ratio-based analysis must be derived from delta measurements instead of cumulative statistics. Many of the global ratios that a DBA will examine come from the `v$sysstat` performance view. This view maintains a count of all the occurrences (in the VALUE column) of a particular database incident (in the NAME column) since the database was brought up. For databases that are kept up for long periods of time, these values can grow quite large and impact how a particular ratio that a DBA may be looking at is interpreted. However, if delta statistics are used (taking, for a specified sampling period, the before and after counts of each statistic that make up a ratio), then an accurate and current portrayal of the various ratios can be had.

A final thing to remember about using ratio-based analysis is that, while there are a number of rules of thumb that can be used as starting points to begin the evaluation of database performance, DBAs must determine each database's individual "personality"with respect to the key performance ratios. Some hard and fast rules simply do not apply to every database. The danger in using blanket ratio standards is that they can lead the DBA to haphazardly take action, which can at times contribute nothing to the situation, and sometimes even degrade performance.

The following ratios are used on the Performance Analyst RAC Overview page to succinctly communicate the general overall performance levels of the monitored database:

- Cross Instance Cache Flush ratio
- Ping Ratio
- Local Buffer Access Percent
- Local Read Percent
- Global Buffer Busy Wait Percent

Instance Analysis pane

As a DBA managing an Oracle RAC, you will oftentimes want a quick performance snapshot of the efficiency and activity across all your Oracle RAC nodes. The Instance Analysis section provides a fast picture of which nodes of your Oracle RAC appear overworked and those that are underutilized.

The following statistics are provided:

ID: The Oracle RAC instance ID

·**Name**: The instance name

Sessions: The total number of sessions logged onto a RAC node

Physical I/O: A summation of all physical reads and physical writes for the RAC node

Logical I/O: A summation of all logical I/O (buffer gets, consistent gets, etc.) for the RAC node

Hit Ratio: An indicator of how often user requests for data are satisfied.

Memory Usage: The session's total memory usage against all session memory usage on the database.

Parses: The total number of parses experienced by the RAC node

Total CPU: The total CPU usage for the RAC node

Parse CPU: The amount of process CPU for the RAC node

Recursive CPU: The amount of process CPU caused by recursive calls on the RAC node

Other CPU: Total CPU usage minus parse and recursive activity. This can be a negative number if the Oracle kernel is not properly reporting total CPU usage

Disk Sorts: The total number of disk sorts for the RAC node

Memory Sorts: The total number of memory sorts for the RAC node

Rows Sorts: The total number of rows sorted on the RAC node

Commits: The total number of commits executed for the RAC node

Rollbacks: The total number of rollbacks executed for the RAC node

Executions: The total number of calls made on the RAC node

Physical Reads: The total number of physical reads produced on the RAC node

DB Block Gets: The total number of block gets produced on the RAC node

Consistent Gets: The total number of consistent gets produced on the RAC node

Consistent Changes: The total number of consistent changes produced on the RAC node

Bottleneck Analysis pane

When an Oracle database is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting can mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. And this is where wait-based or bottleneck analysis comes into play. DBAs used this form of performance analysis to determine if perceived bottlenecks in a database are contributing to a performance problem.

Bottleneck analysis is a valid method of measuring performance because it helps a DBA track where a database has been spending its time. If latch contention or heavy table scan activity has been dragging a database's performance down, a DBA can use bottleneck analysis to confirm the actual root cause. Once one or more wait events or other bottlenecks have been pinpointed as possible performance vampires, the DBA can drill down and oftentimes discover a fair amount of detail about what sessions and objects are causing the problem.

For wait event analysis to be properly carried out, it is imperative that the *timed_statistics* initialization parameter be set to TRUE. By default this parameter is set to FALSE, which disallows the collection of wait times for each wait event defined to the Oracle engine. For one to really understand the impact of wait events on database performance, the DBA not only needs to discover what the database is or has been waiting on, but the durations of the waits. Having both measures in hand allows a complete picture to be formed regarding the magnitude of wait-initiated performance degradations. Almost all Oracle experts now agree that allowing the collection of timing statistics adds little if anything to database overhead, so setting *timed_statistics* to TRUE should not be a worry. The Performance Analyst home page identifies the top system and session waits that are currently a cause of lessened performance.

When using bottleneck analysis, you cannot rely only on information contained in the wait event views that Oracle provides. For example, an object may attempt to extend into another extent of space in a tablespace and yet be denied if no such free space exists. Such a failure will not be reflected in any wait event, but still represents a very real bottleneck to the database. In the same way that you cannot depend on only a few ratios to properly carry out ratio-based performance analysis, an administrator must include other statistical metrics in their overall bottleneck analysis framework to obtain an accurate performance risk assessment. Performance Analyst helps identify bottlenecks in your database that fall outside of pure wait events so you can get a total picture of all stoppages in your system.

For an Oracle RAC, Performance Analyst highlights the top RAC-related waits as well as conflicts that may be slowing down response times. Here you see statistics for Top Rac-Related Waits and Conflicts. Within each section, you find:

Top Rac-Related Waits

When viewing wait statistics, there are several levels of detail that a DBA can drill down into. The first level is the system view, which provides a global, cumulative snapshot of all the waits that have occurred on a system. Viewing these numbers can help a DBA determine what wait events have caused the most commotion in a database thus far. The Top RAC-related bottlenecks section display identifies the top waits that have occurred on the node that Performance Analyst is currently connected to.

Appendix A in the Oracle Reference manual contains a listing and description of every current wait event defined in Oracle. DBAs unfamiliar with what each event represents should keep this listing close by as they examine wait-based event metrics. For example, a “db file scattered read” event is typically indicative of table scan operations. If many of these events are observed, then the DBA should begin to see if large table scans are occurring in the database. Like the “db file scattered read” event, each wait event has its own meaning and individual end-resolution diagnosis.

After looking at system-level wait activity, a DBA can drill down further to discover what current connections may be responsible for any reported waits that are being observed at the system level. Performance Analyst reports on historical and current wait events at the session level, making this investigation easy to accomplish.

Conflicts

There are a number of conflicts that can occur over an Oracle RAC that can reduce response times. The Conflicts section highlights some of the more common conflicts and provides a count of how many times each has occurred.

The conflicts listed include:

Library Cache: How many library cache invalidations have occurred on the node that Performance Analyst is connected to.

Data Dictionary Cache: How many data dictionary cache invalidations have occurred on the node that Performance Analyst is connected to.

Forced Writes (Oracle 9.2 and above): Number of times a block was written to cache because an instance had used the block, but another instance had requested the lock on the block in a conflicting mode.

Pings: How many cross instance writes occurred

Cache Timeouts: The number of global cache convert timeouts that have occurred.

Workload Analysis pane

When your phone starts ringing with complaints of performance slowdowns, one of the first things you should get a handle on is:

- Who is currently connected to the database
- What resources are they using
- What are they currently executing

The Workload Activity Analysis section of the Performance Analyst home page provides insight into the leading resource hogs across all instances of an Oracle RAC. Drill down’s are available so you can easily get detailed information into what each leading session is currently involved with.

If you are seeing a slowdown in your database, and can't seem to find a root cause, one thing to examine is the resource consumption of the leading sessions on a system. Often you will find one or a handful of users making life miserable for everyone else. They accomplish this by overwhelming the I/O capabilities of Oracle (through untuned queries or runaway batch jobs) or hammering the CPU or memory structures of the database and server.

The Workload Analysis section presents:

- Top Physical I/O Process
- Top Logical I/O Process
- Top Memory Process
- Top CPU Process

For each of these, you see the SID, Instance ID, Username, %Used.

If any one session appears to be using more than 50% of a total resource (CPU, memory, etc.), then you should drill down into the session to find out what it is currently executing.

Memory

The Memory tab of the RAC Detail view is divided into the following sections:

Key Ratio Analysis pane	Memory Analysis pane	Library Cache Analysis pane
Data Dictionary Cache Analysis pane		

Key Ratio Analysis pane

Database performance analysts typically use one of two methods for examining the performance levels of a database: Ratio-based or wait/bottleneck-based analysis. Ratio-based analysis involves examining a number of key database ratios that can be used to indicate how well a database is running. Performance ratios serve as very good roll-up mechanisms for busy DBAs to use for at-a-glance performance analysis. Many DBAs have large database farms to contend with and can't spend time checking detailed wait-based analysis outputs for each and every database they oversee. Succinctly presented performance ratios can assist in such situations by giving DBAs a few solid indicators that can be quickly scanned to see if any database needs immediate attention.

While there aremany opinions as to what rules to follow, there are some standards that should always be adhered to. To start with, many of the formulas that make up ratio-based analysis must be derived from delta measurements instead of cumulative statistics. Many of the global ratios that a DBA examines come from the `v$sysstat` performance view. This view maintains a count of all the occurrences (in the VALUE column) of a particular database incident (in the NAME column) since the database was brought up. For databases that are kept up for long periods of time, these values can grow quite large and will impact how a particular ratio that a DBA may be looking at is interpreted. However, if delta statistics are used (taking, for a specified sampling period, the before and after counts of each statistic that make up a ratio), then an accurate and current portrayal of the various ratios can be had.

A final thing to remember about using ratio-based analysis is that, while there are a number of rules of thumb that can be used as starting points to begin the evaluation of database performance, DBAs must determine each database's individual "personality" with respect to the key performance ratios. Some hard and fast rules simply do not apply to every database. The danger in using blanket ratio standards is that they can lead the DBA to haphazardly take action, which can at times contribute nothing to the situation, and sometimes even degrade performance.

The following ratios are used in the Memory Key Ratio Analysis section to succinctly communicate the general overall performance levels of the monitored database:

- **Buffer Cache Hit Ratio**: How often requested information is found in memory versus on disk.
- **Library Cache Hit Ratio**: Shows how often SQL code is being reused by other database users versus the number of times a SQL statement is broken down, parsed, and then loaded (or reloaded) into the shared pool.
- **Latch Miss Ratio**: Defines the number of times a process obtained a willing-to-wait latch versus missing the attempt.

Memory Analysis pane

The Memory Analysis section displays the memory configurations for each node currently up in the RAC. The following information is shown:

Name: The name of the instance.

DB Buffers (MB): The default memory cache that maintains data blocks when they are read from the database. If the DBA doesn't specifically place objects in another data cache (which will be covered next), then any data requested by clients from the database will be placed into this cache. This memory area is controlled by the *db_block_buffers* parameter in Oracle8i and below and *db_cache_size* in Oracle9i and above.

Log Buffer Size (MB): This area buffers modifications that are made to the database before they are physically written to the redo log files. The *log_buffer* configuration parameter controls this memory area.

Shared Pool Size (MB): This familiar area holds object structure as well as code definitions, and other metadata. Setting the proper amount of memory in the shared pool assists a great deal in improving overall performance with respect to code execution.

DB Caches (MB): Beginning in Oracle9i, a DBA can create tablespaces whose blocksize differs from the overall database blocksize. When data is read into the SGA from these tablespaces, their data has to be placed into memory regions that can accommodate their special block size. Oracle9i and above has memory settings for 2K, 4K, 8K, 16K, and 32K caches. The configuration parameter names are in the pattern of *db_nk_cache_size*. For each RAC node, you see the following statistics:

- Buffer Cache Hit Ratio
- Library Cache Hit Ratio
- Latch Miss Ratio

Library Cache Analysis pane

The Library Cache Analysis section provides a drill down view into invalidations that have occurred across the cluster. The following information is shown:

ID: The instance ID.

Name: The instance name.

Namespace: The area of the library cache.

DLM Lock Requests: The number of get requests for instance locks.

DLM Pin Requests: The number of pin requests for instance locks.

DLM Pin Releases: The number of release requests for pin instance locks.

DLM Invalidation Requests: The number of get requests for invalidation instance locks.

DLM Invalidations: The number of invalidation pings received from other RAC instances.

Data Dictionary Cache Analysis pane

The Data Dictionary Cache Analysis section provides a drill-down view into invalidations and conflicts that have occurred across the cluster. The following information is shown:

ID: The instance ID.

Name: The instance name.

Parameter: The specific object area of the dictionary cache.

DLM Requests: The number of lock manager requests.

DLM Conflicts: The number of lock manager conflicts encountered.

DLM Releases: The number of lock manager releases.

Ping Activity

Pinging is the process where one Oracle Instance requests another to write a set of blocks from its SGA to disk so it can obtain it in exclusive mode. To move a data block from one instance's SGA to another is a slow process. The Ping Activity tab of the RAC Detail view displays several levels of specificity, starting from general to specific. The Ping Activity tab is divided into the following sections:

Ping by Instance - shows how often pinging has occurred across all the monitored instances.

Ping by Datafile - shows the number of blocks pinged from the shared set of RAC datafiles.

Ping by Object - provides a drill down view into the amount of specific ping activity across the shared objects and datafiles of a RAC:

- Owner
- Object Name
- Partition Name
- Block Number
- Segment Type
- Kind
- Forced Reads
- Forced Writes
- Datafile

Locks

The Locks tab of the RAC Detail view is divided into the following sections:

Lock Convert Summary - The Lock Convert Summary view provides efficiency statistics regarding the speed at which the RAC acquires and converts locks across the instances. You see the average get and convert times for each RAC node.

DLM Lock Activity - The DLM (Distributed Lock Manager) Lock Activity view provides detailed information on the number of lock conversions that have occurred across the RAC.

PCM Lock Summary - The PCM (Parallel Cache Management) Lock Summary provides a summary count of PCM locks that are releasing, acquiring, or are invalid.

Latch Detail - Protecting the memory structures in Oracle's SGA are latches. Latches ensure that one and only one process at a time can run or modify any memory structure at the same instant. Much more restrictive than locks (which at least allow for some collective user interaction), latches have no queuing mechanism so either you get it or you don't and are forced to continually retry.

The latch detail tab of the contention detail view presents a detailed view of latch activity. Information presented includes:

- **Instance Name**: The name of the instance

- **Name**: The name of the latch
- **Gets**: The total number of times the latch was requested by a process
- **Misses**: The total number of failed attempts to acquire the latch on the first attempt
- **Immediate Gets:** The total number of no-wait requests for a latch
- **Immediate Misses**: The total number of failed no-wait attempts to acquire the latch on the first attempt
- **Immediate Sleeps:** The total number of requests that "paused" while waiting for a latch

Common indicators of latch contention are a latch miss ratio (which records willing-to-wait mode latch requests) and a latch immediate miss ratio (which records no-wait mode latch requests). These statistics reflect the overall health of how often latch requests were made and satisfied without waiting.

Here are a few latch contention situations that you should recognize and get to the bottom of quickly:

- **Cache buffer chain latch:** This latch is responsible for protecting paths to database block buffers in the buffer cache. Very high I/O loads tend to cause contention for this latch. You can alleviate contention somewhat by adding more buffers to the cache (through the db_block_buffers parameter) or by adding more LRU latch chain latches with the db_block_lru_latches parameter.
- **Library cache latches**: Protects cached SQL statements in the library cache area of the Oracle shared pool. Contention for this latch is the usual result of literals being used for SQL statements instead of bind variables.

Other routine latch contention problems used to include the redo allocation and redo copy latches, but these have pretty much been made obsolete in Oracle 8.1.5 and higher.

Sessions

The Sessions tab of the RAC Detail view is divided into the following sections:

User Session by Instance - The User Session by Instance section provides a graphical count of the number of connected sessions for each node in the RAC.

Blocks by Instance - The Blocks by Instance section provides a graphical count of the number of blocked sessions for each node in the RAC.

User Waits by Instance - The User Waits by Instance section provides a graphical count of the number of historical session waits (non-idle) for each node in the RAC.

Top Sessions - When a system experiences heavy activity, sometimes you'll find that all the user connections are contributing somewhat equally to the overall load. Often, however, one or two user connections are causing most of the activity. It may be that a large batch load or other typical process is running that is perfectly OK for your system. Or it may be a runaway process or other rogue connection that needs to be tracked down and possibly eliminated.

The top sessions control displays information regarding all key metrics for all current sessions on a database. Information presented includes:

- **SID**: The system ID of the process
- **Instance:** The instance the session is connected to
- **Username**: The username of the process
- **O/S ID**: The operating system ID of the process
- **Machine Name**: The workstation where the process originated
- **Logon Time:** The date/time the process logged on to Oracle
- **Tot Physical I/O**: A summation of all physical reads and physical writes for the process (includes direct)

>

- **Tot Physical I/O**: A summation of all logical I/O (buffer gets, consistent gets, etc.) for the process
- **Parses**: The total number of parses the process has caused
- **Total CPU:** The total CPU usage for the process
- **Parse CPU**: The amount of process CPU caused by parse activity
- **Recursive CPU**: The amount of process CPU caused by recursive calls
- **Other CPU:** Total CPU usage minus parse and recursive activity. This can be a negative number if the Oracle kernel is not properly reporting total CPU usage
- **Disk sorts:** The total number of disk sorts caused by the process
- **Memory sorts:** The total number of memory sorts caused by the process
- **Rows sorted:** The total number of rows sorted by the process
- **Commits**: The total number of commits executed by the process
- **Rollbacks:** The total number of rollbacks executed by the process
- **Physical reads:** The total number of physical reads produced by the process
- **DB block gets:** The total number of block gets produced by the process
- **Consistent gets:** The total number of consistent gets produced by the process
- **Consistent changes:** The total number of consistent changes produced by the process

Response Time Analysis View (Oracle 10G Only)

The Response Time Analysis View displays statistics that are a part of the Oracle 10g wait enhancements. By analyzing response times, you can review historical data and pinpoint who was on a database and what they were doing during a specific period of time. This allows you to drill deep into the root causes, if any of response time problems, up to 24 hours after the fact.

The main Response Time Analysis View displays the following tabbed pages:

Response Time	Active Time	System Bottlenecks
Session Bottlenecks	Wait Histogram	Historical Session Analysis

Response Time

The Response Time tab of the Response Time Analysis View displays the following statistics:

Current Workload	Current Response Times	Last Hour
Last Hour Detail		

Current Workload

The Current Workload section provides a breakdown of active and idle time for the Oracle database engine in terms of percentage used.

Location	Oracle Performance Analyst statistics > Response Time Analysis View (Oracle 10G Only) > Response Time
Metrics	Seeing consistently high percentages for the Database CPU Time (greater than 75-90%) should indicate a possibly overworked server.
Correction	If you see high numbers for the Database CPU Time, then begin looking into these statistics: *SQL Statements with high elapsed times and CPU times (use the Top SQL view to find these).* *Sessions with high wait times (especially in the User I/O wait class).*

Current Response Times

The Current Response Times section provides a quick view into the current average response time that a transaction experiences on the system along with the current average SQL service time. Both statistics are expressed in seconds.

Location	Oracle Performance Analyst statistics > Response Time Analysis View (Oracle 10G Only) > Response Time
Metrics	Defining ‘good’ and ‘bad’ metrics for these two measurements can only be accomplished by viewing these statistics over time and determining what ‘normal’ is for your system. You can get help by checking the Last Hour section, which details the minimum, maximum, and average values for these statistics as well as other key metrics that provide insight into how well your database is running.

Last Hour

The Last Hour section provides a window into various key performance metrics. Each metric's minimum, maximum, and average value is displayed.

Location	Oracle Performance Analyst statistics > Response Time Analysis View (Oracle 10G Only) > Response Time
Metrics	Key metrics to keep an eye on include:**CPU usage per Txn**, **Database CPU Time Ratio**, **Response Time Per Txn**, and **SQL Service Response Time**.
Correction	If you see high numbers for the Database CPU Time, then begin looking into these statistics: *SQL Statements with high elapsed times and CPU times (use the Top SQL view to find these).* *Sessions with high wait times (especially in the User I/O wait class).*

Last Hour Detail

The Last Hour Detail section provides a graphical look back over the last hour of a selected performance metric. Oracle10g and higher automatically maintains a running history of selected metrics so you can observe trends and patterns in performance.

Location	Oracle Performance Analyst statistics > Response Time Analysis View (Oracle 10G Only) > Response Time

Active Time

The Active Time tab of the Response Time Analysis detail view displays the following statistics:

System Time Summary	System Time Detail	Session Time Detail

System Time Summary

The System Time Summary section allows you to see where, in broad terms, the database engine has been spending the largest percentages of its time.

Location	Oracle Performance Analyst statistics > Response Time Analysis View (Oracle 10G Only) > Active Time

Metrics

The metrics shown here include (*note*, this table is largely from Oracle 10g documentation):

Metric	Description
DB Time	You see the amount of elapsed time (in microseconds) spent performing Database user-level calls. Instance background processes such as PMON are not included.
DB CPU	Amount of CPU time (in microseconds) spent on database user-level calls. Instance background processes such as PMON are not included.
background cpu time	Amount of CPU time (in microseconds) taken up by database background processes.
sequence load elapsed time	Amount of elapsed time spent getting the next sequence number from the data dictionary. If a sequence is cached, then this is the amount of time spent replenishing the cache when it runs out. No time is charged when a sequence number is found in the cache. For non-cached sequences, some time is charged for every nextval call.
parse time elapsed	Amount of elapsed time spent parsing SQL statements. Both soft and hard parse times are included.
hard parse elapsed time	Amount of elapsed time spent hard parsing SQL statements.
sql execute elapsed time	Amount of elapsed time SQL statements are executing. **NOTE**: For select statements this also includes the amount of time spent performing fetches of query results.
connection management call elapsed time	Amount of elapsed time spent performing session connect and disconnect calls.
failed parse elapsed time	Amount of time spent performing SQL parses that ultimately fail with some parse error.
hard parse (sharing criteria) elapsed time	Amount of elapsed time spent performing SQL hard parses when the hard parse resulted from not being able to share an existing cursor in the SQL cache.
hard parse (bind mismatch) elapsed time	Amount of elapsed time spent performing SQL hard parses when the hard parse resulted from bind type or bind size mismatch with an existing cursor in the SQL cache.
PL/SQL execution elapsed time	Amount of elapsed time spent running the PL/SQL interpreter. This does not include time spent recursively executing/parsing SQL statements or time spent recursively executing the Java VM.
PL/SQL compilation elapsed time	Amount of elapsed time spent running the PL/SQL compiler.
inbound PL/SQL rpc elapsed time	Time inbound PL/SQL remote procedure calls have spent executing. It includes all time spent recursively executing SQL and JAVA, and therefore is not easily related to "PL/SQL execution elapsed time".
Java execution elapsed time	Amount of elapsed time spent running the Java VM. This does not include time spent recursively executing/parsing SQL statements or time spent recursively executing PL/SQL.

System Time Detail

The System Time Detail section allows you to see where the database engine has been spending its time. Both raw time and percentage of total metrics are included.

Location	Oracle Performance Analyst statistics > Response Time Analysis View (Oracle 10G Only) > Active Time

Metric

The metrics shown here include (*note*, this table is largely from Oracle 10g documentation):

Metric	Description
DB Time	Amount of elapsed time (in microseconds) spent performing Database user-level calls. This does not include the time spent on instance background processes such as PMON.
DB CPU	Amount of CPU time (in microseconds) spent on database user-level calls. This does not include the CPU time spent on instance background processes such as PMON.
background cpu time	Amount of CPU time (in microseconds) consumed by database background processes.
sequence load elapsed time	Amount of elapsed time spent getting the next sequence number from the data dictionary. If a sequence is cached, then this is the amount of time spent replenishing the cache when it runs out. No time is charged when a sequence number is found in the cache. For non-cached sequences, some time will be charged for every nextval call.
parse time elapsed	Amount of elapsed time spent parsing SQL statements. It includes both soft and hard parse time.
hard parse elapsed time	Amount of elapsed time spent hard parsing SQL statements.
sql execute elapsed time	Amount of elapsed time SQL statements are executing. Note that for select statements this also includes the amount of time spent performing fetches of query results.
connection management call elapsed time	Amount of elapsed time spent performing session connect and disconnect calls.
failed parse elapsed time	Amount of time spent performing SQL parses which ultimately fail with some parse error.
hard parse (sharing criteria) elapsed time	Amount of elapsed time spent performing SQL hard parses when the hard parse resulted from not being able to share an existing cursor in the SQL cache.
hard parse (bind mismatch) elapsed time	Amount of elapsed time spent performing SQL hard parses when the hard parse resulted from bind type or bind size mismatch with an existing cursor in the SQL cache.
PL/SQL execution elapsed time	Amount of elapsed time spent running the PL/SQL interpreter. This does not include time spent recursively executing/parsing SQL statements or time spent recursively executing the Java VM.
PL/SQL compilation elapsed time	Amount of elapsed time spent running the PL/SQL compiler.
inbound PL/SQL rpc elapsed time	Time inbound PL/SQL remote procedure calls have spent executing. It includes all time spent recursively executing SQL and JAVA, and therefore is not easily related to "PL/SQL execution elapsed time".
Java execution elapsed time	Amount of elapsed time spent running the Java VM. This does not include time spent recursively executing/parsing SQL statements or time spent recursively executing PL/SQL.

Session Time Detail

This section allows you to see where sessions that are currently logged on are spending their time.

Location	Oracle Performance Analyst statistics > Response Time Analysis View (Oracle 10G Only) > Active Time

>

Metrics

The metrics shown here include (*note*, this table is largely from Oracle 10g documentation):

Metric	Description
DB Time	Amount of elapsed time (in microseconds) spent performing Database user-level calls. This does not include the time spent on instance background processes such as PMON.
DB CPU	Amount of CPU time (in microseconds) spent on database user-level calls. This does not include the CPU time spent on instance background processes such as PMON.
background cpu time	Amount of CPU time (in microseconds) consumed by database background processes.
sequence load elapsed time	Amount of elapsed time spent getting the next sequence number from the data dictionary. If a sequence is cached, then this is the amount of time spent replenishing the cache when it runs out. No time is charged when a sequence number is found in the cache. For non-cached sequences, some time will be charged for every nextval call.
parse time elapsed	Amount of elapsed time spent parsing SQL statements. It includes both soft and hard parse time.
hard parse elapsed time	Amount of elapsed time spent hard parsing SQL statements.
sql execute elapsed time	Amount of elapsed time SQL statements are executing. Note that for select statements this also includes the amount of time spent performing fetches of query results.
connection management call elapsed time	Amount of elapsed time spent performing session connect and disconnect calls.
failed parse elapsed time	Amount of time spent performing SQL parses which ultimately fail with some parse error.
hard parse (sharing criteria) elapsed time	Amount of elapsed time spent performing SQL hard parses when the hard parse resulted from not being able to share an existing cursor in the SQL cache.
hard parse (bind mismatch) elapsed time	Amount of elapsed time spent performing SQL hard parses when the hard parse resulted from bind type or bind size mismatch with an existing cursor in the SQL cache.
PL/SQL execution elapsed time	Amount of elapsed time spent running the PL/SQL interpreter. This does not include time spent recursively executing/parsing SQL statements or time spent recursively executing the Java VM.
PL/SQL compilation elapsed time	Amount of elapsed time spent running the PL/SQL compiler.
inbound PL/SQL rpc elapsed time	Time inbound PL/SQL remote procedure calls have spent executing. It includes all time spent recursively executing SQL and JAVA, and therefore is not easily related to "PL/SQL execution elapsed time".
Java execution elapsed time	Amount of elapsed time spent running the Java VM. This does not include time spent recursively executing/parsing SQL statements or time spent recursively executing PL/SQL.

System Bottlenecks

The System Bottlenecks tab of the Response Time Analyst detail view displays the following information:

Bottleneck Summary	Bottleneck Detail	Bottleneck History (One Hour)

Bottleneck Summary

Waits on a system generally occur for three reasons:

1. A process waits because it has no work to do.
2. A process waits because a requested resource is not available.

3 A process waits for Oracle to perform a prerequisite task for its given operation.

Idle waits (processes waiting because they have no work) are not normally a problem, however the other two wait causes are the ones worth your time and investigation. From a global database level, there are many different types of waits and sources of contention. The Bottleneck Summary section presents the broad categories of waits (called classes) so you can see if, for example, User I/O is causing high wait times. The percentage of each wait class is shown in the pie chart.

Location	Oracle Performance Analyst statistics > Response Time Analysis View (Oracle 10G Only) > System Bottlenecks
Metrics	Details can be viewed for the wait classes by drilling down into the System Waits view.

Bottleneck Detail

Waits on a system generally occur for one of three reasons:

1 A process waits because it has no work to do.

2 A process waits because a requested resource is not available.

3 A process waits for Oracle to perform a prerequisite task for its given operation.

Idle waits (processes waiting because they have no work) are not normally a problem, however the other two wait causes are the ones worth your time and investigation. From a global database level, there are many different types of waits and sources of contention. The Bottleneck Detail section presents the broad categories of waits (called classes) so you can see if, for example, User I/O is causing high wait times. The wait class, total waits, percentage of total waits, wait time (in seconds) and percentage of wait time is shown. Note that all these metrics are cumulative since the instance was started.

Location	Oracle Performance Analyst statistics > Response Time Analysis View (Oracle 10G Only) > System Bottlenecks

Metrics

Details can be viewed for the wait classes by drilling down into the System Waits view.

Bottleneck History (One Hour)

Waits on a system generally occur for three reasons:

1 A process waits because it has no work to do.

2 A process waits because a requested resource is not available.

3 A process waits for Oracle to perform a prerequisite task for its given operation.

Idle waits (processes waiting because they have no work) are not normally a problem, however the other two wait causes are the ones worth your time and investigation. From a global database level, there are many different types of waits and sources of contention.

The Bottleneck History section allows you to view wait counts, wait time, or a count of user sessions waiting over the past hour for all the specific wait classes (User I/O, System I/O, etc.)

Location	Oracle Performance Analyst statistics > Response Time Analysis View (Oracle 10G Only) > System Bottlenecks

>

Session Bottlenecks

The Session Bottlenecks tab of the Response Time Analysis detail view displays the following information:

Session Bottleneck Summary	Session Bottleneck Detail	Session Wait History (One Hour)

Session Bottleneck Summary

Waits on a system generally occur for three reasons:

1 A process waits because it has no work to do.

2 A process waits because a requested resource is not available.

3 A process waits for Oracle to perform a prerequisite task for its given operation.

Idle waits (processes waiting because they have no work) are not normally a problem, however the other two wait causes are the ones worth your time and investigation. From a global database level, there are many different types of waits and sources of contention.

The Session Bottleneck Summary section shows where currently logged on sessions have been spending the largest percentages of their time (over broad wait classes).

Location	Oracle Performance Analyst statistics > Response Time Analysis View (Oracle 10G Only) > Session Bottlenecks
Metrics	To view the actual wait events that comprise the highest percentages of time viewed in each wait class, drill down into the Session Waits view.

Session Bottleneck Detail

Waits on a system generally occur for three reasons:

1 A process waits because it has no work to do.

2 A process waits because a requested resource is not available.

3 A process waits for Oracle to perform a prerequisite task for its given operation.

Idle waits (processes waiting because they have no work) are not normally a problem, however the other two wait causes are the ones worth your time and investigation. From a global database level, there are many different types of waits and sources of contention.

The Session Bottleneck Detail section shows where currently logged on sessions have been spending the largest percentages of their time (over broad wait classes). Each session is shown along with its SID, user name, wait class, total waits, time waited (in seconds), and percent of total time.

Location	Oracle Performance Analyst statistics > Response Time Analysis View (Oracle 10G Only) > Session Bottlenecks
Metrics	To view the actual wait events that comprise the highest wait times/percentages of time viewed in each wait class, drill down into the Session Waits view.

>

Session Wait History (One Hour)

The Session Wait History allows you to look back over the last hour to see when sessions where waiting, how many waits occurred, and the amount of wait time associated during a specific period of time. Analyzing such data over time will give you an idea of how often waits are occurring and their overall impact on system performance.

Location	Oracle Performance Analyst statistics > Response Time Analysis View (Oracle 10G Only) > Session Bottlenecks

Wait Histogram

The Wait Histogram view allows you to view a histogram of wait times and wait counts for a variety of wait events. Using the wait histogram view, you can see, for example, that a db file scattered read event had 20 waits that took an average of 1 second and 10 waits that took an average of 2 seconds, and so on. This allows you to get a better idea of the actual impact of various wait events than viewing broad percentages or averages.

Historical Session Analysis

The Historical Session Analysis tab of the Response Time Analysis detail view shows data for:

- Historical Session Analysis
- Historical Session Analysis Detail - All Waits

Historical Session Analysis

Oracle10g and above allow you to review various historical performance metrics for your database. The Historical Analysis graph allows you to input a time period and review wait times to see when your server experienced high volumes of wait activity. You can choose to look at global wait classes or specific waits events for wait classes.

You can also select a specific time period in the graph by drawing a box around part of the graph with your mouse. The graph will change to match your selection.

Location	Oracle Performance Analyst statistics > Response Time Analysis View (Oracle 10G Only) > Historical Session Analysis

Historical Session Analysis Detail - All Waits

Oracle10g and above allows you to review various historical performance metrics for your database. The Historical Analysis Detail section allows you to review specific details regarding times when your database server experienced unacceptable wait times. You can review historical response time data for sessions, SQL statements, files, and objects.

This section allows you, for example, to see what wait events long running SQL statements experienced and what objects caused excessive wait times.

Location	Oracle Performance Analyst statistics > Response Time Analysis View (Oracle 10G Only) > Historical Session Analysis

>

Session Details

The following tabbed pages are available on the Session Details page:

Current Work Tab	Locks Tab	Overview Tab
SQL Tab	Waits Tab	

Overview Tab

Several sections are available on the Overview tab of the Session Details page. The table below describes the sections:

Section	Description
Memory Usage	The Memory Usage section defines the session's total memory usage against all session memory usage on the database.
I/O Activity	The I/O section defines the session's total I/O consumption against all I/O activity on the database. **NOTE** Bugs in some versions of the Oracle data dictionary cause negative numbers to display if the database has experienced large amounts of I/O activity.
CPU Activity	The CPU Activity section defines the session's total CPU consumption against all CPU usage on the database. **NOTE** Bugs in some versions of the Oracle data dictionary cause the total CPU usage for a session to not be calculated correctly.
Statistical Detail	The Statistical Detail section details every performance metric available for the session. **NOTE** Each measure continue to increment as long as the session stays connected to Oracle.

Current Work Tab

The Current Work tab of the Session Details page includes the following sections:

- Open Transactions

The following statistic, available on this tab, duplicates a statistic available on the Key User Analysis pane

- Open Cursors

Open Transactions

This Open Transactions section displays details regarding any open transaction that the session currently has. The table below describes the information available in this section:

Column	Description
SQL Text	The SQL that contains the transaction.
Start Time	The time the statement was issued.
Logical I/O	The amount of memory I/O the transaction has used.
Physical I/O	The amount of disk I/O the transaction has used.
Gets	A subset of logical I/O.
Changes	Any data modifications made by the transaction.

>

Location	Oracle Performance Analyst statistics > Session Details > Current Work Tab

Metrics

Seeing a long running, persistent transaction might indicate a session caught in a loop or other coding snafu. Watching a transaction sit idle could clue you into a lock contention problem that exists.

Waits Tab

The Waits tab of the Session Details page includes the following sections:

Time Summary (Oracle 10g only)	Last Ten Waits (Oracle 10g only)	

The following statistics, available on this tab, are duplicates of statistics available on the System Waits Tab of the Users page:

Historical Waits	Current Waits	

Time Summary (Oracle 10g only)

The Time Summary section shows you where the session has been spending its time. Using the Time Summary section, you can easily determine if a session has been bottlenecked by SQL calls, PL/SQL work, or other activities.

Location	Oracle Performance Analyst statistics > Session Details > Waits Tab

Last Ten Waits (Oracle 10g only)

The Last Ten Waits section shows you, in order, the last 10 wait events experienced by the session, along with other important metrics like the wait time, and so forth.

Location	Oracle Performance Analyst statistics > Session Details > Waits Tab

Locks Tab

The following statistics, available on this tab, are duplicates of statistics availale on the Locks/Blocking Locks Tab of the Users page:

All Locks	Blocking Locks	

SQL Tab

The SQL tab of the Session Details page includes the following sections:

Current SQL	Current SQL Statistics	Other SQL

Current SQL

The Current SQL section displays the most current SQL issued by the session.

NOTE: You can run an EXPLAIN PLAN on the SQL or (if you have Embarcadero SQL Tuner installed) tune the SQL for better performance.

Location	Oracle Performance Analyst statistics > Session Details > SQL Tab

Current SQL Statistics

The Current SQL Statistics section displays the performance execution metrics for the current SQL statement in graphical format. Physical and logical I/O is displayed along with rows processed and number of executions.

Location	Oracle Performance Analyst statistics > Session Details > SQL Tab

Other SQL

The Other SQL section displays other SQL that has been parsed by the Oracle username. The table below describes the information available in this section:

Column	Description
SQL Text	A snippet of the SQL code issued.
Disk Reads	The number of total disk reads for the statement.
Buffer Gets	The number of total buffer gets for the statement.
Rows Processed	The number of total rows processed for the statement.
Executions	The number of times the statement has been executed.
Sorts	The number of sorts for the statement.

NOTE: Bugs in some versions of Oracle will cause negative numbers to display for various execution metrics.

>

Performance Analyst for Microsoft SQL Server Tutorial

Embarcadero Performance Analyst provides key insight into the complex world of database performance. Running inside the Embarcadero DBArtisan console, or standalone in a Web browser, Performance Analyst helps you pinpoint performance problems that threaten a database's availability or response times. The powerful analysis facilities will satisfy advanced experts, yet novices will find Performance Analyst's interface easy to use and the expert help a welcome aid to troubleshooting difficult performance issues.

Proceed to About This Performance Analyst for Microsoft Tutorial

About This Performance Analyst for Microsoft Tutorial

This tutorial is intended to help you get started with Embarcadero Performance Analyst.

After completing this tutorial, you will have the foundation you need to explore the many features and benefits of Performance Analyst. You will have learned how to accurately monitor your database using the most popular performance methodologies in use today. You will know how to drill down into the many layers of database performance and also be able to locate performance issues in your database in a fast and effective manner. Finally, you'll understand how to fix performance problems using the integration that exists between Performance Analyst and Embarcadero DBArtisan.

This tutorial is divided into 8 sections.

Session 1: Getting Started with Performance Analyst

Session 2: Using Ratio-based Performance Analysis

Session 3: Using Bottleneck/Wait-based Performance Analysis

Session 4: Using Workload Performance Analysis

Session 5: Using SQL Performance Analysis

Session 6: Using Operating System Performance Analysis

Session 7: Identifying and Getting Help with Performance Problems

Session 8: Fixing Identified Problems Using Performance Analyst for Microsoft

You can use this basic tutorial as a roadmap of product highlights.

After starting this tutorial, you can select **Help** from the menu to find many additional resources that complement and build on many of the activities shown in this brief guide.

For more information, see Additional Evaluation Resources.

Session 1: Getting Started with Performance Analyst for Microsoft

Download and Install

You can obtain the latest version of Performance Analyst by downloading the DBArtisan software from the Embarcadero Web site at http://www.embarcadero.com/downloads/downloaddbartisan.asp. Performance Analyst runs fully contained within DBArtisan so you will need the DBArtisan product to use Performance Analyst (although Performance Analyst can also run in standalone fashion within a Microsoft Internet Explorer Web browser).

Click **Download**, and follow the steps indicated. Save the file on your computer and then double-click to launch the self-extracting file that will guide you through the installation process.

When you first install an evaluation copy of DBArtisan, you can use the tool for 14 days. After that time, a permanent license is needed.

Proceed to Overview of Performance Analyst for Microsoft

Overview of Performance Analyst for Microsoft

When DBArtisan is connected to a Microsoft SQL Server database, you can open Performance Analyst two ways:

1 On **Utilities** menu, select **Performance Analyst**.

For Performance Analyst to work, several procedures need to be installed by the SA account. When Performance Analyst is first invoked, it will check for the presence of this package and if it has not been installed, the product will prompt you to install it. Once created, any account with access to these procedures may monitor databases using Performance Analyst. No other objects need to be installed on a monitored database.

Proceed to Performance Analyst for Microsoft Product Navigation

Performance Analyst for Microsoft Product Navigation

Navigating through Performance Analyst is easily accomplished through two main mechanisms:

1 The main Performance Analyst toolbar contains buttons that invoke the major performance analysis views within the product as well as frequently used actions like Refresh and Performance Threshold Customizations.

2 Positioning the mouse pointer on any area within Performance Analyst and clicking the right mouse button will bring up a navigation menu that allows you to easily move between any of the main Performance Analyst home pages and drill down views. The right mouse menu also contains context sensitive options like drilling down further into a particular statistic, copying the display to a text file, and more.

Proceed to Session 2: Using Ratio-Based Performance Analysis in Performance Analyst for Microsoft

Session 2: Using Ratio-Based Performance Analysis in Performance Analyst for Microsoft

Ratio-based analysis has been around for many years, and used to be the only technique database administrators used when they were called on to diagnose the cause of a database slowdown. Performance ratios are very good roll-up mechanisms for busy DBAs, making possible the analysis-at-a-glance approach, however ratios must be computed correctly for them to be meaningful. Key ratios in Performance Analyst are calculated using *delta* statistics, which is superior to computing ratios with cumulative values. Delta statistics, which are values derived by using the before and after measures in each statistical refresh done in Performance Analyst, lets you see changes in performance metrics that would otherwise be invisible if cumulative statistics were used. Performance Analyst makes it easy to see the ratios that are computed with delta statistics vs. cumulative. A triangle symbol denotes a delta-based statistic, while square symbols represent a cumulative-based measure.

Ratio-based analysis is available for all performance categories (memory, etc.) where ratios can assist in determining the performance level of a database. For example, if you would like to see memory-specific ratios, on the Performance Analyst toolbar click **Memory**. Performance Analyst displays the Memory home page that showcases all important memory-related ratios.

Proceed to Obtaining Detail Information for a Ratio Using Performance Analyst for Microsoft

Obtaining Detail Information for a Ratio Using Performance Analyst for Microsoft

A ratio gives you a global picture for a particular performance measure. However, you often need more information regarding a statistical measure so you can make an informed decision. Performance Analyst has easy-to-use drill down features that make obtaining detailed information for various ratios a breeze.

For example, to view more information for the Buffer Cache Hit Ratio statistic, do the following:

1 On either the **Main** Home page or the **Memory** Home page, locate the **Buffer Cache Hit Ratio** statistic.

2 Position the pointer over the statistic.

The pointer changes to a hand symbol.

3 Double-click the left mouse button.

Performance Analyst opens the Buffer Cache Tab of the memory drill down view.

OR

4 Right-click the statistic, select **Details**.

OR

5 Select **Navigate to**, **Memory**, **Buffer Cache**.

Performance Analyst opens the Buffer Cache Tab of the memory drill down view.

Proceed to Session 3: Using Bottleneck/Wait-Based Performance Analysis in Performance Analyst for Microsoft

Session 3: Using Bottleneck/Wait-Based Performance Analysis in Performance Analyst for Microsoft

When a Microsoft SQL Server database is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting may mean nothing in the overall scheme of things, or it can be an indicator that a database bottleneck exists. Bottleneck analysis is a valid method of measuring performance because it helps you track where a database and user sessions have been spending their time. If latch contention or heavy table-scan activity has been dragging a database's performance down, you can use bottleneck analysis to confirm the actual root cause.

Performance Analyst uses two broad areas of metrics to form its bottleneck analysis. The first is the Microsoft SQL Server wait event interface, which displays where the database and user sessions have been bottlenecked. But, when using bottleneck analysis you can not rely *only* on the information contained in the wait event views that Microsoft SQL Server provides. For example, an object may attempt to extend into another extent of space in a tablespace and yet, be denied if no such free space exists. Such a failure will not be reflected in any wait event, but still represents a very real bottleneck to the database. For this reason, Performance Analyst also shows non-wait event based bottlenecks that represent a very real threat to a database's availability and performance.

Performance Analyst prominently displays bottleneck analysis on every performance home page. The Main home page, displays a summary of all key bottlenecks threatening the database. You can also see specific bottlenecks for memory, I/O, space, objects, and users by navigating to each home page and viewing the Bottleneck Analysis section for that home page. For example, to see bottlenecks that specifically relate to I/O, do the following

1 On the toolbar, click the **I/O** button.

2 Locate the Bottleneck Analysis section at the lower-left portion of the screen.

The I/O bottleneck analysis section displays the Top I/O Wait Events by wait count and wait time, along with the Hottest Database and Log Files.

Proceed to Obtaining Detail on System Bottlenecks Using Performance Analyst for Microsoft

Obtaining Detail on System Bottlenecks Using Performance Analyst for Microsoft

You can easily obtain more information regarding global bottlenecks that are plaguing a database by drilling down into a particular bottleneck metric. For example, if you would like to see more data regarding the current top system bottlenecks, do the following:

1. Open the **Main** home page.
2. Double-click the **Top System Bottlenecks** area.

   Performance Analyst opens the System Waits view.
3. Or right-click anywhere in Performance Analyst, point to **Users**, and then click **System Waits**.

   Performance Analyst opens the System Waits view.

The System Waits view displays critical details regarding all the wait events that Microsoft SQL Server has recorded. You can easily see what events that occurred the most often as well as the events that have accumulated the most wait time.

> **TIP:** If you would like to see an expanded view of the System waits grid shown on the top of the System Waits view, click the Expand button located at the top right corner of the section. All section views in Performance Analyst can be maximized/minimized using this control.

Proceed to Viewing Detail for Session Bottlenecks Using Performance Analyst for Microsoft

Viewing Detail for Session Bottlenecks Using Performance Analyst for Microsoft

Not only can you obtain detail information for system bottlenecks, but you can also use Performance Analyst to get drill down information regarding session-level bottlenecks. There are several layers to session-level details you can view.

To see more information regarding session-level waits, do the following:

1. Open the **Main** home page.
2. Double-click the **Top Session Bottlenecks** area.

   OR
3. Right-click anywhere in Performance Analyst, point to **Users**, and then click **Session Waits**.

   Performance Analyst opens the Session Waits drill down view.

When you are in the Session Waits drill down view, to drill further into a session's details, do the following:

1. In the **Historical Waits** or **Current Waits** grid, select a session that appears to be experiencing critical waits.
2. Double-click the session or right-click, and then select **Details**.

   Performance Analyst opens a drill down view that shows key statistics for the selected session, including its current work activities, wait events, and SQL statements.

Proceed to Session 4: Using Workload Performance Analysis in Performance Analyst for Microsoft

Session 4: Using Workload Performance Analysis in Performance Analyst for Microsoft

Key ratios give you a global perspective on database activity. Bottleneck analysis gives you insight into things that are holding up user activity and throughput. But another technique is necessary if a database professional is to really get a handle on what's occurring inside a badly performing database.

>

Workload analysis involves the investigation of two critical areas of a database's performance:

- Session resource consumption and activity
- SQL execution

Without looking at these two key performance categories, you can miss a lot of what could be responsible for perceived performance problems. When performance on a database takes a sudden nosedive, it is not uncommon to find one or two sessions that are causing the bulk of the workload. Finding these processes can be easily accomplished by viewing session metadata coupled with resource consumption and statistical execution statistics.

Proceed to Pinpointing Top Resource Consumers Using Performance Analyst for Microsoft

Pinpointing Top Resource Consumers Using Performance Analyst for Microsoft

Performance Analyst identifies top resource consumers in a couple of different ways. On the main home page, Performance Analyst highlights the top resource sessions across I/O, memory usage, and CPU consumption. The percentage used across all statistical categories are displayed so a DBA can immediately pinpoint a session that is using all or most of a particular resource. The top resource sessions are also displayed on the memory, I/O, and other home pages as well with a listing for the top sessions for that category (for example, the top memory users appear on the memory home page, etc.)

Proceed to Getting Workload Details Using Performance Analyst for Microsoft

Getting Workload Details Using Performance Analyst for Microsoft

Performance Analyst offers several details on top resource sessions in addition to the summaries provided on each home page. A detailed and flexible top sessions view lets you view a summary of all the major resource sessions on the database as well as every critical detail regarding each session connected to the database.

Further drill downs are available that highlights specific details regarding each session's activities and submitted SQL. To see detailed information regarding a session identified in one of Performance Analyst's Workload Analysis sections, do the following:

1. In a **Workload Analysis** display, position the pointer over one of the processes.
2. Double-click the process or right-click, and then select **Details**.

If you want like to see all key workload analysis metadata for sessions, do the following:

1. On the toolbar, click the **Top Sessions** button.

   OR

2. Right-click anywhere in Performance Analyst, point to **Users**, and then click **Session Waits**.

Proceed to Session 5: Using SQL Performance Analysis in Performance Analyst for Microsoft

Session 5: Using SQL Performance Analysis in Performance Analyst for Microsoft

SQL-based performance analysis is really a subset of Workload Analysis. Understanding current and historical SQL execution patterns lets you to have the second set of data points necessary to properly perform workload analysis. You may find that optimizing SQL code produces some of the best performance-enhancing boosts available for a database.

Proceed to Uncovering Resource-Intensive SQL Using Performance Analyst for Microsoft

Uncovering Resource-Intensive SQL Using Performance Analyst for Microsoft

Performance Analyst offers plenty of insight into SQL performance and execution patterns. Each one of the home pages depicts the most costly SQL statements as determined by a variety of different metrics for that performance category (most physical, logical I/O, etc.) For example, Performance Analyst might show that a single SQL statement is responsible for almost 60 percent of all the CPU consumption on the database.

Proceed to Getting Details on Top SQL with Performance Analyst for Microsoft

Getting Details on Top SQL with Performance Analyst for Microsoft

Getting details on identified SQL statements is easy in Performance Analyst. For example, if you would like to see further information on a SQL statement identified on the Performance Analyst Home page, do the following:

1. In the **SQL Analysis** section, position the pointer over **Hash Values**.
2. Double-click or right-click, and then select **Details**.

   Performance Analyst opens the Top SQL summary page that shows a breakdown of all top SQL statements along with their execution statistics.

If you would like to see the full SQL text of a SQL statement, do the following

1. In the **Top SQL Details** section, double-click a statement.

   Performance Analyst opens the SQL Details Tab of the Top SQL view to display the entire SQL statement and the execution metrics. While on this pane, you can also perform an EXPLAIN of the SQL statement and examine its access path information.

To automatically tune a problem SQL statement if you have Embarcadero SQL Tuner installed on your workstation, do the following:

1. On the **SQL Details** Tab, select a problem statement, and then click the **Tune SQL** button.

   SQL Tuner opens and you can begin a tuning session.

Proceed to Session 6: Using Operating System Performance Analysis in Performance Analyst for Microsoft

Session 6: Using Operating System Performance Analysis in Performance Analyst for Microsoft

In many scenarios, an optimally tuned database may not perform well because there are constraints imposed by the system where the database is running. These constraints may include processes competing with the database sever for resources (CPU, I/O or Memory), a slow CPU, insufficient I/O devices, and insufficient memory. Performance Analyst allows the DBA to examine the relevant operating system metrics for Windows Platforms (NT /2000/XP/2003) as they relate to SQL Server.

To view the operating system metrics of the machine hosting the SQL Server you are monitoring, do the following:

1. Click the **OS** Performance Analyst toolbar button.

   OR

2. Navigate to the OS Home Page by going to **OS**, and then **Home**.

   NOTE: The first time you invoke the OS feature you will be prompted for an operating system login. Please specify the correct operating system, host name, username & password and protocol.

The Operating System home page has been designed in a similar fashion as the datasource home page. On this page you can see the pertinent operating system metrics displayed in five quadrants broken down into the following categories:

- Resource Usage
- Memory Analysis
- Disk Analysis
- Workload Analysis
- Bottleneck Analysis

Proceed to Obtaining Details on the Operating System Using Performance Analyst for Microsoft

Obtaining Details on the Operating System Using Performance Analyst for Microsoft

You can easily obtain more information regarding operating system problems that are degrading database performance by drilling down into a particular analysis quadrant. For example if you'd like to see more information regarding the processes running on the host machine, do the following:

1. Navigate to the Performance Analyst **Operating System** home page.
2. Double-click the **Top Memory Processes** in the **Workload Analysis** quadrant.

   Performance Analyst opens the Processes Drilldown view.

   OR
3. Right-click anywhere in Performance Analyst and select the **Navigate to**, **OS**, and then **Processes**.

The Processes detail tab shows all the processes on the host machine. From this view you can easily determine which processes are showing high marks in I/O or CPU, for example, by simply double clicking on the column headers and sorting the processes in ascending or descending order.

Proceed to Session 7: Identifying and Getting Help with Performance Problems Using Performance Analyst for Microsoft

Session 7: Identifying and Getting Help with Performance Problems Using Performance Analyst for Microsoft

You do not have to be an advanced expert in performance tuning to spot trouble in a database if you use Performance Analyst. There are a number of mechanisms in the tool that help you quickly locate areas of concern within a database.

Proceed to Using the Alarm Log in Performance Analyst for Microsoft

Using the Alarm Log in Performance Analyst for Microsoft

Performance Analyst has a central and flexible alarm log that lists every performance violation found in all monitored databases. Unlike some products that only show you alarms for a particular database, Performance Analyst lists all alarms for each database you are monitoring in one central log so you can see every problem in one place. To view the Performance Analyst Alarm log, do the following:

1 On the toolbar, click the **Alarms** button.

Once in the Alarm log, you can perform a number of actions including:

- Selecting an alarm and viewing the application page that shows details regarding the metric in violation.
- Filtering the types of alarms shown through the Options dialog.
- Clearing all current alarms.
- Copying or saving the Alarm log into a different format.

To select any of these options, do the following:

1 Position the pointer over any area of the **Alarm Log**.

2 Right-click and In the shortcut menu, select an option.

Proceed to Recognizing Performance Violations Using Performance Analyst for Microsoft

Recognizing Performance Violations Using Performance Analyst for Microsoft

Performance Analyst makes it easy to visually pinpoint performance violations in a database. Any statistic that does not conform to your predefined performance standards is colorized in any view that it appears so you can easily see potential troublespots.

The colors used to highlight a particular performance violation depend on the severity of the problem. Performance Analyst uses several layers of performance standards/threshold violations with each layer being represented by a different color. You can control everything regarding threshold violations, including how thresholds are determined, their color representation in Performance Analyst's display, and more by using the Threshold Template feature of Performance Analyst.

How can Performance Analyst inform you of performance violations in your database if you minimize the product to your desktop and are working on other things? Even if Performance Analyst is not right in front of you, it can still let you know that a problem has been found in your database through the systray alarm icon. When Performance Analyst detects a problem in one of your databases, one of the things the product does is enable an alarm icon in your workstation's systray area that communicates the presence of alarms in your database. So even if the Performance Analyst interface is not in sight, the product can let you know that something in your database requires attention.

Proceed to Creating Threshold Templates for Performance Analyst for Microsoft

Creating Threshold Templates for Performance Analyst for Microsoft

How does Performance Analyst know what to signal as a problem in a database? It uses predefined performance standards that you define and enforce. All of this is accomplished by using Performance Analyst's Threshold templates. A template is a set of performance metrics and their thresholds that can be reused to enforce performance standards across any number of databases. Embarcadero provides a default performance template that contains nearly all of the most-watched performance statistics so you can begin monitoring with confidence right out of the box. However, you can create your own Threshold templates so you only are alerted to performance violations that are important to you and your particular environment.

To view or create Threshold templates, do the following:

1 On the toolbar, click the **Thresholds** button.

Performance Analyst opens the Threshold Template Properties dialog box.

The Threshold Template Properties dialog box lets you select the options to add or edit templates, as well as other maintenance options. For example, you can copy the original Embarcadero Threshold template and modify it to fit your own needs. You can control what statistics are sampled for performance violations as well as the levels of severity you would like reported. Once you are satisfied with your template, you can save it and then apply it to one or more datasources in one action to immediately begin enforcing your performance standards.

Proceed to Getting Help with Performance Problems Using Performance Analyst for Microsoft

Getting Help with Performance Problems Using Performance Analyst for Microsoft

As databases become more complex, so do the problems that plague them. Knowing what to do for a particular performance problem can be hard in some cases, but Performance Analyst helps by providing a complete set of expert help that you can use as a guide when performance problems occur. To obtain expert help for any statistic shown in Performance Analyst, do the following:

1 Position the pointer over a statistic that you need help.

2 Right-click and select **Help**.

The Expert Help gives you a general description of the statistic, the range of values that are considered optimal, and some suggestions on how to fix a problem situation. Expert help can also be invoked for any of the Ratio, Bottleneck, SQL, or Workload Analysis sections shown on every home page by clicking on the help icon for that section, which is located at the top right corner of the section.

Proceed to Session 8: Fixing Identified Problems Using Performance Analyst for Microsoft

Session 8: Fixing Identified Problems Using Performance Analyst for Microsoft

Finding performance problems is one thing, but fixing them is another. Many performance monitors point out a particular problem in a database, but then offer no way to actually correct the situation. Such is not the case with Performance Analyst. Because of its tight integration with DBArtisan, certain problems found in Performance Analyst can be visually corrected by DBArtisan.

A Space Problem Correction Example

If an Microsoft SQL Server tablespace is about to run out of free space, Performance Analyst can issue an alarm that a space shortage exists. Once the space problem has been identified, do the following:

1 On the **Space** home page toolbar, click the **Space** button.

2 In the **Space Analyst** grid, right-click a tablespace, and then select **Manage with DBArtisan**.

DBArtisan opens an editor where you can perform corrective actions, such as enabling Autoextend for the tablespace, resizing the tablespace's existing datafile(s), or adding a new datafile for the tablespace.

Microsoft SQL Server Performance Analyst statistics

This section includes expert help for all Microsoft SQL Server statistics in the Embarcadero Performance Analyst views and pages:

Home	Memory	I/O
Space	Users	Database
OS	Top SQL	

Home

The Embarcadero Performance Analyst Home page lets you review availability and overall performance of all monitored databases from a single window. Statistics on the Home view are organized into the following categories:

Key Ratio Analysis pane	Bottleneck Analysis pane	SQL Analysis pane
Storage Analysis pane	Workload Analysis pane	

Key Ratio Analysis pane

Database performance analysts typically use one of two methods for examining the performance levels of a database - ratio-based or wait/bottleneck-based. Ratio-based analysis involves examining a number of key database ratios that can be used to indicate how well a database is running. Performance ratios serve as very good roll-up mechanisms for busy DBA's to use for at-a-glance performance analysis. Many DBAs have large database farms to contend with and cannot spend time checking detailed wait-based analysis outputs for each and every database they oversee. Succinctly presented performance ratios can assist in such situations by giving DBAs a few solid indicators that can be quickly scanned to see if any database needs immediate attention.

While there are certainly many opinions as to what rules to follow, there are some standards that should always be adhered to. To start with, many of the formulas that make up ratio-based analysis must be derived from delta measurements instead of cumulative statistics. Many of the global ratios that a DBA will examine come from the master.sysperfinfo performance table. This table maintains a count of various statistics for the server and particular databases since the server was brought up. For servers that are kept up for long periods of time, these values can grow quite large and will impact how a particular ratio that a DBA might be looking at is interpreted. However, if delta statistics are used (taking, for a specified sampling period, the before and after counts of each statistic that make up a ratio), then an accurate and current portrayal of the various ratios can be had.

A final thing to remember about using ratio-based analysis is that, while there are a number of rules of thumb that can be used as starting points to begin the evaluation of database performance, DBAs must determine each SQL Server's individual 'personality' with respect to the key performance ratios. Some hard and fast rules simply do not apply to every database. The danger in using blanket ratio standards is that they can lead the DBA to haphazardly take action, which can at times contribute nothing to the situation, and sometimes even degrade performance.

The following ratios are used on the Performance Analyst Home page to succinctly communicate the general overall performance levels of the monitored database. The statistics available in this pane, are duplicates of statistics available on the Key Ratio Analysis pane of the Home page:

Buffer Cache Hit Ratio

Data read from memory will produce end-user response times many times faster than when that same data is read from disk. Keeping physical I/Os to an absolute minimum is one of the purposes of the SQL Server buffer/data cache.

>

The buffer cache hit ratio is a terrific indicator of how often user requests for data are satisfied through memory vs. being physically read from disk.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Key Ratio Analysis pane Microsoft SQL Server Performance Analyst statistics > Memory > Buffer Cache Tab
Metrics	To help ensure excellent performance, you want to keep your cache hit ratio in the neighborhood of 90% or higher. However, you should be aware that every server has its own 'personality' and might exhibit excellent performance with below average readings for the cache hit ratio. You should also be aware that excessive logical I/O activity can produce a very high cache hit ratio while actually degrading overall database performance. Consistently viewed low readings (60% or lower) might require tuning attention on the part of the DBA.
Correction	Ensure SQL Server is configured to use as much physical memory as possible by checking the Max Server memory configuration option. Also, consider increasing your SQL Server Min. Memory parameter to allocate more memory to SQL Server. (Note that to obtain optimal values for these parameters, an option is to install more physical RAM to your server.) Check for any large objects that are pinned in memory that could possibly be removed. Often a user process is taking a large amount of memory due to an inordinate amount of I/O.

Procedure Plan Hit Ratio

The SQL Server procedure cache is used to hold the execution plans for all Transact-SQL statements currently executing in the server. When a user executes a Transact-SQL statement, SQL Server looks in the procedure cache for a query plan to use.

This statistic is the percentage of query plan requests generated by stored procedures that are found in the procedure cache area. The percentage of times that a statement's plan and definition can be referenced in memory, the better the procedure execution time.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Key Ratio Analysis pane Microsoft SQL Server Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane
Metrics	A high procedure cache hit rate is a desirable thing. You should strive for a hit ratio between 95-100%, with 95% being a good performance benchmark for code reference. Note that when a database is first started, the procedure cache hit rate will not be at an optimal level because all code being used will be relatively new, and as such, must be read in from disk and placed into the cache. If, however, after a solid hour or two of steady database time, the procedure cache hit rate has not increased to desirable levels, you should look into the possibility of increasing the amount of memory allocated to the cache.
Correction	First, ensure SQL Server is configured to use as much physical memory as possible by checking the Max Server Memory configuration option. Also, consider increasing your SQL Server Min Memory parameter to allocate more memory to SQL Server. (Note that to obtain optimal values for these parameters, an option is to install more physical RAM to your server.) Check for any large objects that are pinned in memory that could possibly be removed.

Ad Hoc SQL Hit Ratio

When an ad hoc SQL statement is issued, the query plan is then stored in the SQL Server procedure cache area. If the identical ad hoc statement is launched in the future, SQL Server uses the query plan already stored in the procedure cache if it is still there. The Ad Hoc SQL Hit Ratio statistic defines the percentage of times that a query plan for an ad hoc SQL statement is found in the procedure cache.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Key Ratio Analysis pane Microsoft SQL Server Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane

Metrics	A high ad hoc hit rate is desirable, but is harder to maintain at a high level than something like a procedure cache hit rate. Therefore, an 80% or greater ad hoc cache hit rate is a good performance benchmark for code reference. Note that when a database is first started, the ad hoc cache hit rate will not be at an optimal level because all code being used will be relatively new, and as such, must be read in from disk and placed into the cache. If, however, after a solid hour or two of steady database time, the ad hoc cache hit rate has not increased to desirable levels, you should look into the possibility of increasing the amount of memory allocated to the cache.
Correction	First, ensure SQL Server is configured to use as much physical memory as possible by checking the Max Server Memory configuration option. Also, consider increasing your SQL Server Min Memory parameter to allocate more memory to SQL Server. (To obtain optimal values for these parameters, an option is to install more physical RAM to your server.) Check for any large objects that are pinned in memory that could possibly be removed.

Log Cache Hit Ratio

The Log Cache Hit Ratio represents the percentage of log cache reads satisfied from the log cache.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Key Ratio Analysis pane Microsoft SQL Server Performance Analyst statistics > Memory > Log Cache Tab
Correction	A low percentage on this statistic is not necessarily a bad sign, as it is possible that the information needed from the log will not be readily available in memory.

CPU Busy

The CPU Busy statistic represents the percentage of time that the CPU has spent working since the last refresh.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Key Ratio Analysis pane
Metrics	Seeing consistently high numbers might signal an overworked server.
Correction	If you consistently see high numbers for CPU busy, then drill down into session activity to see what processes are consuming the most CPU and look at their activity.

I/O Busy

The I/O Busy statistic represents the percentage of time that SQL Server has spent performing input and output operations since the last refresh.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Key Ratio Analysis pane Microsoft SQL Server Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane

Server Idle

The Server Idle statistic represents the percentage of time that SQL Server has spent idle since the last refresh.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Key Ratio Analysis pane

Bottleneck Analysis pane

When SQL Server is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting might mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. And this is where wait-based or bottleneck analysis comes into play. DBAs use this form of performance analysis to determine if perceived bottlenecks in a database are contributing to a performance problem.

Bottleneck analysis is a valid method of measuring performance because it helps a DBA track where a database has been spending its time. If lock contention or heavy table scan activity has been dragging down database performance, a DBA can use bottleneck analysis to confirm the actual root cause. Once one or more wait events or other bottlenecks have been pinpointed as possible performance vampires, the DBA can drill down and oftentimes discover a fair amount of detail about which sessions and objects are causing the problem. The Performance Analyst Home page identifies the top system and session waits that might be currently a cause of lessened performance.

When using bottleneck analysis, you cannot rely only on the information contained in the wait events that SQL Server provides. For example, a database might be suspect or offline. Such a failure will not be reflected in any wait event, but still represents a very real bottleneck to the database. In the same way that a DBA cannot depend on only a few ratios to properly carry out ratio-based performance analysis, an administrator must include other statistical metrics in their overall bottleneck analysis framework to obtain an accurate performance risk assessment. Performance Analyst works to identify bottlenecks in your database that fall outside of pure wait events so you can get a total picture of all stoppages in your system. The Performance Analyst Home page displays space-related bottlenecks as well as other sources of contention/bottlenecks that might be dragging down the overall performance of your system.

The following bottleneck indicators are used on the Performance Analyst Home page to succinctly communicate the general overall performance level of the monitored database:

Current Object Blocks	DB's Low on Space	Deadlocks
Errors in Current Log	Logs Low on Space	Offline Databases
Suspect Databases	Table Lock Escalations	Top System Bottlenecks

Top System Bottlenecks

The Top System Bottlenecks display identifies the top waits that have occurred on SQL Server based on the wait time.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Bottleneck Analysis pane
Correction	Wait events can be hard to interpret at times. If you see a particular event that has caused a lot of wait time, you can review the information in this link (Microsoft Knowledge Base Article - 244455) to help understand the cause and potential remedy: http://support.microsoft.com/default.aspx?scid=kb;en-us;Q244455

Current Object Blocks

A single blocking user has the potential to stop work for nearly all other processes on a small system, and can cause major headaches even on large systems. Blocks are most often caused by user processes holding exclusive locks and not releasing them via a proper COMMIT frequency. Unless a process times out via an application timeout mechanism, or the process has specified a timeout period via the SET LOCK_TIMEOUT command, a process waiting for a lock will wait indefinitely.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Bottleneck Analysis pane
Metrics	You should immediately investigate any indicator above zero, before the situation has a chance to mushroom.

Correction	Once discovered, a blocking lock situation can normally be quickly remedied - the DBA issues a KILL against the offending process, which eliminates the user's stranglehold on the objects they were accessing. Then other user processes, almost always, complete in an instant. Discovering the blocked lock situation is made easier by using tools like Performance Analyst, but preventing the blocking lock situation in the first place is where it gets tricky. The DBA can drill down into Users detail and see all current blocking locks, learning exactly which sessions are holding the currently restrictive locks. The culprit of blocking lock scenarios is usually the application design, or the SQL being used within the application itself. Properly coding an application to reference database objects in an efficient order, and then using the right SQL to get the job done, is an art. The key to avoiding lock contention is to process user transactions in the quickest and most efficient manner possible - something not always easy to do. By default, all processes wait indefinitely for locks in SQL Server. You can change this behavior by using the SET LOCK_TIMEOUT command, which limits the number of seconds that a process will wait for a lock before timing out.

Deadlocks

This metric shows the number of deadlocks per second detected by SQL Server. Deadlocks occur when processes cannot proceed because they are waiting on a set of resources held by each other or held by other processes.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Bottleneck Analysis pane
Metrics	Consistently seeing deadlock counts greater than zero will indicate that some user processes are experiencing delays in completing their work. When SQL Server identifies a deadlock, it resolves the situation by choosing the process that can break the deadlock. This process is termed the deadlock victim. SQL Server rolls back the deadlock victim's transaction, and then notifies the process' application by returning an error message. It also cancels the process' request and allows the transactions of the remaining processes to continue. SQL Server always attempts to choose the least expensive thread running the transaction as the deadlock victim.
Correction	Because SQL Server automatically resolves deadlock situations, you should proactively work to prevent them in the first place. The culprit of most blocking lock and deadlock scenarios is usually the application design, or the SQL being used within the application itself. Properly coding an application to reference database objects in an efficient order, and then using the right SQL to get the job done, is an art. The key to avoiding lock contention is to process user transactions in the quickest and most efficient manner possible - something not always easy to do. You can change default deadlock behavior by using the SET DEADLOCK_PRIORITY command, which reprioritizes a process' position in a deadlock situation.

Table Lock Escalations

The Table Lock Escalations statistic represents the number of times locks on a table were escalated.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Bottleneck Analysis pane Microsoft SQL Server Performance Analyst statistics > Users > Home > The following key statistics are used on the Performance Analyst Users home page to succinctly communicate the general session-based activity levels of the monitored database:
Metrics	Many table lock escalations could indicate contention problems. If increasing numbers of table lock escalations are viewed at the same time as blocking or deadlock problems, then the application design might be at fault.

Suspect Databases

The Suspect Databases statistic represents the number of databases SQL Server has marked as suspect. Databases are marked suspect by SQL Server if they fail during automatic recovery, which is performed at server startup. If serious damage is experienced by a database during regular uptime, SQL server will also mark a database suspect.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Bottleneck Analysis pane Microsoft SQL Server Performance Analyst statistics > Database > Home > Bottleneck Analysis pane
Metrics	You should not have any suspect databases found on any production server. You should immediately investigate any non-zero numbers for this statistic.
Correction	The steps to handling a suspect database will vary from one installation to another. However, there are some general guidelines you can use to troubleshoot a suspect database. *Begin by examining the SQL Server error log for clues as to what caused the database to be marked as suspect.* *It is not unusual for a server to run out of physical disk space on drives used by SQL Server. When this happens, recovery for databases can sometimes fail with the end result being SQL Server marking a database as suspect. To remedy this situation, you should free up space on the identified drives or add files to the newly marked suspect database. For SQL Server 2000, this can be accomplished by utilizing the two procedures - sp_add_data_file_recover_suspect_db and sp_add_log_file_recover_suspect_db stored. For version 7.0 of SQL Server, you must use the sp_resetstatus stored procedure to reset the suspect status flag for the database in question, use the alter database command to add new datafiles to the database, and then stop/start the SQL Server.* *Many times, suspect databases are caused by SQL Server not being able to access a database or log file. This happens if a particular physical hard drive has become unavailable, but also can occur if another operating system process has obtained exclusive access to a file. If this scenario proves to be true, once you have ensured that the file(s) are available once again to the operating system, you can use the sp_resetstatus stored procedure to reset the suspect status flag for the database and then stop/start the SQL Server.* If none of these solutions are possible, you will likely have to restore your database using the last full and transaction log backups.

Offline Databases

The Offline Databases statistic represents the number of databases SQL Server has offline, meaning that no database modifications can occur.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Bottleneck Analysis pane Microsoft SQL Server Performance Analyst statistics > Database > Home > Bottleneck Analysis pane
Metrics	You should not have any offline databases found on any production server. You should immediately investigate any non-zero numbers for this statistic.
Correction	Should an offline database be found by Performance Analyst, you can easily place it back online by utilizing either the sp_dboption stored procedure or the alter database command.

Errors in Current Log

SQL Server records various system events in its system or error log. The majority of messages placed into the log are informational in nature, however since some contain critical messages, you should immediately review them. Performance Analyst indicates the number of actual error messages in the most recent error log so you know if there are potential events that require your attention.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Bottleneck Analysis pane
Correction	Drill down into the Error Log Tab within Performance Analyst to view the actual error messages.

DB's Low on Space

A SQL server contains many databases, some of which are devoted to system-level activities (the master and tempdb databases, for example) and others that hold user data. The Databases Low on Space statistic indicates databases that have fallen below a recommended percentage of free space.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Bottleneck Analysis pane Microsoft SQL Server Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	This might or might not be a problem. Some databases are not dynamic in nature (meaning they are not expected to grow in size) and are sized so very little free space is present. However, growing databases are another situation and might require DBA intervention if their free space amounts run low.
Correction	If the percent used amount of a database is approaching problematic levels, there are three ways a DBA can rectify the situation: *The DBA can resize the current file(s) used by the database via an ALTER DATABASE … MODIFY FILE command* *The DBA can add a new file to the database via the ALTER DATABASE … ADD FILE command.* *The DBA can modify the file(s) used by the database to automatically grow by using the ALTER DATABASE … MODIFY FILE … FILEGROWTH command. You should also ensure that the MAXSIZE setting for each file is set appropriately.* You should also ensure that enough physical space exists on the server to accommodate additional database space.

Logs Low on Space

Each database in SQL Server has a transaction log, which is a serial record of all modifications that have occurred in a database as well as the transactions that caused each change. The Logs Low on Space statistic indicates transaction logs that have fallen below a recommended percentage of free space.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Bottleneck Analysis pane Microsoft SQL Server Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	If any log's used space exceeds the Performance Analyst recommended thresholds, then the DBA should take action to ensure that the log does not run out of available free space.
Correction	There are several things a DBA can do to ensure that a database's log does not run out of available free space: *First, most transactional-oriented databases should have their logs assigned to a separate physical drive than the database. This prevents competition for space between the log and the database itself, allows the log to be monitored for space more effectively, and improves performance.* *If the database is not critical in nature, you can set the truncate log on checkpoint option trunc log on chkpt, which will eliminate any non-active space in the log when a database checkpoint occurs.* *Critical databases needing higher levels of recovery should have schedules established that regular perform transaction log dumps. Doing so ensures better recovery scenarios as well as a reduced risk of the transaction log running out of space.* *If a critical transaction log becomes full, it might be impossible to use standard procedures to dump transactions and reclaim space. The dump operation will likely have to incorporate the no log or truncate only options.* *If a transaction log continuously approaches dangerously low levels of free space, then the DBA should allow the underlying file(s) of the log to automatically grow to meet the demand. This can be accomplished by using the ALTER DATABASE … MODIFY FILE … FILEGROWTH command. You should also ensure that the MAXSIZE setting for each file is set appropriately.* *The DBA should also be on the lookout for large load or data modification operations that do not make use of prudently timed commit points. A single, large transaction has the ability to overwhelm any transaction log since only non-active space in the transaction log is removed from log dumps or truncation operations.*

SQL Analysis pane

A lot of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured server down in terms of end-user response times.

Before you can identify problem SQL in your database, you have to ask the question - "What is bad SQL?" What criteria do you use when you begin the hunt for problem SQL in your critical systems? Understand that even seasoned experts disagree on what constitutes efficient and inefficient SQL; so there is no way to sufficiently answer this question to every SQL Server professional's satisfaction. The table below lists some general criteria you can use when evaluating the output from various database monitors or personal diagnostic scripts:

Criteria	Description
Reads	This is often used as the major statistic in terms of identifying good vs. bad SQL, this is a measure of how many disk reads the query caused to satisfy the user's request. While you certainly want to control disk I/O where possible, it is important that you not focus solely on physical I/O as the single benchmark of inefficient SQL. Make no mistake, disk access is slower than memory access and also consumes processing time making the physical to logical transition, but you need to look at the entire I/O picture of a SQL statement, which includes looking at a statements' logical I/O as well.
Writes	These can be caused by DML activity (INSERTs, etc.), but writes can also be indicative of heavy disk sort activity.
CPU	This is how much CPU time the query took to parse, execute, and fetch the data needed to satisfy the query.
Elapsed Time	This is how much time the query took to parse, execute, and fetch the data needed to satisfy the query. It should not include the network time needed to make the round trip from the requesting client workstation to the database server.

There are other criteria that you can examine like sort activity or access plan statistics (that show items like Cartesian joins and the like), but more often than not, these measures are reflected in the criteria listed above.

The SQL Analysis section provides you with a quick overview of the percentage of Reads, Writes, CPU, and Elapsed Time the most expensive SQL statements tracked by Performance Analyst have used. Depending on the page you are on, you might just see information regarding physical I/O activity or DML statements. For example, you might see that a SQL statement has caused 60% of all physical reads on the system, or that a procedure is responsible for 90% of all the CPU usage. To get information regarding the actual SQL text or stored procedure calls, drill down into the Top SQL details views.

NOTE: SQL Profiling is turned on by default in Performance Analyst so you can automatically collect SQL traffic. If you do not want to collect SQL execution information, use the options inside of Performance Analyst to disable SQL profiling.

Metrics

When you begin to look for inefficient SQL in a database, there are two primary questions you should answer:

- What has been the worst SQL that has historically been run in my database?
- What is the worst SQL that is running right now in my database?

When Correction a slow system, you should be on the lookout for any query that shows an execution count that is significantly larger than any other query on the system. It might be that the query is in an inefficient Transaction SQL loop, or other problematic programming construct. Only by bringing the query to the attention of the application developers will you know if the query is being mishandled from a programming standpoint.

There is the possibility that the SQL statement just is not tuned well. To determine that, you can drill down further into the Top SQL view and begin working with the query through EXPLAIN plan analysis and other techniques inside Embarcadero's DBArtisan.

Storage Analysis pane

While DBAs focus on memory settings and tuning SQL, they frequently forget how dangerous and insidious storage problems can be. This is not a good mindset because storage headaches can play a major role in wrecking an otherwise well-running database.

Storage problems generally take one of two forms:

- The ‘hit-the-wall’ variety that can bring things to a complete standstill.
- The ‘performance vampire’ kind that slowly drains the performance of a database over time.

Storage problems can bring the curtain down on a database very quickly, as in the case of a transaction log running out of free space. But storage problems can also silently work behind the scenes to slowly, but surely rob a database of its performance. For example, a hub table in a busy database might be accessed very quickly when an application is first given life, but over time, if it develops a heavy forwarded row problem, it can cause things to run very differently.

The Storage Analysis section of the Performance Analyst Home page displays the Total Used/Total Free Space of all tablespaces in the database. It also provides a count of the number of databases, files, and file groups that exist on the server. For more detailed information on database space, you can use the Performance Analyst Space details.

Total Used/Total Free Space

These statistics represent the total used and free space available in all databases on the server. Although good to know, a more detailed listing by database is needed to determine where any actual space shortages exist on the server. This information can be viewed in the Performance Analyst Space.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Storage Analysis pane
Metrics	If any one database or transaction log begins to approach the Performance Analyst recommended thresholds (and the objects contained within it are dynamic and growing as opposed to static), then the DBA should take action to prevent any future space allocation errors.

Correction	Here are a few things a DBA can do to prevent free space shortages: For Databases: *The DBA can resize the current file(s) used by the database via an ALTER DATABASE … MODIFY FILE command.* *The DBA can add a new file to the database via the ALTER DATABASE … ADD FILE command.* *The DBA can modify the file(s) used by the database to automatically grow by using the ALTER DATABASE … MODIFY FILE … FILEGROWTH command. You should also ensure that the MAXSIZE setting for each file is set appropriately.* For Transaction Logs: *First, most transactional-oriented databases should have their logs assigned to a separate physical drive than the database. This prevents competition for space between the log and the database itself, allows the log to be monitored for space more effectively, and improves performance.* *If the database is not critical in nature, you can set the truncate log on checkpoint option (trunc log on chkpt), which will eliminate any non-active space in the log when a database checkpoint occurs.* *Critical databases needing higher levels of recovery should have schedules established that regular perform transaction log dumps. Doing so ensures better recovery scenarios as well as a reduced risk of the transaction log running out of space.* *If a critical transaction log becomes full, it might be impossible to use standard procedures to dump transactions and reclaim space. The dump operation will likely have to incorporate the no log or truncate only options.* *If a transaction log continuously approaches dangerously low levels of free space, then the DBA should allow the underlying file(s) of the log to automatically grow to meet the demand. This can be accomplished by using the ALTER DATABASE … MODIFY FILE … FILEGROWTH command. You should also ensure that the MAXSIZE setting for each file is set appropriately.* You should also be on the lookout for large load or data modification operations that do not make use of prudently timed commit points. A single, large transaction has the ability to overwhelm any transaction log since only non-active space in the transaction log is removed from log dumps or truncation operations.

Workload Analysis pane

When experiencing performance slowdowns, some of the first questions you should answer are:

- Who is currently connected to the database?
- What resources are they using?
- What are they currently executing?

The Workload Analysis section of the Performance Analyst Home page provides insight into the leading resource hogs of a server, as well as a count of total active and inactive processes. The Users page lets you easily get detailed information into what each leading session is currently doing.

Workload Analysis statistics include:

- Active User Processes
- Inactive User Processes

Metrics

If any one session appears to be using more than 50% of a total resource (CPU, memory, etc.), then you should drill down into the session to find out what they are currently executing.

>

Active User Processes

The Active User Processes statistic represents the total number of active and open threads reported on the server. This number displays the number of processes actively performing work.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Workload Analysis pane

Inactive User Processes

This metric represents the total number of threads logged on to the server that are idle at the current time.

NOTE: This statistic is also available on the Home page under Users.

Location	Microsoft SQL Server Performance Analyst statistics > Home > Workload Analysis pane
Metrics	A large number of inactive users could indicate user sessions that have mistakenly been left logged on. Since each user thread consumes a portion of memory on the server, to reduce resource usage, you should sever any sessions not needing a connection.
Correction	Drill down into the Session Waits Tab and check sessions that have many seconds idle and/or that have been logged on for very long periods of time, as indicated by the logon time column. After verifying that a session is no longer necessary, you can KILL it.

Memory

In addition to a Home page, the Memory category of SQL Server Performance Analyst includes the following tabbed pages:

Buffer Cache Tab	Log Cache Tab	Procedure Cache Tab

Home

The Home page of the Memory category includes the following sections and statistics:

Key Ratio Analysis pane	Bottleneck Analysis pane	SQL Analysis pane
Memory Analysis pane	Workload Analysis pane	

Key Ratio Analysis pane

Database performance analysts typically use one of two methods for examining the performance levels of a database - ratio-based or wait/bottleneck-based. Ratio-based analysis involves examining a number of key database ratios that can be used to indicate how well a database is running. Performance ratios serve as very good roll-up mechanisms for busy DBA's to use for at-a-glance performance analysis. Many DBAs have large database farms to contend with and cannot spend time checking detailed wait-based analysis outputs for each and every database they oversee. Succinctly presented performance ratios can assist in such situations by giving DBAs a few solid indicators that can be quickly scanned to see if any database needs immediate attention.

>

While there are certainly many opinions as to what rules to follow, there are some standards that should always be adhered to. To start with, many of the formulas that make up ratio-based analysis must be derived from delta measurements instead of cumulative statistics. Many of the global ratios that a DBA will examine come from the master.sysperfinfo performance table. This table maintains a count of various statistics for the server and particular databases since the server was brought up. For servers that are kept up for long periods of time, these values can grow quite large and will impact how a particular ratio that a DBA might be looking at is interpreted. However, if delta statistics are used (taking, for a specified sampling period, the before and after counts of each statistic that make up a ratio), then an accurate and current portrayal of the various ratios can be had.

A final thing to remember about using ratio-based analysis is that, while there are a number of rules of thumb that can be used as starting points to begin the evaluation of database performance, DBAs must determine each SQL Server's individual 'personality' with respect to the key performance ratios. Some hard and fast rules simply do not apply to every database. The danger in using blanket ratio standards is that they can lead the DBA to haphazardly take action, which can at times contribute nothing to the situation, and sometimes even degrade performance.

The following memory ratios are used on the Performance Analyst Home page to succinctly communicate the general overall memory performance levels of the monitored database:

Cache Used	Procedure Cache Used	Used Cache %
Used Procedure Cache %		

The following statistics, available in this pane, are duplicates of statistics available on the Key Ratio Analysis pane of the Home page:

Buffer Cache Hit Ratio	Procedure Plan Hit Ratio	Ad Hoc SQL Hit Ratio
Log Cache Hit Ratio		

Cache Used

The total MB of space that SQL Server is using in the buffer cache. Each instance of SQL Server has its own buffer cache where it stores recently used data pages to reduce physical I/O. The goal is to make the buffer cache large enough to maximize the ratio of logical reads to physical reads, but not so large that excessive memory swapping starts generating physical I/O to the pagefile. (Instances of SQL Server 2000 do this automatically under the default configuration settings.)

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane
Metrics	A percentage consistently remaining close to 100% indicates a deficient amount of memory available to SQL Server.
Correction	First, ensure SQL Server is configured to use as much physical memory as possible by checking the Max Server Memory configuration option. Also, consider increasing your SQL Server Min Memory parameter to allocate more memory to SQL Server. (Note that to obtain optimal values for these parameters, an option is to install more physical RAM to your server.) Check for any large objects that are pinned in memory that could possibly be removed.

Procedure Cache Used

The total MB of space that SQL Server is using in its SQL/procedure cache. Microsoft has begun to transition from the term “procedure cache” to “SQL cache” to define this area of memory. The reason being that in SQL Server's past, this area was devoted exclusively to holding query plans for stored procedures only.

>

The SQL Cache (procedure cache) is the part of the SQL Server memory pool that is used to store execution plans for Transact-SQL batches, stored procedures, and triggers. Execution plans record the steps that SQL Server must take to produce the results specified by the Transact-SQL statements contained in the batches, stored procedures, or triggers.

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane

Used Cache %

The Used Buffer Cache Percent statistic represents the total percentage of memory that is in use in the SQL Server buffer cache. Each instance of SQL Server has its own buffer cache where it stores recently used data pages to reduce physical I/O. The goal is to make the buffer cache large enough to maximize the ratio of logical reads to physical reads, but not so large that excessive memory swapping starts generating physical I/O to the pagefile. (Instances of SQL Server 2000 do this automatically under the default configuration settings.)

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane
Metrics	A percentage used consistently remaining close to 100% indicates a deficient amount of memory available to SQL Server.
Correction	First, ensure SQL Server is configured to use as much physical memory as possible by checking the Max Server Memory configuration option. Also, consider increasing your SQL Server Min Memory parameter to allocate more memory to SQL Server. (Note that to obtain optimal values for these parameters, an option is to install more physical RAM to your server.) Check for any large objects that are pinned in memory that could possibly be removed.

Used Procedure Cache %

The total percentage of memory that SQL Server is using in its SQL/procedure cache. Microsoft has begun to transition from the term “procedure cache” to “SQL cache” to define this area of memory. The reason being that in SQL Server's past, this area was devoted exclusively to holding query plans for stored procedures only.

The SQL Cache is the part of the SQL Server memory pool that is used to store execution plans for Transact-SQL batches, stored procedures, and triggers. Execution plans record the steps that SQL Server must take to produce the results specified by the Transact-SQL statements contained in the batches, stored procedures, or triggers.

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane

Bottleneck Analysis pane

When SQL Server is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting might mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. And this is where wait-based or bottleneck analysis comes into play. DBAs use this form of performance analysis to determine if perceived bottlenecks in a database are contributing to a performance problem.

Bottleneck analysis is a valid method of measuring performance because it helps a DBA track where a database has been spending its time. If lock contention or heavy table scan activity has been dragging down database performance, a DBA can use bottleneck analysis to confirm the actual root cause. Once one or more wait events or other bottlenecks have been pinpointed as possible performance vampires, the DBA can drill down and oftentimes discover a fair amount of detail about which sessions and objects are causing the problem.

>

Memory bottlenecks can cause performance degradation in an otherwise well-running database. Typically, these bottlenecks center around the SQL Server latches, which are lightweight locks used to protect certain resources in memory. To help you identify such problems, the following statistics are presented on the Performance Analyst Memory home page:

Latch Waits/Sec	Total Latch Wait Time	Top Latch Waits

Latch Waits/Sec

This metric represents the number of latches per second that could not be satisfied immediately by SQL Server. Latches are lightweight, mini-locks that are used to protect actions and resources used inside transactions.

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	Unless accompanied by long wait times, latch waits should not pose too much of a performance problem in normal SQL Server installations. If you have a high number of latch waits per second, then it is likely that SQL Server is also sporting a low buffer cache hit ratio and is being forced to perform extra physical I/O.

Total Latch Wait Time

Latches are lightweight, mini-locks that are used to protect actions and resources used inside transactions. The total latch wait time statistic represents the sum total of all latch wait time that has occurred on the system.

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	Unless accompanied by long wait times, latch waits should not pose too much of a performance problem in normal SQL Server installations. If you have a high number of latch waits per second, then it is likely that SQL Server is also sporting a low buffer cache hit ratio and is being forced to perform extra physical I/O.

Top Latch Waits

The Top Latch Waits graph depicts the top latch waits by wait time.

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	Latch waits rarely impact performance, however seeing latch waits with high wait time might indicate an area that needs further investigation.
Correction	Wait events can be hard to interpret. If you see a particular event that has caused a lot of wait time, you can review the information in this Microsoft Knowledge Base Article to help understand the cause and potential remedy: http://support.microsoft.com/default.aspx?scid=kb;en-us;Q244455

SQL Analysis pane

When SQL Server executes a set of SQL (ad hoc, procedure, trigger, etc.), it places the code into its procedure cache. The SQL Analysis section of the Performance Analyst Memory page provides a summary of the total memory pages currently in use by the various SQL types (procedure, ad hoc, etc.)

The SQL Analysis section provides you with a quick overview of the percentage of Reads, Writes, CPU, and Elapsed Time the most expensive SQL statements tracked by Performance Analyst have used. Depending on the page you are on, you might just see information regarding physical I/O activity or DML statements. For example, you might see that a SQL statement has caused 60% of all physical reads on the system, or that a procedure is responsible for 90% of all the CPU usage. To get information regarding the actual SQL text or stored procedure calls, drill down into the Top SQL details views.

NOTE: SQL Profiling is turned on by default in Performance Analyst so you can automatically collect SQL traffic. If you do not want to collect SQL execution information, use the options inside of Performance Analyst to disable SQL profiling.

Memory Analysis pane

SQL Server uses various memory regions to accomplish various tasks (like retain frequently used data, procedures, etc.) and improve overall performance. At start up, SQL Server is allocated memory that will fall between the range of two configurable parameters: Min Memory and Max Memory. These statistics show you the amount of SQL Server memory allocated at present. In a normal SQL Server system, this number fluctuates as SQL Server uses more memory and releases unused memory back to the operating system. SQL Server also works in conjunction with the Windows Memory Manager, deallocating space when it detects that the operating system is unable to satisfy memory requests of other operating system processes. Depending on the version of SQL Server you are using, expect to deal with different memory regions. The table below describes some of the SQL Server memory regions:

Memory Scan	Description
Buffer Cache	The buffer cache is a memory pool of buffer pages into which SQL Server reads data pages for quick access.
Procedure Cache	Beginning with SQL Server version 7.0, you can place other SQL-related code objects in memory for reuse. Once SQL Server parses through and places a set of SQL of program code in memory, response time can be increased for subsequent calls to the same set of SQL or SQL code objects.
Connection Memory	The total amount of dynamic memory the server is using for maintaining connections.
Granted Workspace Memory	The total amount of memory granted to executing processes such as sort, hash, bulk copy operations, and so forth.
Lock Memory	The total amount of dynamic memory the server is using for locks.
Optimizer Memory	The total amount of dynamic memory the server is using for query optimization.
SQL Cache Memory	The total amount of memory used to retrieve stored SQL caches.

Workload Analysis pane

It is not uncommon for one or two users to cause the majority of runtime problems that plague a server. The problem could be a runaway process, an untuned batch procedure, or other user-initiated operation.

If your database server does not have an overabundance of memory, then you should periodically check to see who your heavy memory users are along with the total percentage of memory each takes up. If you see one or two users who have more than 25-50% of the total memory usage, then you should further investigate the sessions to see what activities they are performing.

Buffer Cache Tab

The Buffer Cache Hit Ratio tab includes the following sections:

Page Activity	Total vs. Target Memory	Database Pages

>

Page Life Expectancy		

The following statistic, available on this page, duplicates a statistic on the Key Ratio Analysis pane of the Home page:

- Buffer Cache Hit Ratio

Page Activity

SQL Server uses various memory regions to accomplish various tasks (like retain frequently used data, procedures, etc.) and improve overall performance. At start up, SQL Server is allocated memory that falls between the range of two configurable parameters: Min Memory and Max Memory. The Page Activity chart shows you the amount of SQL Server memory allocated at present. In a normal SQL Server system, this number fluctuates as SQL Server uses more memory and releases unused memory back to the operating system. SQL Server also works in conjunction with the Windows Memory Manager, deallocating space when it detects that the operating system is unable to satisfy memory requests of other operating system processes. Depending on the version of SQL Server you are using, there are several memory regions you should expect to deal with. The table below describes the SQL Server memory regions:

Memory Scan	Description
Default Buffer Cache	The buffer cache is a memory pool of buffer pages into which SQL Server reads data pages for quick access.
Procedure/SQL Cache	Beginning with SQL Server version 7.0, you can place other SQL-related code objects in memory for reuse. Once SQL Server parses through and places a set of SQL of program code in memory, response time can be increased for subsequent calls to the same set of SQL or SQL code objects.
Connection Memory	The total amount of dynamic memory the server is using for maintaining connections.
Lock Memory	The total amount of dynamic memory the server is using for locks.
Optimizer Memory	The total amount of dynamic memory the server is using for query optimization.

The Active and Free Page chart in Performance Analyst displays a count of active memory pages (those being used by SQL Server) and free pages (pages not being used).

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Buffer Cache Tab
Metrics	Seeing consistent high volumes of free pages might indicate that SQL Server is using more memory than necessary.

Total vs. Target Memory

Total memory is the amount of memory that's currently in use by SQL server. Target memory is the amount of memory that could be used. SQL server dynamically adjusts its memory allocations depending either on workload up to the maximum memory configuration set by the DBA or the amount it thinks it should give other (non-SQL) server operations.

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Buffer Cache Tab
Metrics	When total memory approaches or reaches the maximum amount of target (defined) memory, then SQL server is experiencing memory pressure. You can try adding more RAM to the machine or adjusting memory ceiling limits imposed at the SQL server configuration level.

Database Pages

The database pages metric gives the amount of pages in memory devoted to database content compared to all other content.

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Buffer Cache Tab
Metrics	If you see that a large percentage of memory is devoted to database pages, the server is database page centric and not procedure plan centric. This can indicate that too many table scans are taking place, which might be flooding the buffer cache with infrequently used pages. You can determine if this is the case by looking at the page life expectancy metric. Page life expectancy readings of 300 seconds or less can indicate excessive table scan activity.

Page Life Expectancy

The page life expectancy metric indicates how long SQL server believes a page will stay in the buffer cache. Pages served from memory result in much shorter response times than pages read from disk into the cache. It's best if frequently used data is pinned in the buffer cache.

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Buffer Cache Tab
Metrics	Readings of 300 seconds or less frequently mean that too many table scans are occurring and flooding the buffer cache with pages that are used rarely. Check the I/O access patterns on the I/O home page or the User I/O drill-down to confirm if this is the case.

Procedure Cache Tab

The Procedure Cache tab includes the following sections:

Procedure Cache Hit Ratio	Procedure/SQL Cache Detail	Procedure/SQL Cache Summary

Procedure Cache Hit Ratio

The SQL Server procedure cache holds the execution plans for all Transact-SQL statements currently executing in the server. When a user executes a Transact-SQL statement, SQL Server looks in the procedure cache for a query plan to use.

The Procedure Cache Hit Ratio statistic is the percentage of query plan requests generated by stored procedures that are found in the procedure cache area. The greater the percentage of times that a statement's plan and definition can be referenced in memory, the better the procedure execution time.

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Procedure Cache Tab
Metrics	A high procedure cache hit rate is a desirable thing. You should strive for a hit ratio between 95-100%, with 95% being a good performance benchmark for code reference. Note that when a database first starts, the procedure cache hit rate will not be at an optimal level because all code being used will be relatively new, and as such, must be read in from disk and placed into the cache. If, however, after a solid hour or two of steady database time, the procedure cache hit rate does not increase to desirable levels, you should look into the possibility of increasing the amount of memory allocated to the cache.
Correction	First, ensure SQL Server is configured to use as much physical memory as possible by checking the max server memory configuration option. Also, consider increasing your SQL Server min memory parameter to allocate more memory to SQL Server. (Note that to obtain optimal values for these parameters, an option is to install more physical RAM to your server.) Check for any large objects that are pinned in memory that could possibly be removed.

>

Procedure/SQL Cache Summary

The SQL Server procedure/SQL cache holds often-referenced code object definitions and ad-hoc SQL. The Procedure/SQL Cache Summary lists each type of code object, the amount of memory each code type (trigger, ad-hoc SQL, etc.) consumes, and a total count of each type of code object.

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Procedure Cache Tab
Metrics	Seeing a lot of memory devoted to ad-hoc SQL plans might indicate an environment where sessions are submitting a lot of SQL requests. Many SQL Server DBAs like to control code submissions through stored procedures.

Procedure/SQL Cache Detail

Beginning with SQL Server version 7.0, other SQL-related code objects might be placed in memory for reuse. Once SQL Server parses through and places a set of SQL of program code in memory, response time can be increased for subsequent calls to the same set of SQL or SQL code object. The Procedure/SQL Cache Details section displays the top 1000 objects in the procedure/SQL cache. The table below describes the information available in the Procedure/SQL Cache Details section on the Procedure Cache tab of the Memory Detail:

Information	Description
Object Type	The type of object (ad hoc SQL, etc.)
Cache Type	The category of SQL code inside the cache (executable plan, etc.)
Object ID	The numeric ID of the object.
Object Name	The name of the object (if applicable).
Database ID	The numeric ID of the database.
Database	The database where the code originated.
User ID	The numeric ID of the user.
User	The owner of the object or code.
Use Count	The number of uses for the object.
SQL Bytes	The amount of SQL bytes used by the code object.
Size (KB)	The size used by the object in kilobytes.
SQL	The actual SQL statement or code being executed.

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Procedure Cache Tab
Metrics	Seeing many objects of the ad hoc SQL type might indicate an environment where sessions are submitting a lot of SQL requests. Many SQL Server DBAs like to control code submissions through stored procedures.

Log Cache Tab

The Log Cache tab includes the following sections:

Database Log Cache Ratios	Log Cache Details	

The following statistic, available on this page, duplicates a statistic on the Key Ratio Analysis pane of the Home page:

- Log Cache Hit Ratio

>

Database Log Cache Ratios

The Log Cache Hit Ratio statistic represents the percentage of log cache reads satisfied from the log cache. The database log cache ratios display breaks down log cache hit ratios by database. Note that some databases with inactive logs will not display a log cache hit ratio.

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Log Cache Tab
Correction	A low percentage on this statistic is not necessarily a bad sign, as it is possible that the information needed from the log will not be readily available in memory.

Log Cache Details

Before ever writing transactions to disk, the log manager of SQL Server formats everything in memory. This area of memory is known as the log cache. The log writer of SQL Server moves through the log caches when transactions are committed (as well as other events) and flushes each cache out to disk. SQL Server also reads from the cache when log records are needed. The table below describes the information available in the Log Cache Details section on the Log Cache tab of the Memory Detail:

Information	Description
Database	The name of the SQL Server database.
Log Cache Reads	The number of reads from the log cache.
Log Flushes	The number of times data was flushed from the log to disk.
Log Flush Waits	The number of times the log had to wait before flushing data to disk.
Log Flush Wait Time	The amount of time the log waited before flushing data to disk in milliseconds.
Log Growths	The number of times the log had to physically grow in size to meet the need for more space.
Log Shrinks	The number of times the log contracted in physical size.

Location	Microsoft SQL Server Performance Analyst statistics > Memory > Log Cache Tab
Metrics	Seeing high amounts of wait time for log flushes could indicate a bottleneck at the disk level. A log that shows high number of growths likely indicates an undersized log. While automatic growth can alleviate out-of-space conditions, many growth operations can slow down overall operations. It is better to have a properly sized transaction log that allows SQL Server to continually enlarge it in size when needed.
Correction	Consider relocating logs showing high amounts of wait time to faster disks. For logs showing high numbers of growths, permanently enlarge the log(s) via an ALTER DATABASE command that will resize the log files.

I/O

In addition to a Home page, the I/O category of SQL Server Performance Analyst includes the following tabbed pages:

Database I/O Tab	File Tab	System I/O Tab
User I/O Tab		

>

Home

The Home page of the I/O category includes statistics in the following categories:

Bottleneck Analysis pane	I/O Access Patterns	Key System Analysis pane
SQL Analysis pane	Workload Analysis pane	

Key System Analysis pane

When complaints begin to surface about your database's performance, oftentimes the root cause can be traced to one or more issues with I/O. The thing to keep in mind when you begin to monitor the I/O of your database is that you are actually reviewing the success of your physical design model.

All the physical storage characteristics and placements, the table and index designs, and the speed that it all works are on display when I/O is monitored. Because a database's main index of performance is measured by how fast I/O needs are satisfied, it is your responsibility to quickly interrogate SQL Server to determine if a reported database slowdown is I/O related.

The following key I/O values are reported in this section:

Checkpoint Pages	Lazy Writes	Log Cache Reads
Log Flushes	Workload Analysis pane	Page Writes
Page Reads	Read Ahead Pages	

Page Reads

This metric represents that number of physical database page reads that are issued per second by SQL Server.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	Page reads are to be expected, especially after initial server start up. This is because SQL Server must first satisfy requests for data and meta-data by reading information in from physical disk. Numerous page reads can also be expected if the physical server does not contain an adequate amount of memory to hold repetitively requested blocks of information. No hard-and-fast rules exist for how many page reads per second is too much. You can cross-reference this statistic with the physical server disk statistics to see if physical page reads and accompanying physical disk I/O is approaching the server's premium capacity levels. And because logical I/O is always many times faster than physical I/O, you should also evaluate the buffer cache hit ratio to determine overall memory vs. physical read efficiency.
Correction	If you find that the server is becoming overworked from a physical I/O standpoint, there are several courses of action you can take: *Examine index usage to ensure that unnecessary table scans are not occurring.* *Check the physical database design to see if table objects have been over-normalized.* *Ensure that SQL Server is configured to use sufficient amounts of memory. Examine the min server memory (MB) and max server memory (MB) parameters to see if SQL Server is constrained on either end of the memory spectrum.* *Check for large pinned table objects that could be using excessive amounts of space in the buffer cache.* *Last, but not least, investigate the possibility of adding more RAM to the physical server.*

Page Writes

The Page Writes statistic represents the number of physical database page writes issued by SQL Server. Page Writes take place during operations such as checkpoints, lazywriter writes, index creations, and BCP routines.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	Page Writes can give you an idea of overall physical write activity, however there are a number of statistics that pertain specifically to certain write activities like checkpoints, etc., that you can examine to determine the amount of physical writes caused by distinct SQL Server processes. With respect to performance, response times experienced by SQL Server users are normally not impacted by write operations unless the writes are synchronous in nature. These are typically BCPs, database recovery operations, and index creations.

Read Ahead Pages

The Read Ahead Pages statistic represents the number of physical database page read in anticipation of use by SQL Server. If SQL Server senses that database pages are being read in a sequential manner, it will institute a pre-fetch mechanism that moves pages into the buffer cache before they are actually needed by a process.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	If data is accessed sequentially (for example, through the use of a clustered index), the read ahead mechanism of SQL Server can increase performance by needed database pages already in the buffer cache before they are actually requested. However, because the read ahead mechanism is typically triggered by full table or index range scans, if the read ahead pages are actually required to satisfy a user's query, then performance might actually suffer. In these cases, the judicious use of indexes is a better route to take.

Log Flushes

The Log Flushes statistic represents the total number of log pages for all databases written to disk by the log writer process. A log flush occurs when SQL Server writes all changes from the database's log cache out to the database's log files on disk.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	Increasing numbers observed for log flushes should not cause concern unless the I/O subsystem of the server appears overwhelmed. In addition, to minimize I/O contention between a database and its accompanying log, it is wise to place database files and log files on separate disks.
Correction	If you have placed a very active database on the same physical file as its log, you can look into moving the log to a separate physical device by adding new log files to a new drive and subsequently removing the old log files when they are not being used.

Log Cache Reads

The Log Cache Reads statistic represents the reads performed per second through the log manager cache. Before ever writing transactions to disk, the log manager of SQL Server formats them in memory. This area of memory is known as the log cache and only contains log records for SQL Server 2000 and later. The log writer of SQL Server moves through the log caches when transactions are committed (as well as other events) and flushes each cache out to disk.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > Home > Key System Analysis pane

>

Checkpoint Pages

The Checkpoint Pages statistic represents the number of pages flushed to disk per second by a checkpoint or other operation that require all dirty (modified) pages to be flushed.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	Checkpoint operations are used by SQL Server to minimize the amount of work the server must perform when databases are recovered during system startup. Checkpoints periodically write out modified pages that are found in the buffer cache to disk. After, SQL Server records the operation in the log to signify that the operation succeeded. Checkpoints can be explicitly performed by a database owner issuing the checkpoint command. SQL Server also performs checkpoints automatically for databases that have the trunc log on chkpt option set. Large SQL Servers have the potential to generate lots of checkpoint write operations. Although SQL Server will do what it can to minimize checkpoint activity, you can also set the recovery interval server parameter to influence how often checkpoints should run.
Correction	If you believe excessive checkpoint activity is occurring, you can take the following steps: *Set the recovery interval server parameter to a larger value with sp_configure.* *Restart SQL Server so the change will take affect or use the RECONFIGURE option to make the change immediately.*

Lazy Writes

The lazy writer is a SQL Server system process that flushes out batches of dirty and aged buffers (buffers that contain changes that must be written back to disk before the buffer can be reused for a different page) and make them available to user processes. The Lazy Writes statistic represents the number of buffers written by the buffer manager's lazy writer process.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	High values might indicate a thrashing situation with data scanned into the buffer cache (but referenced seldom) and then immediately moved back out.

Bottleneck Analysis pane

When SQL Server is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting might mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. And this is where wait-based or bottleneck analysis comes into play. DBAs use this form of performance analysis to determine if perceived bottlenecks in a database are contributing to a performance problem.

Bottleneck analysis is a valid method of measuring performance: it helps you track where a database has been spending its time. The Performance Analyst I/O home page identifies the top I/O related waits that might be currently causing performance issues. Also included are the current hottest database/log files and a summary of I/O activity.

Bottleneck Analysis includes the following information:

I/O Errors	Hottest Database/Log Files	Total I/O
Top I/O Waits		

The following statistic, available on this pane, duplicates a statistic on the Key Ratio Analysis pane of the Home page:

- I/O Busy

I/O Errors

I/O Error Rate reflects total number of I/O errors (errors during read and write operations) encountered by the server since the last refresh inside Performance Analyst Center. I/O Error Rate is a percentage based on Total I/O (the sum the physical reads and writes).

Location	Microsoft SQL Server Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Metrics	You should observe few, if any errors.
Correction	If you notice any errors, you should check the SQL Server error log for details.

Hottest Database/Log Files

The Hottest Database/Log Files display presents the most active database files as indicated by physical I/O activity (reads and writes).

Location	Microsoft SQL Server Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Metrics	Understanding the storage-level hot spots of a database is beneficial for a couple of reasons. *First, you can get a feel for overworked physical disks by viewing I/O statistics at the tablespace and datafile levels. If a particular disk or set of disks is under too much strain, you attempt to relocate the databases to other less-used devices or create new file groups on different disks and move hot objects to them (assuming, of course, you have extra disks available).* *Second, if you have used standard DBA practice and placed indexes in their own file group, you can view the I/O statistics for that file group and see if the indexes are actually being used.*
Correction	Some areas to consider when viewing the hottest database files display are: *Seeing a lot of activity in the master database and datafiles might indicate a lot of recursive calls.* *The tempdb database (devoted to sort activity) shows higher volumes of physical I/O which could indicate a problem with excessive disk sorts.* *You might want to quickly review all the physical I/O for each drive/file system and get a idea of the overworked disks on your server. If you have under-utilized disk drives (with their own controllers), you should consider relocating some databases or filegroups that exhibit high I/O characteristics to those drives.* *To minimize contention at the physical disk level, and improve performance overall, it is generally recommended that a SQL Server database have its database and log files physically separated onto different drives.*

Total I/O

The Total I/O statistic represents the total number of physical reads and writes.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane

Top I/O Waits

The Top I/O Waits statistics rank the top I/O related waits by wait time.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Correction	Wait events can be hard to interpret at times. If you see a particular event that has caused a lot of wait time, you can review the information in this link (Microsoft Knowledge Base Article - 244455) to help understand the cause and potential remedy: http://support.microsoft.com/default.aspx?scid=kb;en-us;Q244455

SQL Analysis pane

A lot of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured server down in terms of end-user response times.

Before you can identify problem SQL in your database, you have to ask the question - "What is bad SQL?" What criteria do you use when you begin the hunt for problem SQL in your critical systems? Understand that even the seasoned experts disagree on what constitutes efficient and inefficient SQL; so there is no way to sufficiently answer this question to every SQL Server professional's satisfaction. The table below lists some general criteria you can use when evaluating the output from various database monitors or personal diagnostic scripts:

Criteria	Description
Reads (Physical)	This is often used as the major statistic in terms of identifying good vs. bad SQL, this is a measure of how many disk reads the query caused to satisfy the user's request. While you certainly want to control disk I/O where possible, it is important that you not focus solely on physical I/O as the single benchmark of inefficient SQL. Make no mistake, disk access is slower than memory access and also consumes processing time making the physical to logical transition, but you need to look at the entire I/O picture of a SQL statement, which includes looking at a statements' logical I/O as well.
Writes	These can be caused by DML activity (INSERTs, etc.), but writes can also be indicative of heavy disk sort activity.

There are other criteria that you can examine like sort activity or access plan statistics (that show items like Cartesian joins and the like), but more often than not, these measures are reflected in the criteria listed above.

The SQL Analysis section provides you with a quick overview of the percentage of Reads, Writes, CPU, and Elapsed Time the most expensive SQL statements tracked by Performance Analyst have used. Depending on the page you are on, you might just see information regarding physical I/O activity or DML statements. For example, you might see that a SQL statement has caused 60% of all physical reads on the system, or that a procedure is responsible for 90% of all the CPU usage. To get information regarding the actual SQL text or stored procedure calls, drill down into the Top SQL details views.

NOTE: SQL Profiling is turned on by default in Performance Analyst so you can automatically collect SQL traffic. If you do not want to collect SQL execution information, use the options inside of Performance Analyst to disable SQL profiling.

Metrics

When you begin to look for inefficient SQL in a database, there are two primary questions you should answer:

- What has been the worst SQL that has historically been run in my database?
- What is the worst SQL that is running right now in my database?

When Correction a slow system, you should be on the lookout for any query that shows an execution count that is significantly larger than any other query on the system. It might be that the query is in an inefficient Transaction SQL loop, or other problematic programming construct. Only by bringing the query to the attention of the application developers will you know if the query is being mishandled from a programming standpoint.

There is the possibility that the SQL statement just is not tuned well. To determine that, you can drill down further into the Performance Analyst Top SQL view and begin working with the query through EXPLAIN plan analysis and other techniques inside Embarcadero's DBArtisan.

I/O Access Patterns

There are many different routes that SQL Server can take to access various objects contained in a database. The I/O Access Patterns display summarizes the methods SQL Server is using to satisfy end-user requests for data.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > Home > SQL Analysis pane

Metrics

Depending on the version of SQL Server you are using, expect to deal with different access styles. The table below describes some of the SQL Server access styles:

Access Style	Description
Full Scans	Full Scans are the total number of full table or index scans. Full scans occur if a table is inadequately indexed or if SQL Server needs to access all rows in a table or index to satisfy a query. Other operations that can cause full scans include UPDATE STATISTICS calls. Avoid unnecessary scans on large tables is something to avoid, and can be a signal to you as a DBA to investigate the use of more indexes and to review SQL access through EXPLAIN plans. Small table scans are efficient because SQL Server can often cache the entire table in a single I/O operation. Large numbers of index scans are normally desirable too, since it typically indicates the fastest possible resolution to data access requests.
Range Scans	Range Scans are the total number of qualified range scans through indexes. Large numbers of index scans are normally desirable, since it typically indicates the fastest possible resolution to data access requests is being taken.
Probe Scans	Probe scans are used in SQL Server to directly find rows in an index or base table. Large numbers of probe scans are normally desirable, since they typically indicate the fastest possible resolution to data access requests is being taken.
Index Searches	Index searches are normally used to start range scans, for single index record fetches and also might be used to reposition an index.
Forward Record Fetches	The Forwarded Record Fetches statistic represents the total number of records fetched by reading forwarded record pointers. SQL Server will move rows in a table under certain conditions. One situation might arise when you update a row in a table that has a variable-length column to a larger size that will no longer fit on its original page. Another situation would be if SQL Server moves a row when the clustered index column changes. When SQL Server creates a forwarding pointer, it remains in place unless one of two things happens. The first is when a row shrinks enough to move back to its original location. The second is when the entire database shrinks. When a database file shrinks, SQL Server will reassign the row identifiers, which are used as the row locators, so the shrink process never generates forwarded rows. Forwarded records can reduce performance at times because additional I/O is involved to first obtain the record pointer to the relocated row, and then the row itself is read.

Correction

Actions to take after observing the ways SQL Server is accessing data depend on the actual situation you witness. If consistent numbers are present for Forward Record Fetches, you might want to examine your databases to see which tables have forwarded records. You can easily do this by utilizing the Embarcadero Space Analyst to view forwarded record data for tables.

If you do not have Space Analyst, then to see the total count of forwarded records in a table, you can enable trace flag 2509, and then execute the DBCC CHECKTABLE command. The output should display the number of forwarded records in that table. Tables with many forwarded records could be candidates for table reorganization.

Here are some methods you can use to avoid unnecessary full scans:

- Try not to use SQL statements that include the NOT IN, NOT LIKE, <>, IS NULL operators since 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 is not, the index will not be used at all.
- Avoid using functions in WHERE predicates.

Workload Analysis pane

It is not uncommon for one or two users to cause the majority of runtime problems plaguing a database. The problem could be a runaway process, an untuned batch procedure, or other user-initiated operation. The Top I/O Hogs display showcases the top sessions that have caused the most physical I/O on the system (both reads and writes).

>

Metrics

If you see one or two users who have caused more than 25-50% of total I/O, then you should drill down further into the sessions to see what activities they are performing.

System I/O Tab

The System I/O tab of the I/O Detail includes the following sections:

AWE I/O	I/O Activity	Physical I/O
Space I/O		

Physical I/O

SQL Server performs many system-related I/O functions to keep data moving into and out of the server. The System I/O tab of the I/O Detail displays statistics that track various system-related I/O operations. The Physical I/O section details I/O tasks that involve physical disk activity. The table below describes the information available in the Physical I/O section:

Information	Description
Page Reads	The number of physical database page reads that are issued per second by SQL Server. It is collective in nature, meaning that it represents the total page reads per second across all databases that exist on the target SQL Server.
Page Writes	The number of physical database page writes issued by SQL Server. Page Writes take place during operations such as checkpoints, lazy writer writes, index creations, and BCP routines.
Checkpoint Pages	The number of pages flushed to disk per second by a checkpoint or other operation that require all dirty (modified) pages to be flushed.
Lazy Writes	The number of buffers written by the buffer manager's lazy writer. The lazy writer is a system process that flushes out batches of dirty and aged buffers and makes them available to SQL Server processes. The lazy writer eliminates the need to perform frequent checkpoints in order to create available buffers.
Read Ahead Pages	The number of physical database pages read in anticipation of use by SQL Server. If SQL Server senses that database pages are being read in a sequential manner, it institutes a pre-fetch mechanism that moves pages into the buffer cache before they are actually needed by a process.
Page Lookups	The number of requests to find a page in the buffer pool.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > System I/O Tab

Metrics	Page reads are to be expected, especially after initial server start up. SQL Server must first satisfy requests for data and metadata by reading information in from physical disk. Numerous page reads can also be expected if the physical server does not contain an adequate amount of memory to hold repetitively requested blocks of information. No hard and fast rules exist for how many page reads per second is too much. You can cross reference this statistic with the physical server disk statistics to see if physical page reads and accompanying physical disk I/O is approaching the server's premium capacity levels. And because logical I/O is always many times faster than physical I/O, you should also evaluate the buffer cache hit ratio to determine overall memory vs. physical read efficiency. Page Writes can give you an idea of overall physical write activity. There are a number of statistics, however, that pertain specifically to certain write activities like checkpoints, that can be examined to determine the amount of physical writes caused by distinct SQL Server processes. With respect to performance, response times experienced by SQL Server users are normally not impacted by write operations unless the writes are synchronous in nature. These are typically BCPs, database recovery operations, and index creations. With respect to read ahead pages, If data is accessed sequentially (for example, through the use of a clustered index), the read ahead mechanism of SQL Server can increase performance by needed database pages already in the buffer cache before they are actually requested. However, because the read ahead mechanism is typically triggered by full table or index range scans, if the read ahead pages are actually required to satisfy a user's query, performance can actually suffer. In these cases, the judicious use of indexes is a better route to take. Checkpoint operations are used by SQL Server to minimize the amount of work the server must perform when databases are recovered during system startup. Checkpoints periodically write out modified pages that are found in the buffer cache to disk. Afterward, SQL Server records the operation in the log to signify that the operation succeeded. A database owner issuing the checkpoint command can explicitly perform checkpoints. SQL Server also performs checkpoints automatically for databases that have the trunc log on chkpt option set. Large SQL Servers have the potential to generate lots of checkpoint write operations. Although SQL Server minimizes checkpoint activity, you can also set the recovery interval server parameter to influence how often checkpoints can run.
Correction	If you find that the server is becoming overworked from a physical I/O standpoint, here are a few things you can do: *Examine index usage to ensure that unnecessary table scans are not occurring.* *Check the physical database design to see if table objects have been over-normalized.* *Ensure that SQL Server is configured to use sufficient amounts of memory. Examine the min server memory (MB) and max server memory (MB) parameters to see if SQL Server is constrained on either end of the memory spectrum.* *Check for large pinned table objects that could be using excessive amounts of space in the buffer cache.* *Last, but not least, investigate the possibility of adding more RAM to the physical server.* If you believe excessive checkpoint activity is occurring, you can take the following steps: *Set the recovery interval server parameter to a larger value with sp_configure.* *Restart SQL Server or use the RECONFIGURE option to make the change immediately.*

Space I/O

SQL Server performs many system-related I/O functions to keep data moving into and out of the server. The Space I/O section of the I/O Detail details space-related I/O operations. The table below describes the information available in the Space I/O section:

Information	Description
Extents Allocated	The number of space extents that SQL Server allocated. Rapidly increasing numbers for these statistics indicates that SQL Server is receiving large volumes of incoming data and is allocating space to make room.

>

Information	Description
Extent Deallocations	This indicates that SQL Server is reclaiming space from database objects due to shrinking database volumes.
Freespace Page Fetches	The number of pages by free space scans used to satisfy requests to insert record fragments.
Freespace Scans	The number of scans performed by SQL Server to locate free space for an incoming record.
Mixed Page Allocations	The number of pages allocated from mixed extents. These are used for storing the first eight pages that are allocated to an index or table.
Page Deallocations	This indicates that SQL Server is reclaiming space from database objects due to shrinking database volumes.
Page Splits	When data is inserted or updated in a table, SQL Server might reorganize the storage of the data in the table's index pages. When an index page becomes full, but a DML operation demands room on that page, SQL Server moves about half the rows to a new page to accommodate the request. This reorganization is known as a page split. Performance for DML actions can be impaired from page split operations. In addition, more index pages can make for longer index scan times.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > System I/O Tab
Metrics	Increasing numbers for extent and page allocation, and freespace operations likely indicates aggressive volumes of data being inserted or modified in SQL Server. Page splits cause additional overhead in the form of CPU usage and I/O. Observing large numbers of page splits can signal a resource bottleneck in your server.
Correction	To avoid page splits, you can look into tuning the FILLFACTOR property of an index, which controls the percentage of the index page that is filled during creation. The default, 100, tells SQL Server to completely fill each page, whereas lower numbers tell SQL Server to leave room for additional index rows or updates to existing rows.

I/O Activity

SQL Server performs many system-related I/O functions to keep data moving into and out of the server. The I/O Activity section of the I/O Detail details maintenance-related I/O operations. The table below describes the information available in the I/O Activity section

Information	Description
DBCC Logical Scans	The number of logical read scan bytes per second caused by DBCC operations.
Bulk Copy Rows	The number of rows copied either into or out of the database via the BCP utility.
Bulk Copy Throughput	The amount of data (in KB) copied via BCP operations.
Transactions	The number of transactions that have occurred within the database.
Log Flushes	The number of log flushes for the server.
Backup/Restore T-Put	Defines the read/write throughput for backup and restore operations.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > System I/O Tab

>

AWE I/O

SQL Server performs many system-related I/O functions to keep data moving into and out of the server. The AWE I/O section of the I/O Detail details I/O relating to Address Windowing Extensions (the SQL Server mechanism for supporting very large amounts of RAM). The table below describes the information available in the AWE I/O section:

Information	Description
Lookup Maps	The number of times that a database page was requested by the server, found in the buffer pool, and mapped. When it is mapped, it is made a part of the server's virtual address space.
Stolen Maps	The number of times that a buffer was taken from the free list and mapped.
Unmap Cells	The number of calls to unmap buffers. When a buffer is unmapped, it is excluded from the virtual server address space. One or more buffers might be unmapped on each call.
Unmap Pages	The number of SQL Server buffers that are unmapped.
Write Maps	The number of times that it is necessary to map in a dirty buffer so it can be written to disk.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > System I/O Tab

User I/O Tab

The User I/O tab of the I/O Detail includes the following sections:

Lock Activity	SQL Activity	Temporary Object Activity
User Object Activity		

User Object Activity

The User Object Activity section of the I/O Detail displays statistics that track various user-related I/O operations. The I/O function along with its counter value is presented. The User Object Activity section details performance statistics that reflect how SQL Server is performing object access operations. The table below describes the information available in the User Object Activity section:

Information	Description
Forwarded Records	The number of records per second fetched through forwarded record pointers. At times forwarded records can reduce performance because additional I/O is involved to first obtain the record pointer to the relocated row, and then the row itself is read.
Full Scans	Full scans of moderately sized indexes or tables are generally okay. SQL Server can scan and cache a small table much faster than using its index to navigate to any requested data. Full, unrestricted, large table scans, however, are typically not good and degrade overall system performance and response time.
Index Searches	The total number of index searches per second. Index searches are normally used to start range scans, for single index record fetches and can be used to reposition an index.
Probe Scans	The total number of probe scans per second. Probe scans are used in SQL Server to directly find rows in an index or base table.
Range Scans	The total number of qualified range scans through indexes per second.
Skipped Ghosted Records	The number of ghosted records per second skipped during scans.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > User I/O Tab

Metrics	Full scans occur if a table is inadequately indexed or if SQL Server truly needs to access all rows in a table or index to satisfy a query. UPDATE STATISTICS calls can also cause full scans. Unnecessary scans on large tables is something to avoid, and can be a signal to you as a DBA to investigate the use of more indexes and to review SQL access through EXPLAIN plans. Small table scans are actually a good thing because SQL Server can often cache the entire table in a single I/O operation. Large numbers of index scans are normally desirable too, because this typically indicates the fastest possible resolution to data access requests. When SQL Server creates a forwarding pointer, it remains in place unless one of two things happens. The first is when a row shrinks enough to move back to its original location. The second is when the entire database shrinks. When a database file shrinks, SQL Server reassigns the row identifiers, which are used as the row locators, so the shrink process never generates forwarded rows. Forwarded records can reduce performance at times because additional I/O is involved to first obtain the record pointer to the relocated row, and then the row itself is read. Large numbers of index searches and probe scans are normally desirable because they typically indicate the fastest possible resolution to data access requests is being taken.
Correction	Here are some methods you can use to avoid unnecessary large table scans: *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 is not, the index will not be used at all.* *Avoid using functions in WHERE predicates.* If consistent numbers are present for Forward Record Fetches, examine your databases to see which tables have forwarded records. This can easily be done with the Embarcadero Space Analyst component. If you do not have Space Analyst, then to see the total count of forwarded records in a table, enable trace flag 2509, and then execute the DBCC CHECKTABLE command. The output should display the number of forwarded records in that table. Tables with many forwarded records could be candidates for table reorganization.

SQL Activity

The SQL Activity section of the User I/O tab of the I/O Detail displays statistics that track various user-related I/O operations. The I/O function along with its counter value is presented. The SQL Activity section details performance statistics that reflect SQL I/O-related operations. The table below describes the information available in the SQL Activity section:

Information	Description
Auto-Param Attempts	Auto-parameterization occurs when an instance of SQL Server attempts to reuse a cached plan for a previously executed query that is similar to, but not the same as, the current query. The Auto-param Attempts statistic shows the number of auto-parameterization attempts per second and includes failed, safe, and unsafe auto-parameterizations.
Batch Requests	A batch is a collection of one or more SQL statements sent in one unit by the client. Each batch is compiled into a single execution plan. If the batch contains multiple SQL statements, all of the optimized steps needed to perform all the statements are built into a single execution plan. The Batch Requests statistic shows the number of batch requests per second that are processed by the SQL Server instance.
Failed Auto-Params	Auto-parameterization occurs when an instance of SQL Server attempts to reuse a cached plan for a previously executed query that is similar to, but not the same as, the current query. The Failed Auto-params statistic shows the number of failed auto-parameterization attempts per second.
Safe Auto-Params	Auto-parameterization occurs when an instance of SQL Server attempts to reuse a cached plan for a previously executed query that is similar to, but not the same as, the current query. The Safe auto-params statistic shows the number of auto-parameterization attempts per second and includes only safe auto-parameterizations (ones where the cached plan can be shared).

Information	Description
Scan Point Revalidations	The number of times per second that the scan point had to be revalidated to continue the scan.
SQL Compilations	The number of SQL compilations performed, indicating the number of times the compile code path is entered. This also includes compiles due to recompiles. When SQL Server user activity levels become stable, this value reaches a steady state.
SQL Re-Compilations	The total number of recompiles triggered per second in a SQL Server instance. Recompiles occur when SQL Server determines that the currently defined execution plan for an executing stored procedure might no longer be the best possible plan. SQL Server pauses the query execution and recompiles the stored procedure.
Unsafe Auto-Params	Auto-parameterization occurs when an instance of SQL Server attempts to reuse a cached plan for a previously executed query that is similar to, but not the same as, the current query. The Unsafe auto-params statistic shows the number of auto-parameterization attempts per second and includes only safe auto-parameterizations (ones where the cached plan cannot be shared).

Location	Microsoft SQL Server Performance Analyst statistics > I/O > User I/O Tab
Metrics	SQL Server's ability to match new SQL statements with existing, unused execution plans is increased when parameters or parameter markers are used in Transact-SQL statements. If an SQL statement is executed without parameters, SQL Server parameterizes the statement internally to increase the possibility of matching it against an existing execution plan. A high number for auto-param attempts shows that SQL Server is efficiently reusing existing cached plans. A small number for failed auto-param attempts shows that SQL Server is efficiently reusing existing cached plans. Because compilation is a significant part of a query's turnaround time, you should strive to have as many compilations stored in the cache as possible. If this number does not stabilize in direct proportion to user activity stabilizing, you should investigate your SQL Cache to see if it has adequate memory assigned to it. Recompiles slow down the process that is executing the procedure and increases the load on the CPU. By extension, the more recompiles that are occurring on your system, the more overall load increases resulting in poor performance. In general, you want to keep the number of recompiles low. The most common reasons SQL Server would issue a recompile are: *Running sp_recompile against any table referenced in the stored procedure.* *Significant data changes in a referenced table.* *Schema changes to referenced objects.* *The use of the WITH RECOMPILE clause in the CREATE PROCEDURE or EXECUTE statement.* *A plan no longer available in the system cache.*
Correction	For failed auto-param attempts, you can increase the ability of the relational engine to match complex SQL statements to existing, unused execution plans, by explicitly specify the parameters using either sp_executesql or parameter markers in your T-SQL code. For high numbers of SQL re-compilations, try to practice coding standards that eliminate the most frequent causes detailed above. Also, try to: *Use temporary tables only in the stored procedure that created them.* *Minimize creating temporary tables in control block structures.* *Use the KEEP PLAN option on references to static temporary tables.* *Issue the CREATE TABLE statement before any other references to the created table.* *Minimize the use of temporary tables.*

Lock Activity

The Lock Activity section of the User I/O tab of the I/O Detail displays statistics that track various user-related I/O operations. The I/O function along with its counter value is presented. The Lock Activity section details performance statistics that reflect how SQL Server is handling lock operations. The table below describes the information available in the Lock Activity section:

Information	Description
Table Lock Escalations	The number of times locks on a table were escalated.
Locks	The total number of locks acquired on the server.
Blocked Users	The users currently being blocked by other processes.
Deadlocks	The number of deadlocks detected by SQL Server. Page Deadlocks occur when processes cannot proceed because they are waiting on a set of resources held by each other or held by other processes.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > User I/O Tab
Metrics	Many table lock escalations could indicate contention problems. If increasing numbers of table lock escalations are viewed at the same time as blocking or deadlock problems, the application design could be at fault. Consistently seeing positive numbers for the blocked statistic should also clue you into the fact that a bottleneck exists for some processes. You can easily drill down and discover the exact process(es) holding locks that are blocking out other user activity. Another situation to look for with respect to locking, is when the total number of acquired locks reaches the maximum lock limit currently set on SQL Server. Consistently seeing page deadlock counts greater than zero indicates that some user processes are experiencing delays completing their work. When SQL Server identifies a page deadlock, it resolves the situation by choosing the process that can break the deadlock. This process is termed the deadlock victim. SQL Server rolls back the deadlock victim's transaction, and then notifies the process' application by returning an error message. It also cancels the process' request and allows the transactions of the remaining processes to continue. SQL Server always attempts to choose the least expensive thread running the transaction as the deadlock victim.
Correction	Once discovered, a blocking lock situation can normally be quickly remedied - the DBA issues a KILL against the offending process, which eliminates the stranglehold on the objects the user was accessing. Other user processes then almost always complete in an instant. Discovering the blocked lock situation is made easier by using tools like Embarcadero Performance Analyst, but preventing the blocking lock situation in the first place is tricky. The culprit of blocking lock scenarios is usually the application design, or the SQL being used within the application itself. Properly coding an application to reference database objects in an efficient order, and then using the right SQL to get the job done, is an art. The key to avoiding lock contention is to process user transactions in the quickest and most efficient manner possible - something not always easy to do. By default, all processes wait indefinitely for locks in SQL Server. You can change this behavior by using the SET LOCK_TIMEOUT command, which limits the number of seconds that a process waits for a lock before timing out. Because SQL Server automatically resolves deadlock situations, you should work proactively to prevent them in the first place. You can change default deadlock behavior by using the SET DEADLOCK_PRIORITY command, which reprioritizes a process' position in a deadlock situation.

Temporary Object Activity

The Temporary Object Activity section of the User I/O tab of the Users Detail displays statistics that track various user-related I/O operations. The I/O function along with its counter value is presented. The Temporary Object Activity section details performance statistics that reflect how SQL Server is handling the creation of temporary objects, such as temporary tables used in stored procedures. The table below describes the information available in the Temporary Object Activity section:

Information	Description
Workfiles Created	The number of workfiles created by SQL Server.
Worktables Created	The total number of work tables created. Worktables are used many times by SQL Server to perform a logical operation specified in an end-user SQL statement. GROUP BY, ORDER BY, or UNION queries can cause worktables to be created as can specific CREATE statements used in Transact SQL processing. Worktables are built in the tempdb database and are dropped automatically at the end of the statement or procedure run.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > User I/O Tab
Metrics	With respect to worktables, keep in mind that the tempdb database should be large enough to hold large worktables.

Database I/O Tab

The Database I/O tab of the I/O Detail includes the following sections:

Database Bytes Read/Write Summary	Location	Database Read/Write Summary

Database Read/Write Summary

The Reads statistic represents that number of physical database page reads that are issued per second by SQL Server. The Writes statistic represents the number of physical database page writes issued by SQL Server. Writes take place during operations such as checkpoints, lazywriter writes, index creations, and BCP routines. Both statistics are collective in nature, meaning that they represents the total page reads and writes across all databases that exist on the target SQL Server.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > Database I/O Tab
Metrics	Page reads are to be expected, especially after initial server start up. This is because SQL Server must first satisfy requests for data and metadata by reading information in from physical disk. Numerous page reads can also be expected if the physical server does not contain an adequate amount of memory to hold repetitively requested blocks of information. No hard-and-fast rules exist for how many page reads per second is too much. You can cross reference this statistic with the physical server disk statistics to see if physical page reads and accompanying physical disk I/O is approaching the server's premium capacity levels. And because logical I/O is always many times faster than physical I/O, you should also evaluate the buffer cache hit ratio to determine overall memory vs. physical read efficiency. Page Writes can give you an idea of overall physical write activity. There are a number of statistics, however, that pertain specifically to certain write activities like checkpoints, that can be examined to determine the amount of physical writes caused by distinct SQL Server processes. Regarding performance, response times experienced by users are normally not impacted by write operations unless the writes are synchronous. These are typically BCPs, database recovery operations, and index creations.

>

Database Bytes Read/Write Summary

The Database Bytes Read/Write Summary statistics represent the number of bytes read and written by SQL Server.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > Database I/O Tab

Database I/O Detail

The Database I/O section of the Database I/O tab of the I/O Detail summarizes I/O activity for each database, letting you quickly spot the “hot” databases on your server. The table below describes the information available in the Database I/O Detail section for SQL Server 2000 and later:

Information	Description
Database	The name of the database.
Reads	The number of reads issued against the database.
Writes	The number of writes issued against the database.
Bytes Read	The total number of bytes read for the database.
Bytes Written	The total number of bytes written for the database.
I/O Stall	The total amount of time that processes have waited for I/O operations to complete, in milliseconds.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > Database I/O Tab
Metrics	Consider moving databases with lots of I/O activity and wait time onto separate drives/devices.

File Tab

The File tab of I/O Detail includes the following sections:

.File Bytes Read/Write Summary	File I/O Detail	File Read/Write Summary

File Read/Write Summary

The Reads statistic represents that number of physical database page reads that are issued per second by SQL Server. The Writes statistic represents the number of physical database page writes issued by SQL Server. Writes take place during operations such as checkpoints, lazywriter writes, index creations, and BCP routines. Both statistics are collective in nature, meaning that they represents the total reads and writes across all databases that exist on the target SQL Server.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > File Tab

>

Metrics	Page reads are to be expected, especially after initial server start up. This is because SQL Server must first satisfy requests for data and metadata by reading information in from physical disk. Numerous page reads can also be expected if the physical server does not contain an adequate amount of memory to hold repetitively requested blocks of information. No hard and fast rules exist for how many page reads per second is too much. You can cross reference this statistic with the physical server disk statistics to see if physical page reads and accompanying physical disk I/O is approaching the server's premium capacity levels. And because logical I/O is always many times faster than physical I/O, you should also evaluate the buffer cache hit ratio to determine overall memory vs. physical read efficiency. Page Writes can give you an idea of overall physical write activity. There are a number of statistics, however, that pertain specifically to certain write activities like checkpoints, that can be examined to determine the amount of physical writes caused by distinct SQL Server processes. Regarding performance, response times experienced by users are normally not impacted by write operations unless the writes are synchronous in nature. These are typically BCPs, database recovery operations, and index creations.

.File Bytes Read/Write Summary

The File Bytes Read/Written statistics communicate the number of bytes read and written by SQL Server.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > File Tab

File I/O Detail

The File I/O Detail section of the File tab of the I/O Detail summarizes I/O activity for each database file, letting you quickly spot the “hot” databases and files on your server. The table below describes the information available in the File I/O Detail section for SQL Server 2000 and later:

Information	Description
Database	The database name.
File ID	The file identifier for the target file.
Logical Name	The name given the file by the DBA.
File Name	The physical file name of the file.
Timestamp	The internal time stamp of when the data was obtained.
Reads	The number of reads issued against the database file.
Writes	The number of writes issued against the database file.
Bytes Read	The total number of bytes read for the database file.
Bytes Written	The total number of bytes written for the database file.
I/O Stall	The total amount of time that processes have waited for I/O operations to complete, in milliseconds.

Location	Microsoft SQL Server Performance Analyst statistics > I/O > File Tab
Metrics	Consider moving databases and/or files with lots of I/O activity and wait time onto separate drives/devices.

Space

In addition to a Home page, the Space category of SQL Server Performance Analyst includes the following tabbed pages:

File Groups Tab	Files Tab	Objects Tab
Virtual Log Files Tab	Server Space Tab	

Home

The Home page of the Space category includes statistics in the following categories:

Bottleneck Analysis pane	Disk Analysis pane	Space Analysis pane
SQL Analysis pane		

Space Analysis pane

A SQL server contains many databases, some of which are devoted to system-level activities (the master and tempdb databases, for example) and others that hold user data. The Key Space Analysis section of the Space home page displays details about the space situation for each database in a SQL server, including their total, used, and free space. The percent used amount for each database is also shown. In addition, each database's transaction log space is analyzed and presented.

Metrics

If a database or transaction log's free space goes below the Performance Analyst recommended threshold, (and either the database or transaction log does not have its automatic growth property enabled or the files have reached their growth limit) then the DBA should take action to ensure that the database or transaction log does not run out of available free space.

Correction

If the percent used amount of a database is approaching problematic levels, there are three ways a DBA can rectify the situation:

1. The DBA can resize the current file(s) used by the database via an ALTER DATABASE … MODIFY FILE command.
2. The DBA can add a new file to the database via the ALTER DATABASE … ADD FILE command.
3. The DBA can modify the file(s) used by the database to automatically grow by using the ALTER DATABASE … MODIFY FILE … FILEGROWTH command. You should also ensure that the MAXSIZE setting of each file is set appropriately.

Of course, the DBA should also ensure that enough physical space exists on the server to accommodate additional database space. There are several things a DBA can do to ensure that a database's log does not run out of available free space:

1. First, most transactional-oriented databases should have their logs assigned to separate physical drives other than the database. Reasons for doing this include:
   - It prevents competition for space between the log and the database itself.
   - It allows the log to be monitored for space more effectively.
   - It improves performance.

2 If the database is not critical in nature, you can set the truncate log on checkpoint option (trunc log on chkpt), which eliminates any non-active space in the log when a database checkpoint occurs.

3 Critical databases needing higher levels of recovery should have schedules established that regular perform transaction log dumps. Doing so ensures better recovery scenarios as well as a reduced risk of the transaction log running out of space.

4 If a critical transaction log becomes full, it might be impossible to use standard procedures to dump transactions and reclaim space. The dump operation will likely have to incorporate the no log or truncate only options.

5 If a transaction log continuously approaches dangerously low levels of free space, then the DBA should allow the underlying file(s) of the log to automatically grow to meet the demand. This can be accomplished by using the ALTER DATABASE … MODIFY FILE … FILEGROWTH command. You should also ensure that the MAXSIZE setting of each file is set appropriately.

The DBA should also be on the lookout for large load or data modification operations that do not make use of prudently timed commit points. A single, large transaction has the ability to overwhelm any transaction log since only non-active space in the transaction log is removed from log dumps or truncation operations.

Bottleneck Analysis pane

When using bottleneck analysis to troubleshoot a database, a DBA cannot rely only on the information contained in the wait event views that SQL Server provides. For example, a database might attempt to allocate another segment of space and yet be denied if no such free space exists on the server or if the underlying file has reached its growth limit. Such a failure will not be reflected in any wait event, but still represents a very real bottleneck to the database. In the same way that a DBA cannot depend on only a few ratios to properly carry out ratio-based performance analysis, an administrator must include other statistical metrics in their overall bottleneck analysis framework to obtain an accurate performance risk assessment. Performance Analyst works to identify bottlenecks in your database that fall outside of pure wait events so you can get a total picture of all stoppages in your system. The Performance Analyst Space home page displays space-related bottlenecks that might be threatening the overall performance of your system.

The following bottleneck indicators are used on the Performance Analyst Space home page to succinctly communicate the general overall space situation of the monitored database:

Databases Without Auto-Grow	Databases With Auto-Shrink	Database Files Near Max Size
Databases on Same Disk as Log	Logs Without Auto-Grow	Log Files Near Max Size

The following statistics, available on this pane, duplicate statistics available on the Bottleneck Analysis pane of the Home page:

DB's Low on Space	Logs Low on Space	

Database Files Near Max Size

A DBA can specify that one or more files of a SQL Server database can automatically grow by enabling the file's auto-growth property. A DBA can also specify that the file not be permitted to grow beyond a certain limit, and once that limit is reached, the file and/or the database might be permanently out of space. The Database Files Near Max Size statistic provides a count of how many database files are near their maximum file size.

Location	Microsoft SQL Server Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	A DBA should not see any positive indicator that any file has reached its growth limit. While it might be okay from the standpoint that a file cannot simply grow anymore (if the database has other files that can grow or that no more space will be added to the database), the DBA should disable the auto-growth property of the file.

Correction	If a file has reached its growth limit, a DBA can take a number of actions, including: *Eliminate the growth limit for the file either by increasing the limit size or allowing to grow in an unlimited fashion.* *If the file is on a disk or file system that has no free space left, the DBA can add another file to the database to allow for space expansion and alter the problem file to have auto-growth property disabled.*

Databases on Same Disk as Log

To minimize I/O contention, a database should have its database files located on different physical disks than its log files. The Databases on Same Disk as Log statistic provides a count of how many databases have files that are on the same physical disk as their transaction log.

Location	Microsoft SQL Server Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	For critical production systems, all databases should have their database and log files on different physical disks.
Correction	If a DBA has database and log files on the same physical disk, they can create new database filegroups on different disks, relocate objects to the new filegroups, and then delete the old files/filegroups.

Databases Without Auto-Grow

In Microsoft SQL Server, a DBA has the ability to tell SQL Server to automatically grow a database in size when more space is required. This feature can save a critical transaction or other database request from failing due to a lack of free space in the database. It is recommended that critical databases have this feature enabled. The Databases Without Auto-Grow statistic provides a count of databases that do not have their automatic growth property enabled.

Location	Microsoft SQL Server Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	Static databases (those not expected to grow in size) will likely not need their auto-growth property enabled. Growing, dynamic databases should almost always be allowed to automatically grow when needed.
Correction	If any critical, dynamic database is found to not have its auto-growth feature enabled, then the DBA can modify the file(s) used by the database to automatically grow by using the ALTER DATABASE … MODIFY FILE … FILEGROWTH command. You should also ensure that the MAXSIZE setting of each file is set appropriately.

Databases With Auto-Shrink

A database can be set up to automatically shrink in size if SQL Server finds the database with an abundance of unused free space. Setting a database up so that this occurs can help reduce wasted storage on a database server. The Databases With Auto-Shrink statistic provides a count of databases that are allowed to automatically shrink in size.

Location	Microsoft SQL Server Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	Dynamic databases that have the potential to balloon in size and then have large amounts of that same data removed should be considered candidates for auto-shrink.
Correction	If a DBA wants to configure a database to automatically shrink, they can issue a command similar to the following: **EXEC sp_dboption 'mydatabase','autoshrink',true**

Logs Without Auto-Grow

In Microsoft SQL Server, a DBA has the ability to tell SQL Server to automatically grow a database or transaction log in size when more space is required. This feature can save a critical transaction or other database request from failing due to a lack of free space in the database or transaction log. It is recommended that critical databases and their transaction logs have this feature enabled.

The Logs Without Auto-Grow statistic provides a count of transaction logs that do not have their automatic growth property enabled.

Location	Microsoft SQL Server Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	Static databases (those not expected to grow in size) will likely not need their transaction log's auto-growth property enabled. Growing, dynamic databases should almost always have their transaction log be set to automatically grow when needed.
Correction	If any critical, dynamic database is found to not have their transaction log auto-growth feature enabled, then the DBA can modify the file(s) used by the database's transaction log to automatically grow by using the ALTER DATABASE … MODIFY FILE … FILEGROWTH command. You should also ensure that the MAXSIZE setting for each file is set appropriately.

Log Files Near Max Size

A DBA can specify that one or more transaction log files of a SQL Server database can automatically grow by enabling the file's auto-growth property. A DBA can also specify that the file not be permitted to grow beyond a certain limit, and once that limit is reached, the file and/or the transaction log might be permanently out of space. The Log Files Near Max Size statistic provides a count of how many transaction log files are near their maximum file size.

You should not see any positive indicator that any file has reached its growth limit. While it might be okay from the standpoint that a file cannot simply grow anymore (if the database transaction log has other files that can grow or that no more transactions are expected to occur within the database), the DBA should disable the auto-growth property of the file.

Location	Microsoft SQL Server Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Correction	If a file has reached its growth limit, a DBA can take a number of actions: *Eliminate the growth limit for the file either by increasing the limit size or allowing to grow in an unlimited fashion.* *If the file is on a disk or file system that has no free space left, the DBA can add another file to the database's transaction log to allow for space expansion and alter the problem file to have auto-growth property disabled.*

SQL Analysis pane

A lot of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured database down in terms of end-user response times.

The SQL Analysis section of the Space home page shows a count of all SQL statements that have executed INSERT, UPDATE, DELETE, and TRUNCATE commands. Seeing these statistics should help you determine how much space-related SQL activity is occurring on your system.

The SQL Analysis section provides you with a quick overview of the percentage of Reads, Writes, CPU, and Elapsed Time the most expensive SQL statements tracked by Performance Analyst have used. Depending on the page you are on, you might just see information regarding physical I/O activity or DML statements. For example, you might see that a SQL statement has caused 60% of all physical reads on the system, or that a procedure is responsible for 90% of all the CPU usage. To get information regarding the actual SQL text or stored procedure calls, drill down into the Top SQL details views.

NOTE: SQL Profiling is turned on by default in Performance Analyst so you can automatically collect SQL traffic. If you do not want to collect SQL execution information, use the options inside of Performance Analyst to disable SQL profiling.

Metrics

The statistics in this section can be used in many different ways. For example, large numbers of inserts could indicate a growing database and that a DBA should keep a watchful eye on free space. Heavy update numbers could mean trouble for objects containing many indexes (indexes typically undergo splits and fragmentation when large volumes of updates are thrown their way). Many deletes might signal potential wasted space problems.

Disk Analysis pane

The Disk Analysis section of the Performance Analyst Space home page provides a summary of the total database and log space used per server disk drive, as well as a summary of free disk space for each server drive.

Metrics

If any database or transaction log file has been set up to automatically grow, the DBA should ensure there is enough server disk space to accommodate any new, additional requests for space.

Correction

If you see any drive that has reached zero free space (or is close), you might want to add new files, on other disks with abundant free space, to any databases or transaction logs so that no out-of-space errors result.

File Groups Tab

SQL Server manages physical storage space through files and file groups. The file groups tab displays detailed information regarding storage usage in each of the SQL Server file groups. The table below describes the information available in the File Groups Summary section of the File Groups tab:

Information	Description
Database Name	The name of the database.
File Group ID	The numeric ID of the file group.
File Group	The name of the file group.
Can Grow	This indicates if one of more of the files making up the file group is allowed to grow automatically.
File Count	The number of files that make up the file group.
Size (MB)	The total physical size of the file group.
Table Res	The amount of reserved space consumed by tables.
Index Res	The amount of reserved space consumed by indexes.
Free (MB)	The amount of free space for the file group.
Free Percent	The percentage of free space for the file group.

Metrics

Unless server space is tight, it is normally wise practice to allow your files to automatically grow to meet demand for more incoming data. You should also physically separate your database and log files onto separate physical drives.

Correction

To let your files automatically grow until out of space, set the file's growth option in Embarcadero's DBArtisan or Embarcadero's Rapid SQL, or by using the ALTER DATABASE... MODIFY FILE command.

Files Tab

SQL Server manages physical storage space through files and file groups. The Files tab displays detailed information regarding storage usage in each of the SQL Server files. The table below describes the information available in the Files Summary section of the Files tab:

Information	Description
Database Name	The name of the database.
Logical Name	The nickname given to the file by the DBA.
File Group	The name of the file group.
File Name	The name and location of the file.
Size (MB)	The total physical size of the file.
Can Grow	This indicates if the file can automatically grow in size.
Growth Amount	This indicates how much the file will grow in size.
Max File Size	This indicates the maximum file size that the file can grow.

Metrics

Unless server space is tight, it is normally wise practice to allow your files to automatically grow to meet demand for more incoming data. It is also smart to physically separate your database and log files onto separate physical drives.

Correction

To let your files automatically grow until out of space, set the file's growth option in Embarcadero's DBArtisan or Embarcadero's Rapid SQL, or by using the ALTER DATABASE... MODIFY FILE command.

Virtual Log Files Tab

The Virtual Log Files tab of the Space Detail displays an internal structure of sorts for a each database's log. The presented information is helpful when trying to shrink a database's log because you can see how much of the log is active and exactly where the active portion resides. The table below describes the information available on the Virtual Log Files tab of the Space Detail:

Information	Description
Database Name	The name of database.
File Name	The name of the log file.
Status	This indicates if this portion of the log is active or inactive (not being used).
Size	The size of this portion of the log in MBs.

>

Objects Tab

The Objects tab of the Space/Databases Detail includes the following sections:

Database Object Detail	Database Object Summary	

Database Object Summary

Tables and indexes consume the storage in all databases. The Database Object Summary grid displays summary information regarding table and index storage for all databases. The table below describes the information available in the Database Object Summary section of the Objects tab of the Space Detail:

Information	Description
Database Name	The name of the database.
Table Count	The count of all tables in the database.
Index Count	The count of all indexes in the database.
Table Reserved (MB)	The amount of space (in MB) reserved by tables.
Index Reserved (MB)	The amount of space (in MB) reserved by indexes.

Location	Microsoft SQL Server Performance Analyst statistics > Space > Objects Tab Microsoft SQL Server Performance Analyst statistics > Database > Objects Tab
Metrics	Negative numbers viewed for index space information can be caused by inaccuracies contained in the SQL Server data dictionary. Frequently, running a DBCC UPDATEUSAGE command against the database resolves the problem. However, there are bugs in SQL Server that sometimes caused erroneous numbers to the reported for the reserved space amounts used by tables and indexes.

Database Object Detail

Tables and indexes consume the storage in all databases. The Database Object Detail grid displays object space details for the database selected in the Database Object Summary grid. The table below describes the information available in the Database Object Detail section of the Objects tab of the Space Detail:

Information	Description
Object Owner	The owner of the object.
Object Name	The name of the object.
Clustered	Indicates if the index is clustered.
File Group	The file group where the object resides.
Reserved (KB)	The amount of space reserved by the object in kilobytes.
Used (KB)	The amount of space used by the object in kilobytes.
Free (KB)	The amount of free space used by the object in kilobytes.
Percent Used	The percentage of space used by the object.

Location	Microsoft SQL Server Performance Analyst statistics > Space > Objects Tab Microsoft SQL Server Performance Analyst statistics > Database > Objects Tab
Metrics	Negative numbers viewed for index space information can be caused by inaccuracies contained in the SQL Server data dictionary. Frequently, running a DBCC UPDATEUSAGE command against the database resolves the problem. However, there are bugs in SQL Server that sometimes caused erroneous numbers to the reported for the reserved space amounts used by tables and indexes.

Server Space Tab

The Server Space tab of the Space Detail displays the amounts of space used by SQL Server per physical drive and by database per physical drive. The first section displays summary information and the second section contains space information per database.

The table below describes the information available in the Disk Summary By Space section on the Server Space tab of the Space Detail:

Column	Description
SQL Data (MB)	The amount of reserved database space on the drive.
SQL Log (MB)	The amount of reserved log space on the drive.
Free	The total amount of free space that remains on the drive.

The table below describes the information available in the Disk Summary By Database section on the Server Space tab of the Space Detail:

Column	Description
Disk Drive	The physical drive letter.
Database	The database name.
SQL Data (MB)	The amount of reserved database space on the drive.
SQL Log (MB)	The amount of reserved log space on the drive.

Metrics

If you allow your database and/or log files to automatically grow, and you see little or no free space left on their physical drives, an option is to add new files to the database or log on different drives to avoid any out of space errors. It is also smart to physically separate your database and log files onto separate physical drives.

Correction

If you need to add new files to your databases or logs, you can do so easily by using the ALTER DATABASE… ADD FILE and ALTER DATABASE… ADD LOG FILE commands.

Users

In addition to a Home page, the Users category of SQL Server Performance Analyst includes the following tabbed pages:

Locks Tab	Session Waits Tab	System Waits Tab

>

Top Sessions Tab		

Home

The Home page of the Users Performance Analyst statistics includes the following categories:

Bottleneck Analysis pane	Key User Analysis pane	Login Analysis
SQL Analysis pane	Workload Analysis pane	

Key User Analysis pane

User database activity can be examined using key performance statistics and wait/bottleneck-based analysis.

The following key statistics are used on the Performance Analyst Users home page to succinctly communicate the general session-based activity levels of the monitored database:

Active Processes	The Active Processes statistic represents the total number of active and open threads reported on the server. This number displays the number of processes actively performing work.
Active Transactions	The active transactions statistic represents a count of the number of in-process transactions for SQL Server.
Inactive Processes	This metric represents the total number of threads logged on to the server that are idle at the current time.
System Processes	The System Processes statistic represents the total number of threads logged on to the server that are SQL Server internal processes.
Transaction/Sec	The transaction/sec statistic refers to the number of database transactions processed per second by SQL Server.
T-SQL Batches	The T-SQL batches statistic refers to the number of transact SQL batches processed by SQL Server.

Bottleneck Analysis pane

When SQL Server is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting might mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. And this is where wait-based or bottleneck analysis comes into play. DBAs use this form of performance analysis to determine if perceived bottlenecks in a database are contributing to a performance problem.

Bottleneck analysis is a valid method of measuring performance because it helps a DBA track where a database has been spending its time. If lock contention or heavy table scan activity has been dragging down database performance, a DBA can use bottleneck analysis to confirm the actual root cause. Once one or more wait events or other bottlenecks have been pinpointed as possible performance vampires, the DBA can drill down and oftentimes discover a fair amount of detail about which sessions and objects are causing the problem. The Performance Analyst Home page identifies the top system and session waits that might be currently a cause of lessened performance.

When using bottleneck analysis, you cannot rely only on the information contained in the wait events that SQL Server provides. For example, a database might be suspect or offline. Such a failure will not be reflected in any wait event, but still represents a very real bottleneck to the database. In the same way that a DBA cannot depend on only a few ratios to properly carry out ratio-based performance analysis, an administrator must include other statistical metrics in their overall bottleneck analysis framework to obtain an accurate performance risk assessment. Performance Analyst works to identify bottlenecks in your database that fall outside of pure wait events so you can get a total picture of all stoppages in your system. The Performance Analyst Home page displays space-related bottlenecks as well as other sources of contention/bottlenecks that might be dragging down the overall performance of your system.

>

The following bottleneck indicators are used on the Performance Analyst Users home page to succinctly communicate the general overall performance level of the monitored database:

Current Blocked Processes	Failed Auto-Param Attempts	Lock Timeouts
Lock Waits	Lock Wait Time	Scan Point Revalidations
SQL Re-Compilations	Top Bottlenecked User Processes	

The following statistic, available on this pane, dupicates a statistic appearing on the Bottleneck Analysis pane of the Home page:

- Table Lock Escalations

Current Blocked Processes

A single blocking user has the potential to stop work for nearly all other processes on a small system, and can cause major headaches even on large systems. Blocks are most often caused by user processes holding exclusive locks and not releasing them via a proper COMMIT frequency. Unless a process times out via an application timeout mechanism, or the process has specified a timeout period via the SET LOCK_TIMEOUT command, a process waiting for a lock will wait indefinitely.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	You should immediately investigate any indicator above zero, before the situation has a chance to mushroom.
Correction	Once discovered, a blocking lock situation can normally be quickly remedied - the DBA issues a KILL against the offending process, which eliminates the user's stranglehold on the objects they were accessing. Other user processes then almost always complete in an instant. Discovering the blocked lock situation is made easier by using tools like Performance Analyst, but preventing the blocking lock situation in the first place is where it gets tricky. The DBA can drill down into user detail and view all current blocking locks to see exactly which sessions are holding the currently restrictive locks. The culprit of blocking lock scenarios is usually the application design, or the SQL being used within the application itself. Properly coding an application to reference database objects in an efficient order, and then using the right SQL to get the job done, is an art. The key to avoiding lock contention is to process user transactions in the quickest and most efficient manner possible - something not always easy to do. By default, all processes wait indefinitely for locks in SQL Server. You can change this behavior by using the SET LOCK_TIMEOUT command, which limits the number of seconds that a process will wait for a lock before timing out.

Failed Auto-Param Attempts

Auto-parameterization occurs when an instance of SQL Server attempts to reuse a cached plan for a previously executed query that is similar to, but not the same as, the current query. The Failed Auto-Param Attempts statistic shows the number of failed auto-parameterization attempts per second.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	SQL Server's ability to match new SQL statements with existing, unused execution plans increases when parameters or parameter markers are used in Transact-SQL statements. If an SQL statement is executed without parameters, SQL Server parameterizes the statement internally to increase the possibility of matching it against an existing execution plan. A small number for this statistic shows that SQL Server is efficiently reusing existing cached plans.

Correction	You can increase the ability of the relational engine to match complex SQL statements to existing, unused execution plans, by explicitly specify the parameters using either sp_executesql or parameter markers in your T-SQL code. Doing so helps lower this number.

Lock Timeouts

A single blocking user has the potential to stop work for nearly all other processes on a small system, and can cause major headaches even on large systems. Blocks are most often caused by user processes holding exclusive locks and not releasing them via a proper COMMIT frequency. Unless a process times out via an application timeout mechanism, or the process has specified a timeout period via the SET LOCK_TIMEOUT command, a process waiting for a lock will wait indefinitely. The lock timeouts represent the number of locks that have timed out due to the issuing of the aforementioned command.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	Any indicator above zero might indicate the presence of lock contention on the server.
Correction	Once discovered, a blocking lock situation can be quickly remedied - the DBA issues a KILL against the offending process, eliminating the user's stranglehold on the objects they were accessing. Other user processes then almost always complete in an instant. Discovering the blocked lock situation is made easier by using tools like Performance Analyst, but preventing the blocking lock situation in the first place is more difficult. You can drill down into user detail and view all current blocking locks to see exactly which sessions are holding the currently restrictive locks. The culprit of blocking lock scenarios is usually the application design, or the SQL being used within the application itself. Properly coding an application to reference database objects in an efficient order, and then using the right SQL to get the job done, is an art. The key to avoiding lock contention is to process user transactions in the quickest and most efficient manner possible - something not always easy to do. By default, all processes wait indefinitely for locks in SQL Server. You can change this behavior by using the SET LOCK_TIMEOUT command, which limits the number of seconds that a process will wait for a lock before timing out.

Lock Waits

A single blocking user has the potential to stop work for nearly all other processes on a small system, and can cause major headaches even on large systems. Blocks are most often caused by user processes holding exclusive locks and not releasing them via a proper COMMIT frequency. Unless a process times out via a timeout mechanism, or the process has specified a timeout period via the SET LOCK_TIMEOUT command, a process waiting for a lock will wait indefinitely. The Lock Waits statistic represents the number of lock requests that required a process to wait.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	Any indicator above zero might indicate the presence of lock contention on the server.
Correction	Once discovered, a blocking lock situation can be quickly remedied - the DBA issues a KILL against the offending process, which eliminates the user's stranglehold on the objects they were accessing. Other user processes then almost always complete in an instant. Discovering the blocked lock situation is made easier by tools like Performance Analyst, but preventing the blocking lock situation in the first place is more difficult. The DBA can drill down into user detail and view all current blocking locks to see exactly which sessions are holding the currently restrictive locks. The culprit of blocking lock scenarios is usually the application design, or the SQL being used within the application itself. Properly coding an application to reference database objects in an efficient order, and then using the right SQL to get the job done, is an art. The key to avoiding lock contention is to process user transactions in the quickest and most efficient manner possible - something not always easy to do. By default, all processes wait indefinitely for locks in SQL Server. You can change this behavior by using the SET LOCK_TIMEOUT command, which limits the number of seconds that a process will wait for a lock before timing out.

Lock Wait Time

A single blocking user has the potential to stop work for nearly all other processes on a small system, and can cause major headaches even on large systems. Blocks are most often caused by user processes holding exclusive locks and not releasing them via a proper COMMIT frequency. Unless a process times out via an application timeout mechanism, or the process has specified a timeout period via the SET LOCK_TIMEOUT command, a process waiting for a lock will wait indefinitely. The lock wait time statistic represents the wait time (in milliseconds) that processes have had to wait to obtain a lock request.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	Any indicator above zero might indicate the presence of lock contention on the server.
Correction	Once discovered, a blocking lock situation can normally be quickly remedied - the DBA issues a KILL against the offending process, which eliminates the user's stranglehold on the objects they were accessing. Other user processes then nearly almost always complete in an instant. Discovering the blocked lock situation is made easier by using tools like Performance Analyst, but preventing the blocking lock situation in the first place is where it gets tricky. The DBA can drill down into user detail and view all current blocking locks to see exactly which sessions are holding the currently restrictive locks. The culprit of blocking lock scenarios is usually the application design, or the SQL being used within the application itself. Properly coding an application to reference database objects in an efficient order, and then using the right SQL to get the job done, is an art. The key to avoiding lock contention is to process user transactions in the quickest and most efficient manner possible - something not always easy to do. By default, all processes wait indefinitely for locks in SQL Server. You can change this behavior by using the SET LOCK_TIMEOUT command, which limits the number of seconds that a process will wait for a lock before timing out.

Scan Point Revalidations

The Scan Point Revalidations statistic represents the number of times per second that the scan point had to be revalidated to continue the scan.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Home > Bottleneck Analysis pane

SQL Re-Compilations

The SQL re-compilations statistic represents the total number of recompiles triggered per second in a SQL Server instance. Recompiles occur when SQL Server determines that the currently defined execution plan for an executing stored procedure might no longer be the best possible plan. SQL Server pauses the query execution and recompiles the stored procedure.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	Recompiles slow down the process that is executing the procedure and increase the load on the CPU. By extension, the more recompiles that are occurring on your system, the more overall load increases, resulting in poor performance. In general, you want to keep the number of recompiles low. The most common reasons SQL Server would issue a recompile are: Running sp_recompile against any table referenced in the stored procedure, significant data changes in a referenced table, schema changes to referenced objects, the use of the WITH RECOMPILE clause in the CREATE PROCEDURE or EXECUTE statement, and a plan no longer available in the system cache.
Correction	Try to practice coding standards that eliminate the most frequent causes detailed above. Also, try to: *Use temporary tables only in the stored procedure that created them.* *Minimize creating temporary tables in control block structures.* *Use the KEEP PLAN option on references to static temporary tables.* *Issue the CREATE TABLE statement before any other references to the created table.* *Minimize the use of temporary tables.*

Top Bottlenecked User Processes

When viewing wait statistics, there are several levels of detail that a DBA can drill down into. The first level is the system view, which provides a global, cumulative snapshot of all the waits that have occurred on a system. The second level is the session view, which shows details on what events connected sessions are experiencing. Viewing these numbers can help a DBA determine which session-related wait events have caused the most commotion in a database thus far. The Top Bottlenecked User Processes display identifies the processes with the most current wait time along with the event causing their current wait.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Home > Bottleneck Analysis pane

SQL Analysis pane

A lot of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured server down in terms of end-user response times.

Before you can identify problem SQL in your database, you have to ask the question - "What is bad SQL?" What criteria do you use when you begin the hunt for problem SQL in your critical systems? Understand that even the seasoned experts disagree on what constitutes efficient and inefficient SQL; so there is no way to sufficiently answer this question to every SQL Server professional's satisfaction. The table below lists some general criteria you can use when evaluating the output from various database monitors or personal diagnostic scripts:

Criteria	Description
Reads (Physical)	This is often used as the major statistic in terms of identifying good vs. bad SQL, this is a measure of how many disk reads the query caused to satisfy the user's request. While you certainly want to control disk I/O where possible, it is important that you not focus solely on physical I/O as the single benchmark of inefficient SQL. Make no mistake, disk access is slower than memory access and also consumes processing time making the physical to logical transition, but you need to look at the entire I/O picture of a SQL statement, which includes looking at a statements' logical I/O as well.
Writes	These can be caused by DML activity (INSERTs, etc.), but writes can also be indicative of heavy disk sort activity.
CPU Time	This is how much CPU time the query took to parse, execute, and fetch the data needed to satisfy the query.
Elapsed Time	This is how much time the query took to parse, execute, and fetch the data needed to satisfy the query. It should not include the network time needed to make the round trip from the requesting client workstation to the database server.

There are other criteria that you can examine like sort activity or access plan statistics (that show items like Cartesian joins and the like), but more often than not, these measures are reflected in the criteria listed above.

The SQL Analysis section provides you with a quick overview of the percentage of Reads, Writes, CPU, and Elapsed Time the most expensive SQL statements tracked by Performance Analyst have used. Depending on the page you are on, you might just see information regarding physical I/O activity or DML statements. For example, you might see that a SQL statement has caused 60% of all physical reads on the system, or that a procedure is responsible for 90% of all the CPU usage. To get information regarding the actual SQL text or stored procedure calls, drill down into the Top SQL details views.

> **NOTE:** SQL Profiling is turned on by default in Performance Analyst so you can automatically collect SQL traffic. If you do not want to collect SQL execution information, use the options inside of Performance Analyst to disable SQL profiling.

Metrics

When you begin to look for inefficient SQL in a database, there are two primary questions you should answer:

- What has been the worst SQL that has historically been run in my database?

- What is the worst SQL that is running right now in my database?

When Correction a slow system, you should be on the lookout for any query that shows an execution count that is significantly larger than any other query on the system. It might be that the query is in an inefficient Transaction SQL loop, or other problematic programming construct. Only by bringing the query to the attention of the application developers will you know if the query is being mishandled from a programming standpoint.

There is the possibility that the SQL statement just is not tuned well. To determine that, you can drill down further into the Performance Analyst Top SQL view and begin working with the query through EXPLAIN plan analysis and other techniques inside Embarcadero's DBArtisan.

Login Analysis

The Login Analysis section displays the number of successful login and logout operations.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Home > SQL Analysis pane

Workload Analysis pane

When the database population as a whole experiences a system slowdown, it is not uncommon to find one or two users who are responsible for bringing the system to its knees. In the best of worlds, users will have a pretty evenly divided amount of memory usage, disk I/O, CPU utilization, and parse activity. Unfortunately, this usually is not the case, and many times you will find users submitting large batch jobs during peak OLTP activity, or sessions that are firing off untuned queries on a critical system.

If you are seeing a slowdown in your database, and cannot seem to find a root cause, one thing to examine is the resource consumption of the leading sessions on a system. Oftentimes you will find one or a handful of users making things miserable for everyone else. They accomplish this by overwhelming the I/O capabilities of SQL Server (through untuned queries or runaway batch jobs) or hammering the CPU or memory structures of the database and server.

Performance Analyst makes it easy to pinpoint the top sessions by showing the leading processes at the physical I/O, memory, CPU, and transaction submission usage levels.

Metrics

If any one session appears to be using more than 50% of a total resource (CPU, memory, etc.), then you should drill down into the session to find out what they are currently executing.

Top Sessions Tab

The Top Sessions tab of the Users Detail includes the following processes:

Top CPU Process	Top I/O Process	Top Memory Process
Top Sessions		

Top I/O Process

The Top I/O Process statistic identifies the SQL Server process that currently has caused the most I/O usage on the database.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Home > Bottleneck Analysis pane

>

Top Memory Process

The Top Memory Process statistic identifies the SQL Server process that currently is using the highest percentage of memory in the database.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Home > Bottleneck Analysis pane

Top CPU Process

The Top CPU Process statistic identifies the SQL Server process that currently has the highest percentage of CPU usage in the database.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Home > Bottleneck Analysis pane

Top Sessions

When a system experiences heavy activity, sometimes you will find that all the user connections are contributing somewhat equally to the overall load. More often than not, however, one or two user connections are causing most of the activity. It might be that a large batch load or other typical process is running that is perfectly okay for your system. Or it might be a runaway process or other rogue connection that needs to be tracked down and possibly eliminated.

Top Sessions displays information regarding all key statistics for all current sessions on a SQL Server. The table below describes the information available in the Top Sessions section on the Top Sessions tab of the Users Detail:

Information	Description
SPID	The unique identifier for the process.
Login	The login name for the process.
Windows User	The operating system name of the process.
Database	The name of the database the process is working in.
Status	The current status of the process.
Program	This identifies the program being run by the process.
Memory	The number of pages in the procedure/SQL cache allocated to the process.
CPU	The cumulative CPU time for the process.
Physical I/O	The cumulative total of all physical reads and writes for the process.
Blocked	This indicates if the process is blocked by another process.
Host	This identifies the workstation of the process.
Open Transaction	This indicates if the process has an open transaction.
Command	The command being currently issued by the process.
Login Time	This identifies when the process logged into SQL Server.
Last Batch	This indicates the date/time the process executed a remote stored procedure or issued an EXEC command.
Host Process	The process identifier on the SQL Server machine.
Net Address	The network address of the process.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Home > Bottleneck Analysis pane

>

Current SQL

This view shows a snippet of the SQL associated with the user in question.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Home > Bottleneck Analysis pane

System Waits Tab

The System Waits tab of the Users Detail includes the following sections:

System Waits	Wait Percent by Time Waited	Wait Percent by Total Waits

System Waits

When SQL Server is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting may mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. The System Waits section identifies the waits that have occurred on SQL Server. The System Waits section identifies the waits that have occurred on SQL Server. The table below describes the information available in the System Waits section on the System Waits tab of the Users Detail:

Information	Description
Wait Type	The actual wait event name.
Requests	The number of waits.
Percent of Total	The percentage of the occurrences of this wait to all wait requests.
Time Waited	The number of seconds SQL Server spent waiting for this wait.
Percent Time Waited	The percentage of time SQL Server spent waiting on this wait.
Signal Wait Time	The signal wait time for the wait.
Percent Signal Wait Time	The percentage of signal wait time SQL Server spent waiting on this wait.

Location	Microsoft SQL Server Performance Analyst statistics > Users > System Waits Tab
Metrics	Wait events can be hard to interpret at times. If you see a particular event that has caused a lot of wait time, you can review the information in this link (Microsoft Knowledge Base Article - 244455) to help understand the cause and potential remedy: http://support.microsoft.com/default.aspx?scid=kb;en-us;Q244455

Correction

The table below describes some possible solutions to some of the most common wait events:

Wait Event	Description
CURSOR	This indicates SQL is waiting to sync up with asynchronous cursors and can point to the excessive use of cursors.
CMEMTHREAD	This indicates waits for memory to be freed up for use.
CXPACKET	This relates to SQL Server parallelism. Oftentimes the cost of SQL Server using parallelism for a query is high, with the end result being these types of waits. If high, you can adjust SQL Server and turn parallelism off for the server or for just low cost queries.
EXCHANGE	Related to CXPACKET. See information above.

>

Wait Event	Description
IO_COMPLETION	This means certain I/O related actions are being delayed. Bulk insert operations, growth of database or log files, high physical I/O SQL, page splitting, and server paging are potential causes. You can check Performance Analyst's SQL Analysis section to ferret out heavy I/O SQL and the Database and File I/O Detail pages (SQL Server 2000 and later) to examine disk and database hot spots with respect to I/O. The I/O Stall column is especially useful in determining which files are experiencing the most delays.
LOGMGR	This refers to waits for the Log Writer to start writing a transaction. High waits here might warrant transfer transaction logs to faster devices or breaking up long running DML transactions.
OLEDB	This indicates waits for an OLE DB operation to act on its requests. Slow connection speeds or very high transaction rates can cause these.
PAGEIOLATCH	Related to IO_COMPLETION. See information above.
PAGELATCH	Related to IO_COMPLETION. See information above.
PAGESUPP	Related to CXPACKET. See information above.
PSS_CHILD	This is SQL waiting for a child thread within an asynchronous cursor and can point to the excessive use of cursors.
TEMPOBJ	This wait occurs when temp tables and the like are dropped.
WAITFOR	Waits caused by the WAITFOR T-SQL command. Not a cause for concern.
WRITELOG	Related to LOGMGR. Refers to waits for writes to disk of transactions. See information above.
XACTLOCKINFO	Involves the release of escalation of locks during bulk operations.

Wait Percent by Total Waits

When SQL Server is running, every connected process is either busy doing work or waiting for work. A process that is waiting might mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. The Wait Percent by Total Waits section graphically depicts the top waits as a percentage of all waits.

Location	Microsoft SQL Server Performance Analyst statistics > Users > System Waits Tab
Correction	Wait events can be hard to interpret at times. If you see a particular event that has caused a lot of wait time, you can review the information in this link (Microsoft Knowledge Base Article - 244455) to help understand the cause and potential remedy: http://support.microsoft.com/default.aspx?scid=kb;en-us;Q244455

Wait Percent by Time Waited

When SQL Server is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting might mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. The Wait Percent by Time Waited section graphically depicts the top waits by their percentage of time waited to total time waited. 5

Location	Microsoft SQL Server Performance Analyst statistics > Users > System Waits Tab
Correction	Wait events can be hard to interpret at times. If you see a particular event that has caused a lot of wait time, you can review the information in this link (Microsoft Knowledge Base Article - 244455) to help understand the cause and potential remedy: http://support.microsoft.com/default.aspx?scid=kb;en-us;Q24445

>

Session Waits Tab

The Session Waits tab of the Users Detail includes the following sections:

Session Wait Percent by Time Waited	Session Wait Percent by Wait Type	Session Waits

Session Waits

When SQL Server is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting may mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. The Session Waits section shows processes that are currently waiting on SQL Server. The table below describes the information available in the Session Waits section on the Session Waits tab of the Users Detail:

Information	Description
SPID	The unique identifier for the process.
Login Name	The login name for the process.
Windows User	The operating system name of the process.
Status	The current status of the process.
Host	This identifies the workstation of the process.
Program	This identifies the program being run by the process.
Last Wait Type	The actual wait event name.
Wait Time	The number of seconds the process spent waiting for this wait.
Percent Total Wait Time	The percentage of wait time for this process vs. total processes.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Session Waits Tab
Metrics	Wait events can be hard to interpret at times. If you see a particular event that has caused a lot of wait time, you can review the information in this link (Microsoft Knowledge Base Article - 244455) to help understand the cause and potential remedy: http://support.microsoft.com/default.aspx?scid=kb;en-us;Q244455
Correction	he table below describes some possible solutions to some of the most common wait events:

T

Wait Event	Description
CURSOR	This indicates SQL is waiting to sync up with asynchronous cursors and can point to the excessive use of cursors.
CMEMTHREAD	This indicates waits for memory to be freed up for use.
CXPACKET	This relates to SQL Server parallelism. Oftentimes the cost of SQL Server using parallelism for a query is high, with the end result being these types of waits. If high, you can adjust SQL Server and turn parallelism off for the server or for just low cost queries.
EXCHANGE	Related to CXPACKET. See information above.
IO_COMPLETION	This means certain I/O related actions are being delayed. Bulk insert operations, growth of database or log files, high physical I/O SQL, page splitting, and server paging are potential causes. You can check Performance Analyst's SQL Analysis section to ferret out heavy I/O SQL and the Database and File I/O Detail pages (SQL Server 2000 and later) to examine disk and database hot spots with respect to I/O. The I/O Stall column is especially useful in determining which files are experiencing the most delays.

Wait Event	Description
LOGMGR	This refers to waits for the Log Writer to start writing a transaction. High waits here might warrant transfer transaction logs to faster devices or breaking up long running DML transactions.
OLEDB	This indicates waits for an OLE DB operation to act on its requests. Slow connection speeds or very high transaction rates can cause these.
PAGEIOLATCH	Related to IO_COMPLETION. See information above.
PAGELATCH	Related to IO_COMPLETION. See information above.
PAGESUPP	Related to CXPACKET. See information above.
PSS_CHILD	This is SQL waiting for a child thread within an asynchronous cursor and can point to the excessive use of cursors.
TEMPOBJ	This wait occurs when temp tables and the like are dropped.
WAITFOR	Waits caused by the WAITFOR T-SQL command. Not a cause for concern.
WRITELOG	Related to LOGMGR. Refers to waits for writes to disk of transactions. See information above.
XACTLOCKINFO	Involves the release of escalation of locks during bulk operations.

Session Wait Percent by Wait Type

When SQL Server is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting might mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. The Session Wait Percent by Wait Type section graphically depicts the top waits by wait type.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Session Waits Tab
Correction	Wait events can be hard to interpret at times. If you see a particular event that has caused a lot of wait time, you can review the information in this link (Microsoft Knowledge Base Article - 244455) to help understand the cause and potential remedy: http://support.microsoft.com/default.aspx?scid=kb;en-us;Q244455

Session Wait Percent by Time Waited

When SQL Server is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting might mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. The Session Wait Percent by Time Waited section graphically depicts the top session waits by their percentage of time waited to total time waited.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Session Waits Tab
Correction	Wait events can be hard to interpret at times. If you see a particular event that has caused a lot of wait time, you can review the information in this Microsoft Knowledge Base Article to help understand the cause and potential remedy: http://support.microsoft.com/default.aspx?scid=kb;en-us;Q244455

Locks Tab

The Locks tab of the Users Detail includes the following sections:

Locks	Locks by Lock Type	Locks by Database

Locks

To modify database information or structures, a user session must obtain a lock on the object to perform its task. In addition to user locks, SQL Server itself issues lock requests to carry out its internal duties. The Locks section gives information the locks currently on the system and also indicates if any blocking situations are occurring. The table below describes the information available in the Locks section:

Information	Description
SPID	The process id of the process holding the lock.
Login	The login name of the process.
Windows User	The operating system name of the process.
Database	The database in which the process is running.
Table Name	The name of the table involved in a lock. This will be NULL for non-table locks or table locks that take place in the tempdb database.
Index ID	The index ID involved in the lock.
Lock Type	The type of the lock (database, table, row id, etc.)
Lock Mode	The lock's mode (shared, exclusive, etc.)
Lock Status	The lock's status (waiting or granted).
Lock Owner Type	Whether the lock came from a regular session or a transaction.
User Program	The executable the process is using against the server.
Blocking SPID	If zero, the process is not being blocked. If non-zero, this column represents the process ID of the process blocking the requested lock.
Wait Time	The current amount of wait time for the process, in milliseconds.
SPID Status	Indicates if the process is actively performing work, is idle, blocked by another process, etc.
SPID Command	The command the process is currently issuing.
NT Domain	The name of Windows 2000/NT domain.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Locks Tab
Metrics	Once discovered, a blocking lock situation can normally be quickly remedied - the DBA issues a KILL against the offending process, which eliminates the stranglehold on the objects the user was accessing. Other user processes then almost always complete in an instant. Discovering the blocked lock situation is made easier by using tools like Embarcadero Performance Analyst, but preventing the blocking lock situation in the first place is tricky. The culprit of blocking lock scenarios is usually the application design, or the SQL being used within the application itself. Properly coding an application to reference database objects in an efficient order, and then using the right SQL to get the job done, is an art. The key to avoiding lock contention is to process user transactions in the quickest and most efficient manner possible - something not always easy to do. By default, all processes wait indefinitely for locks in SQL Server. You can change this behavior by using the SET LOCK_TIMEOUT command, which limits the number of seconds that a process waits for a lock before timing out.

>

Locks by Lock Type

To modify database information or structures, a user session must obtain a lock on the object to perform its task. In addition to user locks, SQL Server itself issues lock requests to carry out its internal duties. The Locks by Lock Type section of the Users Detail provides a graphical percentage breakdown of the different locks types being issued across all databases.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Locks Tab

Locks by Database

The Locks by Database section of the Lock tab provides a graphical representation of the percentage breakdown of locks by database.

Location	Microsoft SQL Server Performance Analyst statistics > Users > Locks Tab

Database

In addition to a Home page, the Database category of SQL Server Performance Analyst includes the following tabbed pages:

Backups Tab	Configuration Tab	Error Log Tab
Overview Tab	Replication Tab	SQL Agent Tab
Objects Tab		

Home

The Home page of the Database Performance Analyst page includes the following categories:

Bottleneck Analysis pane	Database Analysis pane	SQL Analysis pane
Workload Analysis pane		

Database Analysis pane

SQL Server is broken up into many different databases, all of which have their own levels of activity. The Database Analysis section displays a summary of activities that have occurred within each database. The table below describes the information available in the Database Analysis section:

Information	Description
Database Name	The name of the database.
DBCC Logical Scans	The number of DBCC scans that have occurred within the database.
Transactions	The number of transactions that have occurred within the database.
Active Transactions	The number of transactions currently running in the database.
Bulk Copy Rows	The number of rows inserted per second into the database via BCP operations.
Bulk Copy Throughput	The amount of data bulk (KB) copied into the database per second.
Log Cache Reads	The number of reads per second that have gone through the log manager cache.

>

Information	Description
Log Flushes	The number of transaction log flushes per second. A log flush occurs when SQL Server writes all changes from the database's log cache out to the database's log files on disk.
Log Growths	The number of times the transaction log has grown in size.
Log Shrinks	The number of times the transaction log has reduced itself in size.

Bottleneck Analysis pane

When using bottleneck analysis to troubleshoot a database, a DBA cannot rely only on the information contained in the wait event views that SQL Server provides. For example, a critical database might be offline or be marked suspect. Such a failure will not be reflected in any wait event, but still represents a very real bottleneck to the database. In the same way that a DBA cannot depend on only a few ratios to properly carry out ratio-based performance analysis, an administrator must include other statistical metrics in their overall bottleneck analysis framework to obtain an accurate performance risk assessment. Performance Analyst works to identify bottlenecks in your database that fall outside of pure wait events so you can get a total picture of all stoppages in your system. The Performance Analyst Database home page displays database-related bottlenecks that might be dragging down the overall performance of your system.

The following bottleneck indicators are used on the Performance Analyst Database home page to communicate the general overall conditions of databases on the monitored server:

Databases Requiring Backup	Databases Without Auto-Create Stats	Databases Without Auto-Update Stats
Extents Allocated	Freespace Scans	Page Splits

The following statistics, available on this page, duplicate statistics available on the Bottleneck Analysis pane of the Home page:

Suspect Databases	Offline Databases	

Databases Requiring Backup

This metric represents the number of databases in SQL Server that have not been backed up for more than seven days. This statistic excludes the pubs, tempdb, Northwind, msdb, and model databases.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Home > Bottleneck Analysis pane
Metrics	To ensure proper protection for most databases, it is recommended that even the most static databases be backed up at least once a week. You should frequently backup critical, dynamic databases and include transaction log backups to enable point-in-time recovery ability.
Correction	Any critical databases found with obsolete or no backups should immediately be backed up. Moreover, to ensure proper data protection for each database, you should institute a planned backup schedule. The timing and repetition of the backups depend on the critical recovery needs of each database.

Databases Without Auto-Create Stats

The Databases Without Auto-Create Stats statistic represents the total number of databases defined on SQL Server that do not have the AUTO_CREATE_STATISTICS option enabled.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Home > Bottleneck Analysis pane

Metrics	When the AUTO_CREATE_STATISTICS option is enabled, statistics are automatically created on columns used in a SQL query predicate. Keeping object statistics fresh in the SQL Server data dictionary improves query performance because the optimizer can better determine how to evaluate a query and return the requested data. If the statistics are not used, SQL Server should automatically delete them.
Correction	If possible, a DBA should keep their databases in AUTO_CREATE_STATISTICS mode. If a database is found without this option set, you can easily change it by using the command: EXEC sp_dboption '<database name>','auto create statistics',true

Databases Without Auto-Update Stats

The Databases Without Auto-Update Statistics statistic represents the total number of databases defined on SQL Server that do not have the AUTO_UPDATE_STATISTICS option enabled.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Home > Bottleneck Analysis pane
Metrics	When the AUTO_UPDATE_STATISTICS option is enabled, SQL Server automatically updates existing statistics when the statistics become out-of-date because data in the tables has changed enough to affect the optimizer's decision-making process.
Correction	If possible, a DBA should keep databases in AUTO_UPDATE_STATISTICS mode. If a database is found without this option set, you can easily change its auto-update statistics to true by using the command: EXEC sp_dboption '(database name)',

Extents Allocated

SQL Server groups eight 8KB pages together to form an extent of space (64KB). Space is allocated to database objects (tables, indexes) in the form of extents. The Extents Allocated statistic provides a count of the number of extents allocated to tables and indexes.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Home > Bottleneck Analysis pane
Metrics	None. However, be aware that excessive counts of allocated extents might signal a large load job occurring in the database.

Freespace Scans

The Freespace Scans statistic provides a count of the number of scans SQL Server initiated to search for free space that will accommodate requests to insert new data rows.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Home > Bottleneck Analysis pane
Metrics	None. However, be aware that excessive counts of freespace scans might signal INSERT problems being performed on large tables.

Page Splits

When data is added or modified for a table that has indexes, the indexes must be updated as well. As index pages fill up, free space needed to keep index keys in their proper order can oftentimes run very low. If an index update needs to be made to a page, and no free space exists on that page, SQL Server will perform a page split where it moves approximately half the rows on a page to a new page. The Page Splits statistic provides a count of how many page splits are occurring in SQL Server.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Home > Bottleneck Analysis pane

Metrics	Page splits cause additional overhead in the form of CPU usage and I/O. Observing large numbers of page splits can signal a resource bottleneck in your server.
Correction	To avoid page splits, you can look into tuning the FILLFACTOR property of an index, which controls the percentage of the index page that is filled during creation. The default, 100, tells SQL Server to completely fill each page, whereas lower numbers tell SQL Server to leave room for additional index rows or updates to existing rows.

SQL Analysis pane

Certain databases on a SQL Server will experience more SQL activity than others. The SQL Analysis section displays a count of all statement executions for each database.

The SQL Analysis section provides you with a quick overview of the percentage of Reads, Writes, CPU, and Elapsed Time the most expensive SQL statements tracked by Performance Analyst have used. Depending on the page you are on, you might just see information regarding physical I/O activity or DML statements. For example, you might see that a SQL statement has caused 60% of all physical reads on the system, or that a procedure is responsible for 90% of all the CPU usage. To get information regarding the actual SQL text or stored procedure calls, drill down into the Top SQL details views.

NOTE: SQL Profiling is turned on by default in Performance Analyst so you can automatically collect SQL traffic. If you do not want to collect SQL execution information, use the options inside of Performance Analyst to disable SQL profiling.

Error Log Analysis pane

SQL Server records various system events in its system or error log. The majority of messages placed into the log are informational in nature, however since some contain critical messages, you should immediately review them. This section indicates the number of actual error messages in the most recent error log so you know if there are potential events that require your attention.

NOTE: Similar information is also available on the Bottleneck Analysis pane of the Home page.

Metrics

You should investigate any positive values for errors.

Correction

If non-zero values are found for this statistic, you should drill down into the current error log and examine the detail found that accompanies each error issued by SQL Server.

Workload Analysis pane

SQL Server is broken up into many different databases, all of which have their own levels of activity. The Workload Analysis section displays a summary of resource consumption across all databases. The table below describes the information available in the Workload Analysis section:

Information	Description
Database Name	The name of the database.
Connections	The number of processes that are connected to the database.
Blocked Users	The number of processes in the database experiencing blocks from lock activity.
Total Memory	The total number of memory pages in use by processes for each database.
Total I/O	The total amount of physical I/O being experienced by each database.

>

Information	Description
Total CPU	The total amount of cumulative CPU for all processes connected to the database.
Total Wait Time	The total wait time (in milliseconds) that all processes connected to the database have experienced.
DBCC Running Count	This indicates if any DBCC operations are currently running against the database.
BCP Running	This indicates if any BCP operations are currently running against the database.
Backup Restore Running	This indicates if any backup or restore operations are currently running against the database.

Overview Tab

The Overview tab of the Database Detail includes the following sections:

Database Object Analysis	Database Size Analysis	Database Summary
Database User Analysis		

Database Size Analysis

The Database Size Analysis chart displays each database in terms of its total percentage of size to the total size of all databases on the SQL Server.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Home > Bottleneck Analysis pane

Database Object Analysis

The Database Object Analysis chart displays each database in terms of its total percentage of object to the total objects for all databases on the SQL Server.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Home > Bottleneck Analysis pane

Database User Analysis

The Database User Analysis chart displays each database in terms of its total percentage of user accounts to the total user accounts of all databases on the SQL Server.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Home > Bottleneck Analysis pane

Database Summary

The Database Summary Analysis section displays summarized metadata for all defined databases on SQL Server. The table below describes the information available in the Database Summary section:

Information	Description
Database	The name of the database.
Created	The date/time when the database was created.
Status	The status of the database (online, offline, suspect, etc.)
Database Size (MB)	The total size of the database in MB.

>

Information	Description
Log Size (MB)	The total size of the log in MB.
Tables	The number of all tables in the database.
Indexes	The number of all indexes in the database.
Users	The number of all defined user accounts in the database.
Last Backup	The date/time of the last full backup for the database.
Autoshrink	This indicates if the database will release space automatically back to the operating system when it does not need it.
DB Autogrow	This indicates if one or more of the database's underlying files can automatically grow in size if needed.
Log Autogrow	This indicates if one or more of the log's underlying files can automatically grow in size if needed.
Autocreate Stats	This indicates if object statistics are automatically created for objects in the database.
Autoupdate Stats	This indicates if object statistics are automatically updated for objects in the database.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Home > Bottleneck Analysis pane
Metrics	The following are things to take note of: *Any critical database that shows a last full backup date that is older than the database's required backup needs.* *Any database that shows a status of offline or suspect.* *Any growing database that does not have its database or log files set to automatically grow in size.* *Any dynamic database that does not have its object statistics set to automatically update.*

Correction	Depending on the situation, you should take the following actions: For databases that require a full backup, perform a full backup of the database when appropriate. For suspect databases, the steps to handling will vary from one installation to another. However, here are some general guidelines you can use to troubleshoot a suspect database: *Begin by examining the SQL Server error log for clues as to what caused the database to be marked as suspect.* *It is not unusual for a server to run out of physical disk space on drives used by SQL Server. When this happens, recovery for databases can sometimes fail with the end result being SQL Server marking a database as suspect. To remedy this situation, you should free up space on the identified drives or add files to the newly marked suspect database. For SQL Server 2000, this can be accomplished by utilizing the following stored procedures:* sp_add_data_file_recover_suspect_db *and* sp_add_log_file_recover_suspect_db*. For SQL Server version 7.0, you will need to use the* sp_resetstatus *stored procedure to reset the suspect status flag for the database in question, use the* alter database *command to add new datafiles to the database, and then stop/start SQL Server.* Many times, suspect databases are caused by SQL Server not being able to access a database or log file. This happens if a particular physical hard drive has become unavailable, but also can occur if another operating system process has obtained exclusive access to a file. If this scenario proves to be true, once you have ensured that the file(s) are available once again to the operating system, you can use the sp_resetstatus stored procedure to reset the suspect status flag for the database and then stop/start SQL Server. Should an offline database be found by Performance Analyst, you can easily place it back online by utilizing either the sp_dboption stored procedure or the alter database command. If any critical, dynamic database or log is found to not have their auto-growth feature enabled, then the DBA can modify the file(s) used by the database to automatically grow by using the ALTER DATABASE … MODIFY FILE … FILEGROWTH command. You should also ensure that the MAXSIZE setting for each file is set appropriately. If possible, a DBA should keep their databases in AUTO_CREATE_STATISTICS and AUTO_UPDATE_STATISTICS mode. If a database is found without this option set, you can easily change it by using the following commands: *EXEC sp_dboption '<database name>','auto create statistics',true* *EXEC sp_dboption '<database name>','auto update statistics',true*

Error Log Tab

The Error Log tab of the Database Detail includes the following sections:

SQL Server Error Log	Error Log Content Analysis	Error Log Volume Analysis

SQL Server Error Log

SQL Server logs events pertinent to its operations to the SQL Server error log and the operating system application log. Errors encountered by SQL Server are clearly visible in the log, however large logs can be difficult to navigate. The SQL Server Error Log section displays all events for a selected log or just errors.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Error Log Tab
Metrics	You should investigate any positive values reported for errors.

>

Error Log Volume Analysis

SQL Server logs events pertinent to its operations to the SQL Server error log and the operating system application log. Errors encountered by SQL Server as well as informational message are recorded in the log. The Error Log Volume Analysis section displays a graph that shows the volume amount of messages/errors over the last five days.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Error Log Tab
Metrics	Seeing a disproportionate percentage of messages for one day might indicate that something has, or is occurring that needs investigation.

Error Log Content Analysis

SQL Server logs events pertinent to its operations to the error log and the operating system application log. Errors encountered by SQL Server as well as informational message are recorded in the log. The Error Log Content Analysis section displays a graph that shows percentage of informational messages to errors in the current log.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Error Log Tab
Metrics	You should investigate any indicated errors.

Objects Tab

The following statistics, available on this tab, are duplicates of statistics available on the Objects Tab of the Space page:

Database Object Summary	Database Object Detail	

SQL Agent Tab

The SQL Agent tab of the Database Detail includes the following sections:

Job/Alert Detail	SQL Server Alert Summary	SQL Server Job Summary

Job/Alert Detail

SQL Server provides the ability to submit and run jobs as well as be notified about certain SQL Server-related events. The Job/Alert Detail section displays information regarding the status of the server's SQL Agent as well as details for all jobs and alerts that are defined to the system.

Location	Microsoft SQL Server Performance Analyst statistics > Database > SQL Agent Tab

SQL Server Job Summary

The SQL Server Job Summary section graphically displays the outcome of all jobs for the last two days, as well as current activity.

Location	Microsoft SQL Server Performance Analyst statistics > Database > SQL Agent Tab
Metrics	You should investigate any counts noted for failed jobs.

>

SQL Server Alert Summary

The SQL Server alert summary section graphically displays the number of times that performance and event alerts have fired.

Location	Microsoft SQL Server Performance Analyst statistics > Database > SQL Agent Tab
Metrics	You should investigate any counts noted for any event.

Backups Tab

The Backups tab of the Database Detail includes the following:

Backup Detail	Database Backup Summary	Log Backup Summary

Backup Detail

The Backup section displays the most recent 25 backups for a selected database. The table below describes the information available on the Backups tab of the Database Detail:

Information	Description
Database	The name of the database.
Backup Start	The time stamp when the backup began.
Backup Finish	The time stamp when the backup finished.
Backup Type	The type of backup (FULL, INCREMENTAL, etc.)
Backup Size (KB)	The size of the backup, in kilobytes.
Expiration Date	The expiration date for the backup, if any.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Backups Tab
Metrics	Although the needs of an application determine the frequency and type of backup, it is generally recommended that most dynamic databases have a solid plan in place for full and differential backups. For databases requiring point-in-time recovery, a backup plan should also include log backups.

Database Backup Summary

The Database Backup Summary section graphically displays the backup successes and failures for database backups over the last week.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Backups Tab
Metrics	You should investigate any backup failures.

Log Backup Summary

The Log Backup Summary section graphically displays the backup successes and failures for database log backups over the last week.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Backups Tab
Metrics	You should investigate any backup failures.

Configuration Tab

The Configuration tab of the Database Detail displays metadata regarding the hardware characteristics of the SQL Server as well as the SQL Server configuration parameters. The table below describes the information available on the Configuration tab:

Information	Description
Parameter	The parameter name.
Value	The value set for the parameter.
Run	The value currently running for the parameter.
Minimum	The minimum value allowed for the parameter.
Maximum	The maximum value allowed for the parameter.
Dynamic	Whether the parameter is dynamic (can be changed without requiring SQL Server to be stopped/started) or not.

Replication Tab

The Replication tab of the Database Detail provides basic information regarding replication activities that are occurring on the monitored SQL Server.

Replication Agent Details

The Replication Agent Details section shows information regarding replication. The table below describes the information available on the Replication tab:

Information	Description
Agent	The agent performing the replication work.
Publication	The object being published to a subscribing database.
Publisher	The SQL Server providing the source material.
Publisher Database	The database of the SQL Server providing the source material.
Status	The current status of the replication agent's task.
Subscriber	The SQL Server requesting the source material.
Subscriber Database	The database of the SQL Server requesting the source material.
Type	The type of replication activity (push, pull, etc.)
Start Time	The start time for the replication task.
Duration	The duration of the replication task.
Last Action	The output message for the replication task.

Location	Microsoft SQL Server Performance Analyst statistics > Database > Replication Tab

Replication Throughput Details

The Replication Throughput Details section of the Replication tab displays:

- Replication Type
- Measure

- Instance Name
- Value

Location	Microsoft SQL Server Performance Analyst statistics > Database > Replication Tab

OS

In many scenarios, an optimally tuned database may not perform well because there are constraints imposed by the system where the database is running. These constraints may include processes competing with the database sever for resources (CPU, I/O, or Memory), a slow CPU, insufficient or slow I/O devices, and insufficient memory. The OS Statistics page of Performance Analyst lets you examine operating system metrics for the following platforms:

- AIX
- HP-UX

  NOTE: To view processor info and swap disk info on an HP-UX box, you need to login as ROOT in the OS login.

- Linux
- Solaris
- Unix
- Windows XP and 2000

  NOTE: The statistics available on the OS page depend on the operating system platform.

  TIP: If you magnify a graph to change back to the original size or minimize the size, close the OS Page and then reopen it again. Performance Analyst displays the graph in its original form and size.

In addition to a Home page, the OS category of SQL Server Performance Analyst includes the following tabbed pages:

CPU Tab	I/O Tab	Memory Tab
Network Tab	Processes Tab	Space Tab

Home

The OS home page includes the following sections:

Key Resource Usage pane	Disk Analysis pane	Bottleneck Analysis pane
Memory Analysis pane	Workload Analysis pane	

Key Resource Usage pane

The following ratios are used on the Performance Analyst OS home page to communicate the general overall performance levels of the operating system:

Disk Time	Swap Memory Used	Processor Time
Swap Memory Used (Windows)	Swap Memory Used (AIX, HP-UX, Linux, Solaris, Unix)	

TIP: To open the CPU tab, right-click Processor Time. To open the I/O tab, right-click Disk Time. To open the Memory tab, right-click Paged Memory Used.

Disk Time

The Disk Time statistic is the percentage of elapsed time that the selected disk drive/device was busy servicing read or write requests.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	You should avoid consistently seeing values for this statistic greater then 90%.
Correction	Add more disk drives and partition the files among all of the drives.

Load Average

The Load Average statistic represents the system load averages over the last 1, 5, and 15 minutes.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	High load averages usually mean that the system is being used heavily and response time is correspondingly slow.

Processor Time

The Processor Time statistic indicates the percentage of time the processor is working. This counter is a primary indicator of processor activity.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	If your computer seems to be running sluggishly, this statistic could be displaying a high percentage.
Correction	Upgrade to a processor with a larger L2 cache, a faster processor, or install an additional processor.

Paged Memory Used

The Paged Memory Used statistic is the ratio of Commit Memory Bytes to the Commit Limit. Committed memory is where memory space has been reserved in the paging file if it needs to be written to disk. The commit limit is determined by the size of the paging file. As the paging file increases, so does the commit limit.

NOTE: This statistic is available for the Windows platform.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	This value displays the current percentage value only and not an average. If the percentage of paged memory used is above 90%, you may be running out of memory.
Correction	Increase the size of page file.

Swap Memory Used

The Swap Memory Used statistic is the percentage of swap space currently in use.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	If the percentage of swap memory used is above 90%, you may be running out of memory.
Correction	Increase the size of your swap files.

Bottleneck Analysis pane

The following ratios are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Average Disk Queue Length	Interrupts/Sec	Network Output Queue Length/Network Queue Length
Page Faults/Sec	Processor Queue Length	Processor Speed
Processor		

TIP: To open the I/O tab, right-click any Details menu item. To open the CPU tab, right-click the Item Processor Queues Length Details menu item. To open the Network tab, right-click Network Output Queues Length.

NOTE: The statistics available in this section depend on the platform of operating system.

Average Disk Queue Length

The Average Disk Queue Length statistic is the average number of both read and write requests that were queued for the selected disk during the sample interval.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Bottleneck Analysis pane
Metrics	This metric is useful in identifying I/O related bottlenecks. If the disk queue lengths for certain disks are consistently much higher than others, you may need to redistribute the load among available disks. If the disk queues lengths for all disks are consistently large, and you see a high amount of I/O activity, your disks may be inefficient.
Correction	Some things you can do if you have problems with this statistic include: *Redistribute the data on the disk with the large average disk queue to other disks.* *Upgrade to faster disk(s).*

Interrupts/Sec

Interrupts/Sec is the average rate, in incidents per second, at which the processor received and serviced hardware interrupts. This value is an indirect indicator of the activity of devices that generate interrupts, such as the system clock, the mouse, disk drivers, data communication lines, network interface cards, and other peripheral devices. These devices normally interrupt the processor when they have completed a task or require attention. Normal thread execution is suspended. The system clock typically interrupts the processor every 10 milliseconds, creating a background of interrupt activity. This statistic shows the difference between the values observed in the last two samples, divided by the duration of the sample interval.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Bottleneck Analysis pane Microsoft SQL Server Performance Analyst statistics > OS > CPU Tab > CPU Utilization

Metrics	A high value indicates possible excessive hardware interrupts; justification is dependent on device activity.

Network Output Queue Length/Network Queue Length

The Network Output Queue Length statistic is the number of threads in the processor queue.

NOTE: The name of this statistic depends on the platform of the operating system.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Bottleneck Analysis pane
Metrics	Unlike the disk counters, this counter shows ready threads only, not threads that are running. There is a single queue for processor time even on computers with multiple processors. Therefore, if a computer has multiple processors, you need to divide this value by the number of processors servicing the workload. A sustained processor queue of less than 10 threads per processor is normally acceptable, dependent of the workload.
Correction	A sustained high value in the Processor Queue Length could indicate that a processor bottleneck has developed due to threads of a process requiring more process cycles than are available. If this is the case, you should look at installing a faster (or an additional) processor.

Page Faults/Sec

The Page Faults/Sec statistic is the overall rate faulted pages are handled by the processor. It is measured in numbers of pages faulted per second. A page fault occurs when a process requires code or data that is not in its working set. This counter includes both hard faults and soft faults.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Bottleneck Analysis pane Microsoft SQL Server Performance Analyst statistics > OS > Home > Memory Tab
Metrics	This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
Correction	If the number of page faults remains consistently high, you can check with your Windows System Administrator for further investigation. Often, large numbers of page faults are not a problem so long as they are soft faults. However, hard faults, that require disk access, can cause delays.

Processor Queue Length

The Processor Queue Length statistic is the number of threads in the processor queue.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Bottleneck Analysis pane Microsoft SQL Server Performance Analyst statistics > OS > Home > CPU Tab
Metrics	Unlike the disk counters, this counter shows ready threads only, not threads that are running. There is a single queue for processor time even on computers with multiple processors. Therefore, if a computer has multiple processors, you need to divide this value by the number of processors servicing the workload. A sustained processor queue of less than 10 threads per processor is normally acceptable, dependent of the workload.
Correction	A sustained high value in the Processor Queue Length could indicate that a processor bottleneck has developed due to threads of a process requiring more process cycles than are available. If this is the case, you should look at installing a faster (or an additional) processor.

Processor Speed

The Processor Speed statistic displays the speed of the active processor in MHz. The speed is approximate.

Processor

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Bottleneck Analysis pane

The Processor Statistic displays the type of processor currently in use, for example, GenuineIntel.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Bottleneck Analysis pane

Memory Analysis pane

The following metrics are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Available Paged Memory (Windows)	Available Physical Memory	Available Swap Memory (AIX, HP-UX, Linux, Solaris, Unix)
Total Paged Memory/Total Swap Memory	Total Physical Memory	

TIP: To open the Memory tab, right-click any Details menu item.

NOTE: The statistics available in this section depend on the platform of operating system.

Available Paged Memory

The Available Paged Memory statistic shows the amount of virtual memory available for the processes.

NOTE: This statistic is available for the Windows platform.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	If the available virtual memory is less than 10% of the total virtual memory, your system may run out of memory.
Correction	Increase the size of page file.

Available Physical Memory

The Available Physical Memory statistic represents the amount of RAM available to all processes.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	This counter displays the last observed value only and not an average. Use this value with the Total physical memory and paging metrics (Memory details page). If the available physical memory is very small compared to this value, and the paging activity is high, your system may be running low on memory.
Correction	Some things you can do if you have problems with this statistic include: *Check the running processes to see if there are any memory leaks.* *Stop any services that are not required.* *Install additional RAM.*

Available Swap Memory

The Available Swap Memory statistic represents the amount of virtual memory available for the processes.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	If the available Available Swap Memory is less than 10% of total Swap Memory, your system may run out of memory.
Correction	Increase the size of swap files.

Total Paged Memory/Total Swap Memory

The Total Paged Memory statistic shows the maximum amount of virtual memory available to all processes.

NOTE: The name of this statistic depends on the platform of the operating system.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	It is recommended that this be 1½ - 3 times the amount of RAM on the system.

Total Physical Memory

The Total Physical Memory statistic shows the amount of physical memory installed on your computer.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	This is an informational metric and displays the total amount installed on the machine. Use this value with the available physical memory and paging metrics (Memory details page). If the available physical memory is very small compared to this value, and the paging activity is high, your system may be running low on memory.

Disk Analysis pane

The following ratios are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Free Disk Space	Total Disk Space	Used Disk Space

TIP: To open the Space tab, right-click any Details menu item.

NOTE: The statistics available in this section depend on the platform of operating system.

Free Disk Space

The Free Disk Space statistic shows the unallocated space, in megabytes on all logical disk drives.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Disk Analysis pane

Correction	There are many things you can do to ensure that a database does not encounter a space problem due to physical space limitations: *If a database currently resides on a disk that has little free space, you can add more files to the database. Of course, you should add the new files to other physical hard disks that can accommodate a growing database.* *You should examine hard disks with shrinking disk space to see if you can relocate or delete files to allow more free space.*

Total Disk Space

Total Disk Space displays the total allocated and unallocated space, in megabytes on all logical disk drives.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Disk Analysis pane
Correction	There are many things a DBA can do to ensure that a database does not encounter a space problem due to physical space limitations, here are two: *If a database currently resides on a disk that has little free space, you can add more files to the database. Of course, you should add the new files to other physical hard disks that can accommodate a growing database.* *You should examine hard disks with shrinking disk space to see if you can relocate or delete files to allow more free space.*

Used Disk Space

The Used Disk Space statistic shows the amount of allocated space, in megabytes on all logical disk drives.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Disk Analysis pane
Correction	There are many things a DBA can do to ensure that a database does not encounter a space problem due to physical space limitations: *If a database currently resides on a disk that has little free space, you can add more files to the database. Of course, you should add the new files to other physical hard disks that can accommodate a growing database.* *You should examine hard disks with shrinking disk space to see if you can relocate or delete files to allow more free space.*

Workload Analysis pane

The following ratios are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Number of Logins	Number of Processes	Top CPU Process
Top I/O Process	Location	

TIP: To open the Processes tab, right-click any Details menu item.

NOTE: The statistics available in this section depend on the platform of operating system.

>

Number of Logins

This statistic displays the total number of logins on the server.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Disk Analysis pane

Number of Processes

This statistic displays the total number of processes on the server.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Disk Analysis pane

Top Memory Process

Top Memory Process shows the current process that is consuming the most amount of memory. The information displayed is dependent on the platform of the operating system. Information displayed includes the name of the process, process ID, amount of memory consumed expressed in KB, amount of CPU expressed as a percentage, the amount of Major Page Faults, and the amount of I/O expressed in KB/sec.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Disk Analysis pane
Metrics	If you are running out of memory on the system, this is a quick way to identify the top memory user. If the displayed process is using a significant portion of the total memory, it could be causing the memory issues.

Top CPU Process

Top CPU Process shows the current process that is consuming the most amount of CPU. The information displayed is dependent on the platform of the operating system. Information displayed includes the name of the process, process ID, amount of memory consumed expressed in KB, amount of CPU expressed as a percentage, the amount of Major Page Faults, and the amount of I/O expressed in KB/sec.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Disk Analysis pane
Metrics	If the amount of CPU time used by this process is close to 100% and the CPU usage is very high, this process may be the bottleneck on the server.
Correction	Investigate the process further to see if it is in an inconsistent state. Also, look at minimum requirements for CPU speed for the process. You may need to upgrade your CPU.

Top I/O Process

The Top I/O Process statistic shows the current process that is consuming the most amount of CPU. The information displayed is dependent on the platform of the operating system. Information displayed includes the name of the process, process ID, amount of memory consumed expressed in KB, amount of CPU expressed as a percentage, the amount of Major Page Faults, and the amount of I/O expressed in KB/sec.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Home > Disk Analysis pane

CPU Tab

The CPU tab of the OS Detail includes the following sections:

CPU Events	CPU Utilization	Load Averages

>

The following statistics, available on this tab, duplicate statistics available on the Bottleneck Analysis pane of the OS Home page:

Interrupts/Sec	Processor Queue Length	

NOTE: The sections available on this tab depend on the platform of operating system.

CPU Utilization

The CPU Utilization section includes the following information:

% Privileged Time	% User Time	

% Privileged Time

This metric is the percentage of elapsed time that the process threads spent executing code in privileged mode.

NOTE: For Windows systems, when a Windows system service is called, the service will often run in privileged mode to gain access to system-private data. Such data is protected from access by threads executing in user mode. Calls to the system can be explicit or implicit, such as page faults or interrupts. These kernel commands, are considered privileged to keep the low-level commands executing and prevent a system freeze. Unlike some early operating systems, Windows uses process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. Some work done by Windows on behalf of the application might appear in other subsystem processes in addition to the privileged time in the process.

Location	Microsoft SQL Server Performance Analyst statistics > OS > CPU Tab
Metrics	The ideal range should be 0-40% (less than 40% indicates excessive system activity).
Correction	If your CPU consistently runs at less than 40% you may need to upgrade your system to include a faster processor(s).

% User Time

The % User Time statistic is the percentage of elapsed time the processor spends in the user mode. User mode is a restricted processing mode designed for applications, environment subsystems, and integral subsystems. The alternative, privileged mode, is designed for operating system components and allows direct access to hardware and all memory. The operating system switches application threads to privileged mode to access operating system services. This counter displays the average busy time as a percentage of the sample time.

Location	Microsoft SQL Server Performance Analyst statistics > OS > CPU Tab
Metrics	If the Privileged Time is high in conjunction with Physical Disk Reads, consider upgrading the disk I/O subsystem.

CPU Events

The CPU Events section includes the following information:

Context Switches/Sec	Load Averages	System Calls/Sec

NOTE: The statistics available in this section depend on the platform of operating system.

Context Switches/Sec

The Context Switches/Sec section shows the combined rate at which all processors on the computer are switched from one thread to another. Context switches occur when a running thread voluntarily relinquishes the processor, is preempted by a higher priority ready thread, or switches between user-mode and privileged (kernel) mode to use an Executive or subsystem service.

Location	Microsoft SQL Server Performance Analyst statistics > OS > CPU Tab
Metrics	The ideal range should be between 0-10,000. GA number greater then 10,000 may indicate too many threads contending for resources.

System Calls/Sec

System calls are functions which require the services of the operating system. Usually they involve access to some data that users must not be permitted to corrupt or even change.

Location	Microsoft SQL Server Performance Analyst statistics > OS > CPU Tab
Correction	Try to find a good ratio between the amount of time the operating system runs (fulfilling system calls and doing process scheduling), and the amount of time the application runs. Your goal should be running 60% to 75% of the time in application mode, and 25% to 40% of the time in operating system mode. If you find that the system is spending 50% of its time in each mode, then you should investigate to determine what is wrong.

Load Averages

The Load Average statistic represents the system load averages over the last 1, 5, and 15 minutes.

Location	Microsoft SQL Server Performance Analyst statistics > OS > CPU Tab
Metrics	High load averages usually mean that the system is being used heavily and the response time is correspondingly slow.

Processes Tab

The Processes tab of the OS Detail page succinctly communicates the general overall performance levels of processes. The columns available in this table depend on the platform of operating system. The table below describes the information available in the table on this tab:

Column	Description
Process	The name of the process.
User	The user of the process.
ID	The ID Process is the unique identifier of this process. ID Process numbers are reused, so they only identify a process for the lifetime of that process.
CPU	The CPU is the percentage of elapsed time that all of process threads used the processor to execution instructions.
User Mode	The User Mode is the percentage of elapsed time that the process threads spent executing code in user mode.
Memory **WINDOWS ONLY**	Memory is the current size, in bytes, of the virtual address space the process is using. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is finite, and the process can limit its ability to load libraries.

Column	Description
Memory (MB)	Memory is the current size, in bytes, of the virtual address space the process is using. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is finite, and the process can limit its ability to load libraries.
Memory	Memory is the percentage of the memory used of the total memory.
Active Memory	Active Memory is the amount of committed virtual memory, in bytes for this process. Active memory is the physical memory which has space reserved on the disk paging file(s). There can be one or more paging files on each physical drive. This counter displays the last observed value only; it is not an average.
I/O Data	The rate at which the process is reading and writing bytes in I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/Os.
Elapsed Time	The total elapsed time, in seconds, that this process has been running.
Thread Count	The number of threads currently active in this process. An instruction is the basic unit of execution in a processor, and a thread is the object that executes instructions. Every running process has at least one thread.
Handle Count	The total number of handles currently open by this process. This number is equal to the sum of the handles currently open by each thread in this process.
Priority	The current base priority of this process. Threads within a process can raise and lower their own base priority relative to the process' base priority.
Creating Proc ID	The Creating Process ID value is the Process ID of the process that created the process. The creating process may have terminated, so this value may no longer identify a running process.
Page Faults/Sec	Page Faults/Sec is the rate at which page faults by the threads executing in this process are occurring. A page fault occurs when a thread refers to a virtual memory page that is not in its working set in main memory. This may not cause the page to be fetched from disk if it is on the standby list and hence already in main memory, or if it is in use by another process with whom the page is shared.
Page File	Page File is the current number of kilobytes that this process has used in the paging file(s). Paging files are used to store pages of memory used by the process that are not contained in other files. Paging files are shared by all processes, and the lack of space in paging files can prevent other processes from allocating memory.
Private	Private is the current size, in kilobytes, of memory that this process has allocated that cannot be shared with other processes.

I/O Tab

The I/O tab includes the following sections:

Average Disk Queue/Total Disk Queue	LocBytes per I/O Operation	Disk I/O Time
Disk Transfers/Sec	I/O Details	Transfer Rate

NOTE: The sections available on this tab depend on the platform of operating system.

>

I/O Details

The table below describes the information available in this section:

Column	Description
Disk	The disk number assignment.
Reading (KB/s)	The amount of bytes read from the device.
Writing (KB/s)	The amount of bytes written to the device.
Disk Read Time	Disk Read Time is the percentage of elapsed time that the selected disk drive was busy servicing read requests.
Disk Write Time	Disk Write Time is the percentage of elapsed time that the selected disk drive was busy servicing write requests.
Disk Time	Disk Time is the percentage of elapsed time that the selected disk was busy servicing requests.
Avg. Read Queue	Avg. Disk Read Queue Length is the average number of read requests that were queued for the selected disk during the sample interval.
Avg. Write Queue	Avg. Disk Write Queue Length is the average number of write requests that were queued for the selected disk during the sample interval.
Disk Reads/Sec	Disk Reads/Sec is the rate of read operations on the disk.
Disk Writes/Sec	Disk Writes/Sec is the rate of write operations on the disk.

NOTE: The columns available in this table depend on the platform of operating system.

Location	Microsoft SQL Server Performance Analyst statistics > OS > I/O Tab

Transfer Rate

The Transfer Rate section succinctly communicates the general overall performance levels of the server's I/O. The table below describes the statistics in this section:

Statistic	Description	Metrics
Reading	The average number of bytes transferred from the disk during read operations.	It is measured in number of KBs per second.
Writing	The average number of bytes transferred to the disk during write operations.	It is measured in number of KBs per second.
Transferred	The amount of time required to transfer the data between the device's cache and the actual device media.	N/A

NOTE: The statistics available in this section depend on the platform of operating system.

>

LocBytes per I/O Operation

Location	Microsoft SQL Server Performance Analyst statistics > OS > I/O Tab

The Bytes per I/O Operation section of the OS Detail succinctly communicates the general overall performance levels of the server's I/O. The table below describes the statistics in this section:

Statistic	Description	Metrics
Reading	The average number of bytes transferred from the disk during read operations.	It is measured in number of bytes per second.
Writing	The average number of bytes transferred to the disk during write operations.	It is measured in number of bytes per second.

Location	Microsoft SQL Server Performance Analyst statistics > OS > I/O Tab

Average Disk Queue/Total Disk Queue

The Average Disk Queue section of the OS Detail page succinctly communicates the general overall performance levels of the server's I/O. The table below describes the statistics in this section:

Statistic	Description	Information
Average Read Queue	Average Read Queue Length is the average number of read requests that were queued for a particular disk.	Sustained loads on any of the disk queue lengths could represent problems within your system.
Average Write Queue	Average Write Queue is the average number of write requests that were queued for a particular disk.	Sustained loads on any of the disk queue lengths could represent problems within your system. Contact your Windows System Administrator for assistance in Correction these problems.
Average Queue	Average time in ms. to service I/O requests which for disks includes seek, rotational latency and data transfer time).	Sustained loads on any of the disk queue lengths could represent problems within your system. Contact your Unix System Administrator for assistance in Correction these problems.

NOTE: The statistics available in this section depend on the platform of operating system.

Location	Microsoft SQL Server Performance Analyst statistics > OS > I/O Tab
Correction	Contact your Windows System Administrator for assistance in Correction these problems.

Disk I/O Time

The Disk I/O Time section of the OS Detail page succinctly communicates the general overall performance levels of the disk's I/O. The table below describes the statistics in this section:

NOTE: The statistics available in this section depend on the platform of operating system.

Statistic	Description	Metrics
% Disk Read Time	% Disk Read Time is the percentage of time during the sample interval that the disk is busy servicing read requests.	Use this in conjunction with the % Disk Write Time to identify the disks that may be running at capacity and may require upgrades. This metric may also be used to balance the load between under/over utilized disks.

Statistic	Description	Metrics
% Disk Write Time	% Disk Write Time is the percentage of elapsed time that the selected disk drive was busy servicing write requests.	Use this in conjunction with the % Disk Read Time to identify the disks that may be running at capacity and may require upgrades. This metric may also be used to balance the load between under/over utilized disks.
% Disk Time	% Disk Time is the percentage of time during the sample interval that the disk is busy servicing requests.	Use this statistic to identify the disks that may be running at capacity and may require upgrades. This metric may also be used to balance the load between under/over utilized disks.

Location	Microsoft SQL Server Performance Analyst statistics > OS > I/O Tab

Disk Transfers/Sec

The Disk Transfers/Sec section of the OS Detail page succinctly communicates the general overall performance levels of the disk transfers. The table below describes the statistics in this section

Statistic	Description
Reads/Sec	The the rate of read operations on the disk.
Writes/Sec	The rate of write operations on the disk.
Transfers/Sec	The rate of transfer operations per second.

NOTE: The statistics available in this section depend on the platform of operating system.

Location	Microsoft SQL Server Performance Analyst statistics > OS > I/O Tab

Memory Tab

The Memory tab of the OS Detail page includes the following sections:

Buffer Size	Cache Size	Cache Efficiency
Memory Available	Paging Activity	Page Replacements

The following statistic, available on this pane, duplicates a statistic appearing on the Bottleneck Analysis pane pf the OS Home page:

- Page Faults/Sec

NOTE: The sections available on this tab depend on the platform of operating system.

>

Paging Activity

The Paging Activity section includes the following statistics:

Blocks Input/Sec	Blocks Output/Sec	Paged In
Paged Out	Pages Input/Sec	Pages Output/Sec

NOTE: The statistics available in this section depend on the platform of operating system.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab

Blocks Input/Sec

The Blocks Input/sec statistic is the number of blocks sent to a block device per second.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab
Correction	Although it never hurts to have as much physical memory as your system can handle, here are two things you can check within your system to alleviate the memory bottleneck: *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your swap file. Normally, the bigger the initial size of your swap file, the better, in performance terms.*

Blocks Output/Sec

The Blocks Output/sec statistic is the number of blocks received from a block device per second).

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab
Correction	Although it never hurts to have as much physical memory as your system can handle, here are two things you can check within your system to alleviate the memory bottleneck. *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your swap file. Normally, the bigger the initial size of your swap file, the better, in performance terms*.

Paged In

The number of disk block paged into core from disk which occurs when the server is experiencing a shortage of RAM. Any nonzero value indicates excessive swap activity. In any virtual memory environment, it is not uncommon to see RAM pages moved out to a swap disk. This is a special disk area in UNIX that's reserved for holding memory pages so that the processor is capable of addressing RAM in excess of its full capability. While page-out of operations are a normal part of any server's operations, page-in operations indicate that the real amount of available RAM has been exceeded and that additional RAM pages are required on the server.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab
Correction	High values indicate the possible need to increase the amount of RAM on the system.

Paged Out

The number of disk block paged out of core from disk. In any virtual memory environment, it is not uncommon to see RAM pages moved out to a swap disk. This is a special disk area in UNIX that's reserved for holding memory pages so that the processor is capable of addressing RAM in excess of its full capability. While page-out of operations are a normal part of any server's operations, page-in operations indicate that the real amount of available RAM has been exceeded and that additional RAM pages are required on the server.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab

Pages Input/Sec

The Pages Input/Sec statistic is the number of pages read from disk to resolve hard page faults. Hard page faults occur when a process requires code or data that is not in its working set or elsewhere in physical memory, and must be retrieved from disk.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab
Metrics	This value was designed as a primary indicator of the kinds of faults that cause system-wide delays. It includes pages retrieved to satisfy faults in the file system cache (usually requested by applications) and in non-cached mapped memory files. This counter counts numbers of pages, and can be compared to other counts of pages, such as Memory: Page Faults/sec, without conversion. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
Correction	Although it never hurts to have as much physical memory as your system can handle, there are some things you can check within your system to alleviate the memory bottleneck. *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your page file. Normally, the bigger the initial size of your page file, the better, in performance terms.*

Pages Output/Sec

The Pages Output/Sec statistic is the number of pages written to disk to free up space in physical memory. Pages are written back to disk only if they are changed in physical memory. A high rate of pages output might indicate a memory shortage.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab
Metrics	Windows NT writes more pages back to disk to free up space when low in physical memory. This counter counts numbers of pages, and can be compared to other counts of pages, without conversion. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
Correction	Although it never hurts to have as much physical memory as your system can handle, there are some things you can check within your system to alleviate the memory bottleneck. *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your page file. Normally, the bigger the initial size of your page file, the better, in performance terms.*

Memory Available

The Memory Available section of the OS Detail page succinctly communicates the general overall performance levels of the server's memory. The statistics available in this section depend on the platform of operating system. The table below describes these statistics:

Statistic	Description	Metrics	Correction
Physical Memory	The Physical Memory statistic is the amount of committed virtual memory, in megabytes. Committed memory is where memory space has been reserved in the paging file if it needs to be written to disk.	This counter displays the last observed value only and not an average.	N/A
Physical	The Physical statistic is the total physical memory available.	N/A	N/A
Paged Memory	The Paged Memory statistic is the amount of committed virtual memory, in bytes. Used Memory is the physical memory which has space reserved on the disk paging file(s). There can be one or more paging files on each physical drive.	This counter displays the last observed value only; it is not an average.	N/A
Swap Memory	The Swap Memory statistic is the free, unreserved swap space.	N/A	Increase the size of your swap file or add additional files if you consistently see a shortage of free, unreserved swap space.
Swap	The Swap statistic is the total swap memory being used.	N/A	N/A
Total Physical	The Total Physical statistic is the total physical memory available.	N/A	N/A
Total Swap	The Total Swap statistic is the total swap memory available.	N/A	N/A
Swap-ins	The number of memory pages paged in from swapspace.	N/A	If the system is consistently low on free memory (over a 30 second average), the memory scheduler will start to swap out processes. Increase the size of your swap file or add additional files.
Swap-outs	The number of memory pages paged in from swapspace.	N/A	If the system is consistently low on free memory (over a 30 second average), the memory scheduler will start to swap out processes. Increase the size of your swap file or add additional files.
Buffers	The Buffers statistic represents the amount of buffers used by the kernel in kilobytes.	N/A	N/A

>

Statistic	Description	Metrics	Correction
Page Faults/Sec	The Page Faults/Sec statistic is the overall rate faulted pages are handled by the processor. It is measured in numbers of pages faulted per second. A page fault occurs when a process requires code or data that is not in its working set. This counter includes both hard faults and soft faults.	This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.	If the number of page faults remains consistently high, you can check with your Windows System Administrator for further investigation. Often, large numbers of page faults are not a problem so long as they are soft faults. However, hard faults, that require disk access, can cause delays.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab

Buffer Size

The Buffer Size statistic represents the amount of cache used by the kernel in kilobytes.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab

Cache Size

The Cache Size statistic represents the amount of cache used by the kernel in kilobytes.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab

Page Replacements

The following statistics are available in this section:

Memory Freed (Pages/sec)	Location	

Memory Freed (Pages/sec)

Freed or destroyed (Kb/s).

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab

Clock Algorithm Scans (Pages/sec)

The VMM uses a technique known as the clock algorithm to select pages to be replaced. This technique takes advantage of a referenced bit for each page as an indication of what pages have been recently used (referenced). When the page-stealer routine is called, it cycles through the PFT, examining each page's referenced bit.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab

>

Cache Efficiency

The Cache Efficiency section of the Memory tab succinctly communicates the general overall performance levels of the server's memory. The following statistics are available in this section:

Copy Read Hits %	Data Map Hits %	MDL Read Hits %
Pin Read Hits %		

Copy Read Hits %

The Copy Read Hits % statistic is the percentage of cache copy read requests that hit the cache and does not require a disk read to provide access to the page in the cache.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab
Metrics	When the page is pinned in the memory, the page's physical address in the file system cache will not be altered. A copy read is a file read operation where a page in the cache is copied to the application's buffer. Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.
Correction	Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate

Data Map Hits %

The Data Map Hits % statistic is the percentage of data maps in the file system cache that could be resolved without having to retrieve a page from the disk.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab
Metrics	Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.
Correction	Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate.

MDL Read Hits %

The MDL Read Hits % statistic is the percentage of Memory Descriptor List Read requests to the file system cache that hit the cache and does not require disk access to provide memory access to the pages in the cache.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab
Metrics	Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.
Correction	Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate.

Pin Read Hits %

The Pin Read Hits % statistic is the percentage of pin read requests that hit the file system cache and does not require a disk read in order to provide access to the page in the file system cache.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab
Metrics	Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.
Correction	Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate.

Space Tab

The Space tab of the OS Detail page includes the following sections:

Disk Space Free	Disk Space Detail	Device Summary
Device Detail		

Disk Space Free

The Disk Space Free metric displays the amount of free disk space in megabytes.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab

Disk Space Detail

The Disk Space Detail section of the Space tab succinctly communicates the general overall performance levels of the server's disks and space allotment. The table below describes the statistics in this section:

Statistic	Description
Partition	The drive letter of the disk.
Local Filesystem	The name of the file system.
Type	The type of file system.
Total Space	Total size of the disk/device's capacity expressed in MBs.
Used Space	Amount of MBs currently allocated on the particular disk/device.
Free Space	Amount of MBs currently unallocated and free on the particular disk/device.
Capacity	The percentage of space used on the device.
Mounted On	The mount point of the device.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab

Device Summary

The Device Summary section of the Space tab displays a bar chart for all devices. The Device Summary section succinctly communicates the general overall performance levels of the server's disks and space allotment. The table below describes the statistics in this section:

Statistic	Description
Used	Amount of MBs currently allocated on the particular disk/device.
Free	Amount of MBs currently unallocated and free on the particular disk/device.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab

Device Detail

The Device Detail section of the Space tab succinctly communicates the general overall performance levels of the server's disks and space allotment. The statistics available in this section depend on the platform of operating system. The table below describes the statistics in this section:

Statistic	Description
Partition	The drive letter of the disk.
Local Filesystem	The name of the file system.
Type	The type of file system.
Total Space	Total size of the disk/device's capacity expressed in MBs.
Used Space	Amount of MBs currently allocated on the particular disk/device.
Free Space	Amount of MBs currently unallocated and free on the particular disk/device.
Capacity	The percentage of space used on the device.
Mounted On	The mount point of the device.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Memory Tab

Network Tab

The Network tab of the OS Detail page succinctly communicates the general overall performance levels of the server's networking. The Network tab of the OS Detail page includes the following sections:

Network Details	Packet Collisions	Packet Discards
LPacket Errors	Transmission Rate	Transmission Rate (Bytes)
Transmission Queue Length		

NOTE: The sections available on this tab depend on the platform of operating system.

>

Network Details

The statistics available in this section depend on the platform of operating system. The table below describes the information available in this section:

Column	Description
Network Interface	The name of network interface.
INET Address/Address	The IP address assigned to the network interface.
Pkts Sent/Sec	The number of packets sent per second.
Pkts Received/Sec	The number of packets received per second.
Sent (KB/Sec)	The number of bytes sent per second.
Received (KB/Sec)	The number of bytes received per second.
Out Pkts Discarded	The number of outbound packets discarded.
In Pkts Discarded	The number of inbound packets discarded.
Out Pkt Errors	The number of outbound packet errors.
In Pkt Errors	The number of inbound packet errors.
Queue Length	The queue length.
Collisions	The number of collisions.
Packets Discarded	The number of packets discarded.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Network Tab

Transmission Rate

The Transmission Rate section of the Network tab succinctly communicates the packet transmission rate. The following statistics are available in this section:

- Location
- Packets Sent/Sec

Packets Sent/Sec

The Packets Sent/Sec statistic is the number of packets sent over each network adapter per second.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Network Tab

Packets Received/Sec

The Packets Received/Sec statistic is the number of packets received over each network adapter per second.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Network Tab

>

Transmission Rate (Bytes)

The Transmission Rate (Bytes) section of the Network tab succinctly communicates the packet transmission rate. The following statistics are available in this section:

Received (KB/Sec)	Sent (KB/Sec)	

Sent (KB/Sec)

The Sent (KB/Sec) statistic is the rate at which bytes are sent over each network adapter including framing characters.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Network Tab

Received (KB/Sec)

The Received (KB/Sec) statistic is the rate at which bytes are received over each network adapter, including framing characters.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Network Tab

Transmission Queue Length

The number of pending outgoing packets on either the software transmit queue or the hardware transmit queue

Location	Microsoft SQL Server Performance Analyst statistics > OS > Network Tab

Packet Collisions

A collision occurs when two devices on the network sense that the network is idle and end up trying to send data at the same time. Some collisions are normal on a repeated network, but excess collisions can cause serious performance problems.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Network Tab

Correction

Contact your network administrator if there is high rate of persistent Packet Collisions.

Packet Discards

The Packet Discards section of the Network tab succinctly communicates the packet discard rate. The following statistics are available in this section:

Inbound Packets Discarded	Outbound Packets Discarded	Packet Discards

Outbound Packets Discarded

The Outbound Packets Discarded statistic is the number of outbound packets that were discarded even though no errors had been detected to prevent from being transmitted.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Network Tab

Correction	A possible reason for discarding such a packet could be to free up buffer space.

Inbound Packets Discarded

Received Packets Discarded is the number of received packets that were discarded even though no errors had been detected to prevent from being transmitted.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Network Tab
Correction	A possible reason for discarding such a packet could be to free up buffer space.

Packet Discards

The Packet Discards statistic represents the number of network packets discarded.

LPacket Errors

ocation	Microsoft SQL Server Performance Analyst statistics > OS > Network Tab
Correction	Contact your network administrator if there is high rate of persistent packet discards.

The Packet Errors section of the Network tab succinctly communicates the packet error rate. The following statistics are available in this section:

Inbound Packet Errors	Outbound Packet Errors	

Outbound Packet Errors

The Outbound Packet Errors statistic is the outbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Network Tab

Inbound Packet Errors

The Packets Received Errors statistic is the received packets that contained errors preventing them from being deliverable to a higher-layer protocol.

Location	Microsoft SQL Server Performance Analyst statistics > OS > Network Tab

Top SQL

One or two bad queries can cause a lot of trouble for the remaining sessions in a database, and so can one or two bad queries. It is important to find them before they get into a production system, but sometimes a few sneak through.

By applying custom filters and performance-related thresholds, the Top SQL view locates inefficient SQL. By applying filters to certain I/O and statistical counters, you hope to isolate queries that far exceed their nearest competitors in the same area (like disk reads). When you find them, you should reduce the number of sorts a query performs. Or, for a query that returns only a few records, you should try to minimize the number of rows a query processes.

>

Top SQL displays requested SQL for SQL Server datasources. The following tabs are available on the Top SQL page:

Summary Tab	Criteria Tab	SQL Details Tab

Summary Tab

The Summary tab of the Top SQL page includes the following sections:

Top SQL Statements - SQL Server	Top SQL Details	

Top SQL Statements - SQL Server

A lot of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured database down in terms of end user response times.

Before you can identify problem SQL in your database, you have to ask the question of what 'bad SQL' is. What criteria do you use when you begin the hunt for problem SQL in your critical systems? Understand that even the seasoned experts disagree on what constitutes efficient and inefficient SQL; so there is no way to sufficiently answer this question to every Oracle professional's satisfaction. The table below describes some general criteria you can use when evaluating the output from various database monitors or personal diagnostic scripts:

Criteria	Description
Overall Response (Elapsed) Time	This is how much time the query took to parse, execute, and fetch the data needed to satisfy the query. It should not include the network time needed to make the round trip from the requesting client workstation to the database server.
CPU Time	This is how much CPU time the query took to parse, execute, and fetch the data needed to satisfy the query.
Physical I/O	Often used as the major statistic in terms of identifying good vs. bad SQL, this is a measure of how many disk reads the query caused to satisfy the user's request. While you certainly want to control disk I/O where possible, it is important that you not focus solely on physical I/O as the single benchmark of inefficient SQL. Make no mistake, disk access is slower than memory access and also consumes processing time making the physical to logical transition, but you need to look at the entire I/O picture of a SQL statement, which includes looking at a statements' logical I/O as well.
Logical I/O	This is a measure of how many memory reads the query took to satisfy the user's request. The goal of tuning I/O for a query should be to examine both logical and physical I/O, and use appropriate mechanisms to keep both to a minimum.
Repetition	This is a measure of how often the query has been executed. A problem in this area is not as easy to spot as the others unless you know your application well. A query that takes a fraction of a second to execute may still be causing a headache on your system if it is executed erroneously over and over.

There are other criteria that you can examine like sort activity or access plan statistics (that show items like Cartesian joins and the like), but more often than not, these measures are reflected in the criteria listed above.

Fortunately, Microsoft records some of these statistics and Performance Analyst computes the rest, which makes tracking the SQL that has been submitted against a SQL Server much easier.

Location	Microsoft SQL Server Performance Analyst statistics > Top SQL > Summary Tab

>

Metrics	When you begin to look for inefficient SQL in a database, there are two primary questions you need to answer: *What has been the worst SQL that's historically been run in my database?* *What is the worst SQL that's running right now in my database?* When Correction a slow system, you should be on the lookout for any query that shows an execution count that is significantly larger than any other query on the system. It may be that the query is in an inefficient Transact SQL loop, or other problematic programming construct. Only by bringing the query to the attention of the application developers will you know if the query is being mishandled from a programming standpoint.

Top SQL Details

A lot of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured database down in terms of end user response times.

Before you can identify problem SQL in your database, you have to ask the question of what 'bad SQL' is. What criteria do you use when you begin the hunt for problem SQL in your critical systems? Understand that even the seasoned experts disagree on what constitutes efficient and inefficient SQL; so there is no way to sufficiently answer this question to every SQL Server professional's satisfaction. The Top SQL Details Section displays the top SQL statements across various performance categories. The table below describes the information available in this section:

Information	Description
Top For	Indicates the performance category for the SQL statement.
Usage	This displays the percent usage for this SQL statement in this performance category.
SQL text	A snippet of the actual SQL statement or code object (procedure, etc.)
Total Reads	The total physical reads caused by the statement since monitoring began.
Total Writes	The total physical writes caused by the statement since monitoring began.
Executions	The total number of times the statement or code object has been executed since monitoring began.
Avg CPU Time	The average CPU time used by the statement each time it is executed.
Avg Elapsed Time	The average elapsed response time experienced each time the statement is executed.

Location	Microsoft SQL Server Performance Analyst statistics > Top SQL > Summary Tab
Metrics	When you begin to look for inefficient SQL in a database, there are two primary questions you need to answer: *What has been the worst SQL that's historically been run in my database?* *What is the worst SQL that's running right now in my database?* When Correction a slow system, you should be on the lookout for any query that shows an execution count that is significantly larger than any other query on the system. It may be that the query is in an inefficient Transact SQL loop, or other problematic programming construct. Only by bringing the query to the attention of the application developers will you know if the query is being mishandled from a programming standpoint.

Criteria Tab

The Top SQL view displays all SQL running or executed in the current datasource. You can use the lists and boxes to reorder the display according to your specific needs. The Criteria tab of the Top SQL page includes the following sections that you can mix and match to give you the information you want to find.

- Demographic Criteria
- Statistical Criteria

>

- Sort Criteria

SQL Details Tab

The SQL Details tab of the Top SQL page includes the following sections:

Top SQL Results	Location	

Top SQL Results

Overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements or stored procedures can drag an otherwise well-configured database down in terms of end user response times.

Before you can identify problem SQL in your database, you have to ask the question of what 'bad SQL' is. What criteria do you use when you begin the hunt for problem SQL in your critical systems? Understand that even the seasoned experts disagree on what constitutes efficient and inefficient SQL; so there is no way to sufficiently answer this question to every SQL Server professional's satisfaction. The Top SQL results grid displays the top SQL statements across various performance categories. The table below describes the information available in this section:

Information	Description
SQL Text	A snippet of the full SQL statement. Clicking the statement will cause the full statement to be presented in the Selected SQL Text grid.
Database	The database where the statement or code object's execution originated.
Avg Reads	The average physical reads caused by the statement each time it is executed.
Total Reads	The total physical reads caused by the statement since monitoring began.
Avg Writes	The average physical writes caused by the statement each time it is executed.
Total Writes	The total physical writes caused by the statement since monitoring began.
Executions	The total number of times the statement or code object has been executed since monitoring began.
Avg Elapsed Time	The average elapsed response time experienced each time the statement is executed.
Elapsed Time	The total elapsed response time over all executions.
Avg CPU Time	The average CPU time used by the statement each time it is executed.
Total CPU Time	The total CPU time over all executions.
First Start Time	The date/time when the statement was first executed.

Location	Microsoft SQL Server Performance Analyst statistics > Top SQL > SQL Details Tab

Selected SQLText

The Selected SQL Text window allows you to select any of the Top SQL result rows and get the complete SQL text. From here you can open an Explain Plan or Tune the SQL.

Location	Microsoft SQL Server Performance Analyst statistics > Top SQL > Summary Tab

>

Performance Analyst for Sybase Tutorial

Embarcadero Performance Analyst provides key insight into the complex world of database performance. Running inside the Embarcadero DBArtisan console, or standalone in a Web browser, Performance Analyst helps you pinpoint performance problems that threaten a database's availability or response times. The powerful analysis facilities will satisfy advanced experts, yet novices will find Performance Analyst's interface easy to use and the expert help a welcome aid to troubleshooting difficult performance issues.

About This Performance Analyst for Sybase Tutorial

This tutorial is intended to help you get started with Embarcadero Performance Analyst.

After completion of this tutorial, you will have the foundation you need to explore the many features and benefits of Performance Analyst. You will have learned how to accurately monitor your database using the most popular performance methodologies in use today. You will know how to drill down into the many layers of database performance and also be able to locate performance issues in your database in a fast and effective manner. Finally, you'll understand how to fix performance problems using the integration that exists between Performance Analyst and Embarcadero DBArtisan.

This tutorial is divided into 7 sections.

You can use this basic tutorial as a roadmap of product highlights.

After starting this tutorial, you can select **Help** from the menu to find many additional resources that complement and build on many of the activities shown in this brief guide.

For more information, see Additional Evaluation Resources.

Session 1: Getting Started with Performance Analyst for Sybase

Download and Install

You can obtain the latest version of Performance Analyst by downloading the DBArtisan software from the Embarcadero Web site at http://www.embarcadero.com/downloads/downloaddbartisan.asp. Performance Analyst runs fully contained within DBArtisan so you will need the DBArtisan product to use Performance Analyst (although Performance Analyst can also run in standalone fashion within a Microsoft Internet Explorer Web browser).

Click **Download**, and follow the steps indicated. Save the file on your computer and then double-click to launch the self-extracting file that will guide you through the installation process.

When you first install an evaluation copy of DBArtisan, you can use the tool for 14 days. After that time, a permanent license is needed.

Proceed to Overview of Performance Analyst for Sybase

Overview of Performance Analyst for Sybase

When DBArtisan is connected to a Sybase database, you can open Performance Analyst two ways:

1 On the **Analyst** toolbar click the **Performance Analyst** icon.

OR

2 On **Utilities** menu, select **Performance Analyst**.

For Performance Analyst to work, several procedures need to be installed by the SA account. When Performance Analyst is first invoked, it will check for the presence of this package and if it has not been installed, the product will prompt you to install it. Once created, any account with access to these procedures can monitor databases using Performance Analyst. No other objects need to be installed on a monitored database.

Proceed to Performance Analyst for Sybase Product Navigation

Performance Analyst for Sybase Product Navigation

Navigating through Performance Analyst is easily accomplished through two main mechanisms:

1 The main Performance Analyst toolbar contains buttons that invoke the major performance analysis views within the product as well as frequently used actions like Refresh and Performance Threshold Customizations.

2 Positioning the mouse pointer on any area within Performance Analyst and clicking the right mouse button will bring up a navigation menu that allows you to easily move between any of the main Performance Analyst home pages and drill down views. The right mouse menu also contains context sensitive options like drilling down further into a particular statistic, copying the display to a text file, and more.

Proceed to Session 2: Using Ratio-Based Performance Analysis in Performance Analyst for Sybase

Session 2: Using Ratio-Based Performance Analysis in Performance Analyst for Sybase

Ratio-based analysis has been around for many years, and used to be the only technique database administrators used when they were called on to diagnose the cause of a database slowdown. Performance ratios are very good roll-up mechanisms for busy DBAs, making possible the analysis-at-a-glance approach, however ratios must be computed correctly for them to be meaningful. Key ratios in Performance Analyst are calculated using *delta* statistics, which is superior to computing ratios with cumulative values. Delta statistics, which are values derived by using the before and after measures in each statistical refresh done in Performance Analyst, lets you see changes in performance metrics that would otherwise be invisible if cumulative statistics were used. Performance Analyst makes it easy to see the ratios that are computed with delta statistics vs. cumulative. A triangle symbol denotes a delta-based statistic, while square symbols represent a cumulative-based measure.

Ratio-based analysis is available for all performance categories (memory, etc.) where ratios can assist in determining the performance level of a database. For example, if you would like to see memory-specific ratios, on the Performance Analyst toolbar click **Memory**. Performance Analyst displays the Memory home page that showcases all important memory-related ratios.

Proceed to Obtaining Detail Information for a Ratio Using Performance Analyst for Sybase

Obtaining Detail Information for a Ratio Using Performance Analyst for Sybase

A ratio gives you a global picture for a particular performance measure. However, you often need more information regarding a statistical measure so you can make an informed decision. Performance Analyst has easy-to-use drill down features that make obtaining detailed information for various ratios a breeze.

For example, to view more information for the Data Cache Hit Ratio statistic, do the following:

1 On either the **Main** Home page or the **Memory** Home page, locate the **Data Cache Hit Ratio** statistic.

2 Position the pointer over the statistic.

The pointer changes to a hand symbol.

3 Double-click the left mouse button.

Performance Analyst opens the Data Cache Tab of the memory drill down view.

OR

4 Right-click the statistic, select **Details**.

OR

5 Select **Navigate to**, **Memory**, **Data Cache**.

Performance Analyst opens the Data Cache Tab of the memory drill-down view.

Proceed to Session 3: Using Bottleneck/Wait-Based Performance Analysis in Performance Analyst for Sybase

Session 3: Using Bottleneck/Wait-Based Performance Analysis in Performance Analyst for Sybase

When a Sybase database is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting may mean nothing in the overall scheme of things, or it can be an indicator that a database bottleneck exists. Bottleneck analysis is a valid method of measuring performance because it helps you track where a database and user sessions have been spending their time. If latch contention or heavy table-scan activity has been dragging a database's performance down, you can use bottleneck analysis to confirm the actual root cause.

Performance Analyst uses three broad areas of metrics to form its bottleneck analysis. The first are contentions and conflicts as reported in Sybase's monitoring *sp_sysmon* procedure, as well as overall lock statistics. The second is the Sybase wait event interface (available only in Sybase 12.5.03 and higher), which displays where the database and user sessions have been bottlenecked. A third area of bottleneck information has to do with space related bottlenecks like databases and transaction logs that are running low on available free space.

Performance Analyst prominently displays bottleneck analysis on every performance home page. The Main home page, displays a summary of all key bottlenecks threatening the database. You can also see specific bottlenecks for memory, I/O, space, objects, and users by navigating to each home page and viewing the Bottleneck Analysis section for that home page. For example, to see bottlenecks that specifically relate to I/O, do the following

1 On the toolbar, click the **I/O** button.

2 Locate the Bottleneck Analysis section at the lower-left portion of the screen.

The I/O bottleneck analysis section shows the I/O related headaches, along with the most active devices.

Proceed to Obtaining Detail on System Bottlenecks Using Performance Analyst for Sybase

Obtaining Detail on System Bottlenecks Using Performance Analyst for Sybase

You can easily obtain more information regarding global bottlenecks that are plaguing a database by drilling down into a particular bottleneck metric. For example, if you would like to see more data regarding the current top system bottlenecks, do the following:

1 Right-click anywhere in Performance Analyst, point to **Users**, and then click **System Waits**.

Performance Analyst opens the System Waits view.

The System Waits view (available in Sybase 12.5.03 and higher with the MDA monitoring views installed) displays critical details regarding all the wait events that Sybase has recorded. You can easily see what events that occurred the most often as well as the events that have accumulated the most wait time.

TIP: If you would like to see an expanded view of the System waits grid shown on the top of the System Waits view, click the Expand button located at the top right corner of the section. All section views in Performance Analyst can be maximized/minimized using this control.

Proceed to Viewing Detail for Session Bottlenecks Using Performance Analyst for Sybase

Viewing Detail for Session Bottlenecks Using Performance Analyst for Sybase

Not only can you obtain detail information for system bottlenecks, but you can also use Performance Analyst to get drill down information regarding session-level bottlenecks. There are several layers to session-level details you can view.

To see more information regarding session-level waits, do the following:

1 Right-click anywhere in Performance Analyst, point to **Users**, and then click **Session Waits**.

Performance Analyst opens the Session Waits drill-down view.

When you are in the Session Waits drill down view, to drill further into a session's details, do the following:

1 In the **Historical Waits** or **Current Waits** grid, select a session that appears to be experiencing critical waits.

2 Double-click the session or right-click, and then select **Details**.

Performance Analyst opens a drill-down view that shows key statistics for the selected session, including its current work activities, wait events, and SQL statements.

Proceed to Session 4: Using Workload Performance Analysis in Performance Analyst for Sybase

Session 4: Using Workload Performance Analysis in Performance Analyst for Sybase

Key ratios give you a global perspective on database activity. Bottleneck analysis gives you insight into things that are holding up user activity and throughput. But another technique is necessary if a database professional is to really get a handle on what's occurring inside a badly performing database.

Workload analysis involves the investigation of two critical areas of a database's performance:

- Session resource consumption and activity
- SQL execution (if available)

Without looking at these two key performance categories, you can miss a lot of what could be responsible for perceived performance problems. When performance on a database takes a sudden nosedive, it is not uncommon to find one or two sessions that are causing the bulk of the workload. Finding these processes can be easily accomplished by viewing session metadata coupled with resource consumption and statistical execution statistics.

>

Proceed to Pinpointing Top Resource Consumers

Pinpointing Top Resource Consumers

Performance Analyst identifies top resource consumers in a couple of different ways. On the main home page, Performance Analyst highlights the top resource sessions across I/O, memory usage, and CPU consumption. The percentage used across all statistical categories are displayed so a DBA can immediately pinpoint a session that is using all or most of a particular resource. The top resource sessions are also displayed on the memory, I/O, and other home pages as well with a listing for the top sessions for that category (for example, the top memory users appear on the memory home page, etc.)

Proceed to Getting Workload Details

Getting Workload Details

Performance Analyst offers several details on top resource sessions in addition to the summaries provided on each home page. A detailed and flexible top sessions view lets you view a summary of all the major resource sessions on the database as well as every critical detail regarding each session connected to the database.

Further drill downs are available that highlight specific details regarding each session's activities and submitted SQL. To see detailed information regarding a session identified in one of Performance Analyst's Workload Analysis sections, do the following:

1. In a **Workload Analysis** display, position the pointer over one of the processes.
2. Double-click the process or right-click, and then select **Details**.

If you want like to see all key workload analysis metadata for sessions, do the following:

1. On the toolbar, click the **Top Sessions** button.

   OR

2. Right-click anywhere in Performance Analyst, point to **Users**, and then click **Session Waits**.

Proceed to Session 5: Using Operating System Performance Analysis in Performance Analyst for Sybase

Session 5: Using Operating System Performance Analysis in Performance Analyst for Sybase

In many scenarios, an optimally tuned database may not perform well because there are constraints imposed by the system where the database is running. These constraints may include processes competing with the database sever for resources (CPU, I/O or Memory), a slow CPU, insufficient I/O devices, and insufficient memory. Performance Analyst allows you to examine the relevant operating system metrics for Windows Platforms (NT /2000/XP/2003) as they relate to Sybase.

To view the operating system metrics of the machine hosting the Sybase server you are monitoring, do the following:

1. Click the **OS** Performance Analyst toolbar button.

   OR

2. Navigate to the OS Home Page by going to **OS**, and then **Home**.

   NOTE: The first time you invoke the OS feature you will be prompted for an operating system login. Please specify the correct operating system, host name, username and password, and protocol.

The Operating System home page has been designed in a similar fashion as the datasource home page. On this page you can see the pertinent operating system metrics displayed in five quadrants broken down into the following categories:

- Resource Usage
- Memory Analysis
- Disk Analysis
- Workload Analysis
- Bottleneck Analysis

Proceed to Obtaining Details on the Operating System Using Performance Analyst for Sybase

Obtaining Details on the Operating System Using Performance Analyst for Sybase

You can easily obtain more information regarding operating system problems that are degrading database performance by drilling down into a particular analysis quadrant. For example if you'd like to see more information regarding the processes running on the host machine, do the following:

1. Navigate to the Performance Analyst **Operating System** home page.
2. Double-click the **Top Memory Process** in the **Workload Analysis** quadrant.

   Performance Analyst opens the Processes detail view.

   OR
3. Right-click anywhere in Performance Analyst and select the **Navigate to**, **OS**, and then **Processes**.

The Processes detail tab shows all the processes on the host machine. From this view you can easily determine which processes are showing high marks in I/O or CPU, for example, by simply double clicking on the column headers and sorting the processes in ascending or descending order.

Proceed to Session 6: Identifying and Getting Help with Performance Problems Using Performance Analyst for Sybase

Session 6: Identifying and Getting Help with Performance Problems Using Performance Analyst for Sybase

You do not have to be an advanced expert in performance tuning to spot trouble in a database if you use Performance Analyst. There are a number of mechanisms in the tool that help you quickly locate areas of concern within a database.

Proceed to Using the Alarm Log in Performance Analyst for Sybase

Using the Alarm Log in Performance Analyst for Sybase

Performance Analyst has a central and flexible alarm log that lists every performance violation found in all monitored databases. Unlike some products that only show you alarms for a particular database, Performance Analyst lists all alarms for each database you are monitoring in one central log so you can see every problem in one place. To view the Performance Analyst Alarm log, do the following:

1. On the toolbar, click the **Alarms** button.

Once in the Alarm log, you can perform a number of actions including:

- Selecting an alarm and viewing the application page that shows details regarding the metric in violation.
- Filtering the types of alarms shown through the Options dialog.
- Clearing all current alarms.
- Copying or saving the Alarm log into a different format.

To select any of these options, do the following:

1 Position the pointer over any area of the **Alarm Log**.

2 Right-click and in the shortcut menu, select an option.

Proceed to Recognizing Performance Violations

Recognizing Performance Violations

Performance Analyst makes it easy to visually pinpoint performance violations in a database. Any statistic that does not conform to your predefined performance standards is colorized in any view that it appears so you can easily see potential troublespots.

The colors used to highlight a particular performance violation depend on the severity of the problem. Performance Analyst uses several layers of performance standards/threshold violations with each layer being represented by a different color. You can control everything regarding threshold violations, including how thresholds are determined, their color representation in Performance Analyst's display, and more by using the Threshold Template feature of Performance Analyst.

How can Performance Analyst inform you of performance violations in your database if you minimize the product to your desktop and are working on other things? Even if Performance Analyst is not right in front of you, it can still let you know that a problem has been found in your database through the systray alarm icon. When Performance Analyst detects a problem in one of your databases, one of the things the product does is enable an alarm icon in your workstation's systray area that communicates the presence of alarms in your database. So even if the Performance Analyst interface is not in sight, the product can let you know that something in your database requires attention.

Proceed to Creating Threshold Templates

Creating Threshold Templates

How does Performance Analyst know what to signal as a problem in a database? It uses predefined performance standards that you define and enforce. All of this is accomplished by using Performance Analyst's Threshold templates. A template is a set of performance metrics and their thresholds that can be reused to enforce performance standards across any number of databases. Embarcadero provides a default performance template that contains nearly all of the most-watched performance statistics so you can begin monitoring with confidence right out of the box. However, you can create your own Threshold templates so you only are alerted to performance violations that are important to you and your particular environment.

To view or create Threshold templates, do the following:

1 On the toolbar, click the **Thresholds** button.

Performance Analyst opens the Threshold Template Properties dialog box.

The Threshold Template Properties dialog box lets you select the options to add or edit templates, as well as other maintenance options. For example, you can copy the original Embarcadero Threshold template and modify it to fit your own needs. You can control what statistics are sampled for performance violations as well as the levels of severity you would like reported. Once you are satisfied with your template, you can save it and then apply it to one or more datasources in one action to immediately begin enforcing your performance standards.

Proceed to Getting Help with Performance Problems

Getting Help with Performance Problems

As databases become more complex, so do the problems that plague them. Knowing what to do for a particular performance problem can be hard in some cases, but Performance Analyst helps by providing a complete set of expert help that you can use as a guide when performance problems occur. To obtain expert help for any statistic shown in Performance Analyst, do the following:

1. Position the pointer over a statistic that you need help.
2. Right-click and select **Help**.

The Expert Help gives you a general description of the statistic, the range of values that are considered optimal, and some suggestions on how to fix a problem situation. Expert help can also be invoked for any of the Ratio, Bottleneck, SQL, or Workload Analysis sections shown on every home page by clicking on the help icon for that section, which is located at the top right corner of the section.

Proceed to Session 7: Fixing Identified Problems

Session 7: Fixing Identified Problems

Finding performance problems is one thing, but fixing them is another. Many performance monitors point out a particular problem in a database, but then offer no way to actually correct the situation. Such is not the case with Performance Analyst. Because of its tight integration with DBArtisan, certain problems found in Performance Analyst can be visually corrected by DBArtisan.

A Space Problem Correction Example

If a Sybase Server tablespace is about to run out of free space, Performance Analyst can issue an alarm that a space shortage exists. Once the space problem has been identified, do the following:

1. On the **Space** home page toolbar, click the **Space** button.
2. In the **Space Analyst** grid, right-click a tablespace, and then select **Manage with DBArtisan**.

   DBArtisan opens an editor where you can perform corrective actions, such as enabling Autoextend for the tablespace, resizing the tablespace's existing datafile(s), or adding a new datafile for the tablespace.

>

Sybase ASE Performance Analyst statistics

This section includes expert help for all Sybase ASE categories and statistics in the Embarcadero Performance Analyst views and pages. This help is divided into the following sections:

Home	Memory	I/O
Space	Databases	Users
OS	Top SQL Statistics	

NOTE: **Sybase 12.5.0.3 users** should make sure their MDA tables are properly configured so they can see all available statistics. Please refer to the instructions for Enabling MDA Tables.

Enabling MDA Tables

In ASE 12.5.0.3, a new feature called "MDA tables" was introduced. These proxy tables are mapped to RPCs which provide access to low-level monitoring information about the server.

For more information see: http://www.sybase.com/detail?id=1010039 “The MDA tables – New Monitoring Capabilities in ASE 12.5.0.3”

DBArtisan’s Performance Analyst for Sybase requires these tables be installed to enable several statistics, for example the Top SQL view.

Please follow the instructions below to install and configure these tables.

1 Ensure the ‘enable cis” configuration parameter is set to 1.

```
sp_configure 'enable cis'
go
```

2 Add the ‘loopback’ server name alias.

```
if not exists (select 1 from master..sysservers where srvname = "loopback")
 BEGIN
 exec sp_addserver "loopback", null, <enter your server name>
END
```

3 **Install the MDA tables.**

The script can be found in the following location:

$SYBASE/$SYBASE_ASE/scripts/installmontables.

4 **Assign the ‘mon_role’ to logins allowed MDA access.**

NOTE: The ‘sa’ login requires this grant as well.

```
use master
go
grant role mon_role to sa
go
```

>

5 **Enable the following configuration parameters**.

exec sp_configure "enable monitoring", 1

exec sp_configure "sql text pipe active", 1

exec sp_configure "sql text pipe max messages", 2000

exec sp_configure "plan text pipe active", 1

exec sp_configure "plan text pipe max messages", 2000

exec sp_configure "statement pipe active", 1

exec sp_configure "statement pipe max messages", 2000

exec sp_configure "errorlog pipe active", 1

exec sp_configure "errorlog pipe max messages", 2000

exec sp_configure "deadlock pipe active", 1

exec sp_configure "deadlock pipe max messages", 2000

exec sp_configure "wait event timing", 1

exec sp_configure "process wait events", 1

exec sp_configure "object lockwait timing", 1

exec sp_configure "SQL batch capture", 1

exec sp_configure "statement statistics active", 1

exec sp_configure "per object statistics active", 1

exec sp_configure "max SQL text monitored", 2000

Home

The Embarcadero Performance Analyst Home view lets you review availability and overall performance of all monitored databases from a single window. Statistics on the Home view are organized into the following categories:

Bottleneck Analysis pane	Key Ratio Analysis pane	I/O Analysis pane
Storage Analysis pane	Workload Analysis pane	

Key Ratio Analysis pane

Database performance analysts typically use one of two methods for examining the performance levels of a database - ratio-based or wait/bottleneck-based. Ratio-based analysis involves examining a number of key database ratios that can be used to indicate how well a database is running. Performance ratios serve as very good roll-up mechanisms for busy DBAs to use for at-a-glance performance analysis. Many DBAs have large database farms to contend with and cannot spend time checking detailed wait-based analysis outputs for each and every database they oversee. Succinctly presented performance ratios can assist in such situations by giving DBAs a few solid indicators that can be quickly scanned to see if any database needs immediate attention.

While there are certainly many opinions as to what rules to follow, there are some standards that should always be adhered to. To start with, many of the formulas that make up ratio-based analysis must be derived from delta measurements instead of cumulative statistics. Many of the global ratios that a DBA will examine come from the master.sysperfinfo performance table. This table maintains a count of various statistics for the server and particular databases since the server was brought up. For servers that are kept up for long periods of time, these values can grow quite large and will impact how a particular ratio that a DBA might be looking at is interpreted. However, if delta statistics are used (taking, for a specified sampling period, the before and after counts of each statistic that make up a ratio), then an accurate and current portrayal of the various ratios can be had.

A final thing to remember about using ratio-based analysis is that, while there are a number of rules of thumb that can be used as starting points to begin the evaluation of database performance, DBAs must determine each SQL Server's individual 'personality' with respect to the key performance ratios. Some hard and fast rules simply do not apply to every database. The danger in using blanket ratio standards is that they can lead the DBA to haphazardly take action, which can at times contribute nothing to the situation, and sometimes even degrade performance.

The following ratios are used on the Performance Analyst Home page to succinctly communicate the general overall performance levels of the monitored database:

Data Cache Hit Ratio	Procedure Plan Hit Ratio	Large I/O Hit Rate
Clean Buffer Grab Rate	CPU Busy	I/O Busy
Server Idle		

Data Cache Hit Ratio

Data read from memory produces user response times many times faster than when that same data is read from disk. The Sybase data cache assists with keeping physical I/Os to an absolute minimum.

The data cache hit ratio is an excellent indicator of how often user requests for data are satisfied through memory as opposed to being physically read from disk. The table below describes the three key counters in Sybase used to arrive at this statistic:

Counter	Description
LOGICAL READS	Data read from memory for user requests.
PAGES PER I/O	The number of pages retrieved in a single I/O operation.
PHYSICAL READS	Data read physically from disk.

Location	Sybase ASE Performance Analyst statistics > Home > Key Ratio Analysis pane Sybase ASE Performance Analyst statistics > Memory > Home> Key Ratio Analysis pane Sybase ASE Performance Analyst statistics > Memory > Data Cache Tab
Metrics	To help ensure excellent performance, keep your cache hit rate in the neighborhood of 90% or greater. Lower amounts can be okay for user ad hoc databases where sporadic, large table scan operations occur. However, anything below this general threshold for normal databases can require tuning attention, and the adjustment of the Sybase memory tuning parameters. If you are using named caches, you can drill down into the cache hit rates for each named cache. This helps you understand which objects/operations are depressing the overall cache hit rate for the server.

Correction	If a problem is found in Sybase servers, versions 11-12, you can increase the amount of the total memory configuration parameter or reduce the percentage of memory allocated to the procedure cache (by default, the data cache assumes any free memory left over after Sybase has met its kernel and procedure cache needs). Take care when reducing the procedure cache, as this could reduce performance in the server as it relates to reading procedures in from disk. For Sybase 12.5, the total memory configuration parameter can again be increased to provide more memory for the data cache (and any named caches), but in 12.5, if you want to reduce the size of the procedure cache, note that it is now configured in terms of literal size instead of a percentage of the overall configured memory. Once the data cache has been adjusted, monitor Sybase to see if the cache hit rate improves. If it does not, another increase may be necessary and examination of unnecessary large table scan operations. Also, keep a careful eye on the actual machine's memory limits and swap activity. Increased swap activity can be indicative of too little memory left for the server machine.

Procedure Plan Hit Ratio

The Sybase procedure cache is used to hold the definitions and query plans of stored procedures and triggers. It is used for short-term memory needs like statistics and query plans for parallel queries. When a user executes a stored procedure, Sybase looks in the procedure cache for a query plan to use. If a query plan is available, Sybase places it on the most recently used (MRU) end of the memory chain and the procedure begins to execute. If no execution plan is in memory, or if all copies of the plan are currently being used, the query tree for the procedure is read in again from the data dictionary, optimized, put on the MRU end of the chain, and executed. Note that other operations, like CREATE INDEX, can also use the procedure cache even when no procedure is referenced.

The more often that a procedure's plan and definition can be referenced in memory, the better the procedure execution time.

Location	Sybase ASE Performance Analyst statistics > Home > Key Ratio Analysis pane Sybase ASE Performance Analyst statistics > Memory > Home> Key Ratio Analysis pane Sybase ASE Performance Analyst statistics > Memory > Procedure Cache Tab
Metrics	A high procedure cache hit rate is a desirable thing. You should strive for a hit ratio between 95-100%, with 95% being a good performance benchmark for procedure code reference. Note that when a database is first started, the procedure cache hit rate is not at an optimal level because all code being used is relatively new, and as such, must be read in from disk and placed into the cache. If, however, after a solid hour or two of steady database time, the procedure cache hit rate has not increased to desirable levels, you should look into the possibility of increasing the amount of memory allocated to the cache. You can drill down into the procedure cache to view the procedures currently in memory along with how much memory they are consuming. If there is not enough memory to load a requested procedure, or the maximum number of compiled objects is already in use, Sybase returns an error (normally a 701).
Correction	If a problem is found in Sybase servers, versions 11-12, you can increase the amount of the total memory configuration parameter or increase the percentage of memory allocated to the procedure cache (by default, the data cache assumes any free memory left over after Sybase has met its kernel and procedure cache needs). You should be careful when increasing the procedure cache alone, as this could increase query response times due to more physical I/O being performed. For Sybase 12.5, the total memory configuration parameter can again be increased to provide more memory for the Sybase server, but in 12.5, if you want to increase the size of the procedure cache, note that it is now configured in terms of literal size instead of a percentage of the overall configured memory. Once the procedure cache has been adjusted, monitor Sybase to see if the cache hit rate improves. If it does not, another increase may be necessary. Also, keep a careful eye on the actual machine's memory limits and swap activity. Increased swap activity can be indicative of too little memory left for the server machine.

Large I/O Hit Rate

Large I/O can be enabled by splitting the default or any named cache into pools. By default, Sybase performs I/O operations based on a 2-KB page size. For queries where pages are stored and accessed in a sequential manner, it is possible to read many more data pages in a single I/O operation. Large I/O can greatly reduce disk access time when the right situations exist. Operations that routinely perform large table scans, access image or text data, do bulk copies, scan the leaf level of nonclustered indexes, or initiate DBCC tasks can benefit from large I/O.

If large I/O has been configured and is being used, you should observe a high percentage of hits (the number of times large I/O could be performed vs. the number of times large I/O requests were denied by the server). If large I/O is not configured, no large I/O activity should be present.

Location	Sybase ASE Performance Analyst statistics > Home > Key Ratio Analysis pane Sybase ASE Performance Analyst statistics > Memory > Home> Key Ratio Analysis pane
Metrics	As you might expect, if large I/O is in use, a high hit rate is desirable. You should strive for a hit ratio between 90-100%, with 90% being a good performance benchmark.
Correction	If large I/O is configured, but you see a low hit rate, you should configure more caches for large I/O use.

Clean Buffer Grab Rate

As information is requested from users, buffers are moved into and out of the Sybase data cache. Pages are also modified in the cache (termed dirty buffers) and need to be written out to disk. If Sybase has to wait for a dirty buffer to be written out to disk before a requested buffer is placed into the cache, performance can suffer.

The clean buffer grab rate represents the percentage of time clean buffers were found and referenced in the cache as opposed to Sybase finding dirty buffers.

Location	Sybase ASE Performance Analyst statistics > Home > Key Ratio Analysis pane Sybase ASE Performance Analyst statistics > Memory > Home> Key Ratio Analysis pane
Metrics	Ideally, the clean buffer grab rate should stay at or near 100%.
Correction	A poor clean buffer grab rate for either the default or named caches could indicate that the cache size is too small. You can look into adjusting the total memory configuration parameter higher. Monitor the actual machine's memory limits and swap activity. Increased swap activity can be indicative of too little memory left for the server machine.

CPU Busy

The CPU Busy statistic represents the percent of the sample interval during which the Adaptive Server was busy.

Location	Sybase ASE Performance Analyst statistics > Home > Key Ratio Analysis pane Sybase ASE Performance Analyst statistics > Databases > Home> Server Analysis pane
Correction	If you find that your CPU Busy value is consistently above 70%, consider adding additional hardware.

I/O Busy

This metric represents the number of clock ticks in the sample interval during which the user task performed I/O operations.

Location	Sybase ASE Performance Analyst statistics > Home > Key Ratio Analysis pane Sybase ASE Performance Analyst statistics > I/O > Home> Bottleneck Analysis pane Sybase ASE Performance Analyst statistics > Databases > Home> Server Analysis pane

Metrics	High numbers indicate an I/O-intensive process. If idle time is also high, the application could be I/O bound.
Correction	To improve throughput: *Assign a higher priority to the application.* *Bind the application to a lightly loaded engine or engine group.* *Partition the application's data onto multiple devices.*

Server Idle

This metric represents the number of clock ticks in the sample interval during which the user task was idle.

Metrics

If you see high numbers for idle time and I/O Busy, the application could be I/O-bound.

Correction

To improve throughput:

- Assign a higher priority to the application.
- Bind the application to a lightly loaded engine or engine group.
- Partition the application's data onto multiple devices.

Bottleneck Analysis pane

When Sybase Adaptive Server is running, every connected process is either busy doing work or waiting to work. A process that is waiting might mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. And this is where wait-based or bottleneck analysis comes into play. DBAs use this form of performance analysis to determine if perceived bottlenecks in a database are contributing to a performance problem.

Bottleneck analysis is a valid method of measuring performance because it helps a DBA track where a database has been spending its time. If lock contention or heavy table scan activity has been dragging down database performance, a DBA can use bottleneck analysis to confirm the actual root cause. Once one or more wait events or other bottlenecks have been pinpointed as possible performance vampires, the DBA can drill down and oftentimes discover a fair amount of detail about which sessions and objects are causing the problem. The Performance Analyst Home page identifies the top system and session waits that might be currently a cause of lessened performance.

When using bottleneck analysis, a DBA cannot rely only on the information contained in the wait events that Sybase Adaptive Server provides. For example, a database might be suspect or offline. Such a failure will not be reflected in any wait event, but still represents a very real bottleneck to the database. In the same way that a DBA cannot depend on only a few ratios to properly carry out ratio-based performance analysis, an administrator must include other statistical metrics in their overall bottleneck analysis framework to obtain an accurate performance risk assessment. Performance Analyst works to identify bottlenecks in your database that fall outside of pure wait events so you can get a total picture of all stoppages in your system. The Performance Analyst Home page displays space-related bottlenecks as well as other sources of contention/bottlenecks that might be dragging down the overall performance of your system.

The following bottleneck indicators are used on the Performance Analyst Home page to succinctly communicate the general overall performance level of the monitored database:

Device I/O Contention	Modify Conflicts	Deadlocks
Suspect Databases	Databases Low on Space	Logs Low on Space
Tempdb Percent Full	Task Context Switches	Group Commit Sleeps

Blocked Connections		

Deadlocks

A deadlock occurs when two processes have a lock on a separate page or object and each wants to acquire a lock on the other process' page or object. Each waits for the other to release the necessary lock. Sybase constantly checks for deadlocks and, when found, chooses the transaction that has accumulated the least amount of CPU time and terminates the transaction. The server then rolls back that transaction and issues a notification of the event. The other process gets to move forward.

A deadlock occurs when two processes have a lock on a separate page or object and each wants to acquire a lock on the other process' page or object. Each waits for the other to release the necessary lock. Sybase constantly checks for deadlocks and, when found, chooses the transaction that has accumulated the least amount of CPU time and terminates it (the transaction). The server then rolls back that transaction and issues a notification of the event. The other process gets to move forward.

The deadlock statistic displays the number of current deadlocks in a Sybase server.

Location	Sybase ASE Performance Analyst statistics > Home > Bottleneck Analysis pane Sybase ASE Performance Analyst statistics > Users > Home> Bottleneck Analysis pane
Metrics	A nonzero number should alert you to the fact that application conflicts are likely being experienced by your user community.
Correction	Well-designed applications can minimize deadlocks by always acquiring locks in the same order. You should always do updates to multiple tables in the same order. Once Sybase discovers a deadlock, it takes action and remedies the situation. Embarcadero Performance Analyst makes it easier to discover how prevalent deadlock activity is on a system; preventing deadlocks from occurring in the first place is more difficult. Those responsible for writing systems can minimize deadlocks by ensuring that applications acquire locks in the same order. Likewise, you should always do updates and other DML that act on multiple tables in the same order. You can also shrink the amount of time that Sybase waits to check for deadlocks by modifying the deadlock checking period configuration parameter.

Suspect Databases

It can be hard to recover quickly from a suspect database. There are times when the cache of the database is suspect and not the actual database itself. If you stop and start the Sybase server, you can verify if the cache is the culprit. If the database itself is damaged, there could be a true recovery situation.

Location	Sybase ASE Performance Analyst statistics > Home > Bottleneck Analysis pane Sybase ASE Performance Analyst statistics > Databases > Home> Bottleneck Analysis pane
Correction	The suspect database can be dropped using the DBCC DBREPAIR DROPDB command. You then need to re-create the database and perform a recovery operation using the most recent database dump.

Device I/O Contention

Device I/O Contention reflects the number of times a task or process was put to sleep while waiting for a semaphore for a particular database device.

When a task or process involves physical I/O, Sybase first fills out the block I/O structure and links it to a per engine I/O queue. If two or more Sybase engines request an I/O structure from the same device at the exact same time, the server puts one of them to sleep where it waits for the semaphore it needs.

Location	Sybase ASE Performance Analyst statistics > Home > Bottleneck Analysis pane Sybase ASE Performance Analyst statistics > I/O > Home> Bottleneck Analysis pane
Correction	If consistent high numbers are viewed for Device I/O Contention, you can try redistributing the tables across devices (that are located on separate hard disks) or by adding devices (again, that are located on separate hard disks) and moving tables and indexes to them.

Modify Conflicts

For some database operations, Sybase uses a specialized protection mechanism to get exclusive access to a page without using actual restrictive page locks. Examples include accessing certain system tables and dirty reads on data pages. These actions need exclusive access to the page in question, even though they do not actually modify it. Modify conflicts record how many times these actions conflicted with other processes that need true database locks.

Tempdb Percent Full

The Tempdb Percent Full statistic shows the percentage of the Tempdb that is full.

Metrics

Tempdb is probably the most important resource in Adaptive Server shared by all applications and processes. It is a server-wide resource used for internal processing tasks like work tables, sorts, reformatting, and index builds.

Correction

Problems can be identified by a consistently high percentage registering or receiving a message indicating that your transactions are being suspended until space is made available in the log. In the short term, try using the “dump tran” DBCC command. If you suspect your application requires a very large tempdb, it should be placed on small fast disks. These disks should have no other data on them.

NOTE: On some systems, tempdb performs much faster using an operating system file rather than raw devices. This is because an operating system file uses buffering, which is faster than direct writes to a raw device.

Task Context Switches

Task Context Switches summarizes task-switch activity for all engines on SMP servers. You can use this statistic to observe the effect of reconfigurations.

Location	Sybase ASE Performance Analyst statistics > Home > Bottleneck Analysis pane Sybase ASE Performance Analyst statistics > Databases > Home> Bottleneck Analysis pane
Metrics	You might reconfigure a cache or add memory if tasks appear to block on cache search misses and to be switched out often. Then, check the data to see if tasks tend to be switched out more or less often.

Group Commit Sleeps

For databases needing high throughput, a large log I/O size is very important to prevent disk queuing problems on the transaction log. Group commit sleeps reflect the number of times a task performed a transactional commit operation and was put to sleep by the server until data in the log was written to disk.

Location	Sybase ASE Performance Analyst statistics > Home > Bottleneck Analysis pane Sybase ASE Performance Analyst statistics > I/O > Home> Bottleneck Analysis pane
Metrics	You should examine group commit sleeps in conjunction with the number of committed transactions (found on the Users performance category view).
Correction	A high count for group commit sleeps is not necessarily a problem if the server's transaction rate is low. If there are a significant number of transactions that result in group commit sleeps, and the log I/O size is greater than 2 KB, a smaller log I/O size can help to reduce commit time by causing more frequent page flushes. Other factors that can affect group commit sleeps are the size of the server run queue and the speed of the hard disk where the log is located.

Blocked Connections

A single blocking user has the potential to stop work for nearly all other processes on a small system, and can cause major headaches on large systems. Although Sybase supports flexible locking mechanisms, blocking lock situations do crop up. Blocks are most often caused by user processes holding exclusive locks and not releasing them via a proper COMMIT frequency.

The blocked statistic displays the number of current processes blocked by other processes.

Location	Sybase ASE Performance Analyst statistics > Home > Bottleneck Analysis pane Sybase ASE Performance Analyst statistics > Users > Home> Bottleneck Analysis pane
Metrics	While the lock contention rate is a better measure of the overall lock contention situation, if you consistently see positive numbers for the blocked statistic there may be a bottleneck for some processes. You can easily drill down and discover the exact process(es) holding locks that are blocking out other user activity.
Correction	Once you discover a blocking lock situation, you can normally remedy it by issuing a KILL against the offending process. This eliminates the user's stranglehold on the objects the user was accessing, and usually results in other user processes completing in an instant. Discovering the blocked lock situation is made easier by using tools like Performance Analyst, but preventing the blocking lock situation in the first place is tricky. The culprit of blocking lock scenarios is usually the application design, or the SQL being used within the application itself. Properly coding an application to reference database objects in an efficient order, and then using the right SQL to get the job done, is an art. The key to avoiding lock contention is to process user transactions in the quickest and most efficient manner possible - something not always easy to do. By default, all processes wait indefinitely for locks in Sybase. You can change this behavior by modifying the lock wait period configuration parameter, which limits the number of seconds that a process waits for a lock before timing out.

Databases Low on Space

The Databases Low on Space statistic indicates if any databases in the server are approaching low levels of available free space. Although good to know, you need a more detailed listing by database to determine where any actual space shortages exist in the server.

Location	Sybase ASE Performance Analyst statistics > Home > Bottleneck Analysis pane Sybase ASE Performance Analyst statistics > Space > Home> Bottleneck Analysis pane

Metrics	If any one database begins to approach 90% used space, and is continuing to dynamically grow, you should take action to prevent any future space allocation errors.
Correction	If a database is approaching its limit on space, you can either extend the database onto a new device, or if space exists on the device where the database currently resides, you can allocate more space for the database on the current device.

Logs Low on Space

This statistic indicates if any log for a database in the Sybase server is approaching low levels of available free space. Although good to know, a more detailed listing by database is needed to determine where any actual space shortages exist in the server.

Location	Sybase ASE Performance Analyst statistics > Home > Bottleneck Analysis pane Sybase ASE Performance Analyst statistics > Space > Home> Bottleneck Analysis pane
Metrics	If a database log's used space begins to approach 90%, you should take action to prevent any future space allocation errors. If the transaction log runs out of space, no transactions can take place in the database until you free up space in the log.
Correction	If a database log consistently approaches its limit on used space, there are a few actions a DBA can take to prevent a database from freezing. A backup (dump) of the log can be taken: *If log backups are not required for disaster recovery, the truncate log on checkpoint option can be set for the database. Setting this option deletes any space devoted to transactions in the log that have already completed when a database checkpoint occurs.* *You can extend the log onto a new device. Or, if space exists on the device on which the database log currently resides, you can allocate more space for the log on the current device.*

I/O Analysis pane

I/O statistics gives you a succinct overview of the datasource's overall performance. The ratios on the Home page are:

Server Reads	Server Writes	Transaction Log Writes
I/O Errors		

Server Reads

Total Server Reads reflect the total number of physical reads performed by the database server since the last refresh.

Location	Sybase ASE Performance Analyst statistics > Home > I/O Analysis pane
Metrics	Large numbers of physical reads could reflect a too small data or procedure cache. You should examine the data and procedure cache hit rates to determine the overall effectiveness of logical vs. physical I/O.

Server Writes

This value reflects the total number of physical writes performed by the database server since the last refresh.

Location	Sybase ASE Performance Analyst statistics > Home > I/O Analysis pane

>

Transaction Log Writes

Transaction Log Writes refers to the number of times Sybase wrote a transaction log page to disk since the last refresh. When the current log page becomes full, Sybase writes it out to disk. Sybase also writes transaction log pages to disk after a transaction commits.

Location	Sybase ASE Performance Analyst statistics > Home > I/O Analysis pane

I/O Errors

The I/O Errors value reflects total number of I/O errors (errors during read and write operations) encountered by the database server since the last refresh.

Location	Sybase ASE Performance Analyst statistics > Home > I/O Analysis pane

Storage Analysis pane

While DBAs focus on memory settings and tuning SQL, they frequently forget how dangerous and insidious storage problems can be. This is not a good mindset because storage headaches can play a major role in wrecking an otherwise well-running database.

Storage problems generally take one of two forms:

- The ‘hit-the-wall’ variety that can bring things to a complete standstill.
- The ‘performance vampire’ kind that slowly drains the performance of a database over time.

Storage problems have the capability to bring the curtain down on a database very quickly, as in the case of a transaction log running out of free space. But storage problems can also silently work behind the scenes to slowly, but surely rob a database of its performance. For example, a hub table in a busy database might be accessed very quickly when an application is first given life, but over time, if it develops a heavy forwarded row problem, it can cause things to run very differently.

The Storage Analysis section of the Performance Analyst Home page displays the Total Used and Free Space of all tablespaces in the database. It also provides a count of the number of databases, files, and file groups that exist on the server. For more detailed information on database space, you can use the Performance Analyst Space home page or the Space Detail.

The following ratios succinctly communicate the general overall performance levels of the datasource:

·Total Used	Total Free	Number of Databases
Number of Devices		

Total Used

The Total Used statistic shows the total amount of space currently in use by the Sybase Server.

Location	Sybase ASE Performance Analyst statistics > Home > Storage Analysis pane

Total Free

The Total Free statistic shows the total amount of free space available to the Sybase Server

Location	Sybase ASE Performance Analyst statistics > Home > Storage Analysis pane

>

Correction

A low percentage of free space indicates that additional space should be allocated to your database(s)/devices.

Number of Databases

This value indicates the total number of database currently defined on the Sybase server.

Location	Sybase ASE Performance Analyst statistics > Home > Storage Analysis pane

Number of Devices

This value indicates the total number of devices currently defined on the Sybase server.

Location	Sybase ASE Performance Analyst statistics > Home > Storage Analysis pane

Workload Analysis pane

When your phone starts ringing with complaints of performance slowdowns, some of the first questions you should answer are:

- Who is currently connected to the database?
- What resources are they using?
- What are they currently executing?

The Workload Analysis section of the Performance Analyst Home page provides insight into the leading resource hogs of a server, as well as a count of total active and inactive processes.

Workload Analysis statistics include:

Top I/O Process	Top Memory Process	Top CPU Process
Active Processes	Inactive Processes	

Top I/O Process

The Top I/O Process statistic identifies the current SPID with the highest percentage of I/O activity.

Location	Sybase ASE Performance Analyst statistics > Home > Workload Analysis pane
Metrics	If any one session uses more than 50% of a total resource (CPU, memory, etc.) drill-down into the session to investigate what work is currently executing.

Top Memory Process

The Top Memory Process statistic identifies the current SPID with the highest percentage of memory consumption.

Location	Sybase ASE Performance Analyst statistics > Home > Workload Analysis pane
Metrics	If any one session uses more than 50% of a total resource (CPU, memory, etc.) drill-down into the session to investigate what work is currently executing.

Top CPU Process

The Top CPU Process statistic identifies the current SPID with the highest percentage of CPU consumption.

Location	Sybase ASE Performance Analyst statistics > Home > Workload Analysis pane
Metrics	If any one session uses more than 50% of a total resource (CPU, memory, etc.) drill-down into the session to investigate what work is currently executing.

Active Processes

The Active Processes statistic is the total number of active and open threads reported on the Server. Active Processes displays the number of processes actively performing work.

Location	Sybase ASE Performance Analyst statistics > Home > Workload Analysis pane

Inactive Processes

The Inactive Sessions statistic represents the total number of threads logged onto the server that are currently idle.

Location	Sybase ASE Performance Analyst statistics > Home > Workload Analysis pane
Metrics	A large number of inactive processes could indicate user sessions that have mistakenly been left logged on. Since each user thread consumes a portion of memory on the server, to reduce resource usage, you should sever any sessions that do not need a connection.
Correction	Drill down into the Session Waits page and check sessions that have many seconds idle and/or that have been logged on for very long periods of time, as indicated by the logon time column. After verifying that a session is no longer necessary, you can KILL it.

Memory

In addition to a Home page, the Memory category of Sybase Performance Analyst includes the following tabbed pages:

Data Cache Tab	Procedure Cache Tab	

Home

The Memory performance category view displays the following vital Sybase memory statistics:

Key Ratio Analysis pane	Bottleneck Analysis pane	Metadata Cache Ratio Analysis pane
Memory Analysis pane	Workload Analysis pane	

Key Ratio Analysis pane

Database performance analysts typically use one of two methods for examining the performance levels of a database - ratio-based or wait/bottleneck-based. Ratio-based analysis involves examining a number of key database ratios that can be used to indicate how well a database is running. Performance ratios serve as very good roll-up mechanisms for busy DBA's to use for at-a-glance performance analysis. Many DBAs have large database farms to contend with and cannot spend time checking detailed wait-based analysis outputs for each and every database they oversee. Succinctly presented performance ratios can assist in such situations by giving DBAs a few solid indicators that can be quickly scanned to see if any database needs immediate attention.

While there are certainly many opinions as to what rules to follow, there are some standards that should always be adhered to. To start with, many of the formulas that make up ratio-based analysis must be derived from delta measurements instead of cumulative statistics. Many of the global ratios that a DBA will examine come from the master.sysperfinfo performance table. This table maintains a count of various statistics for the server and particular databases since the server was brought up. For servers that are kept up for long periods of time, these values can grow quite large and will impact how a particular ratio that a DBA might be looking at is interpreted. However, if delta statistics are used (taking, for a specified sampling period, the before and after counts of each statistic that make up a ratio), then an accurate and current portrayal of the various ratios can be had.

A final thing to remember about using ratio-based analysis is that, while there are a number of rules of thumb that can be used as starting points to begin the evaluation of database performance, DBAs must determine each SQL Server's individual 'personality' with respect to the key performance ratios. Some hard and fast rules simply do not apply to every database. The danger in using blanket ratio standards is that they can lead the DBA to haphazardly take action, which can at times contribute nothing to the situation, and sometimes even degrade performance.

The following memory ratios are used on the Performance Analyst Home page to succinctly communicate the general overall memory performance levels of the monitored database:

Dirty Read Restarts	Dirty Read Requests	Procedure Requests
Procedure Reads from Disk		

The following statistics, available on this pane, are duplicates of statistics appearing on the Key Ratio Analysis pane of the Home page:

Data Cache Hit Ratio	Procedure Plan Hit Ratio	Large I/O Hit Rate
Clean Buffer Grab Rate		

Dirty Read Restarts

Sybase allows dirty reads, which are reads of uncommitted data. To accomplish a dirty read, Sybase uses a special lightweight protection mechanism to gain access to an object without using actual page locks. A dirty read restart occurs when a dirty read is active on an object page, and another process makes changes to the page that cause the page to be deallocated in memory. The scan for the dirty read must be restarted. The amount shown for dirty read restarts are the number of restarts that occurred since the last Embarcadero Performance Analyst refresh.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane
Metrics	Values observed much above zero should serve as a signal that application modifications can be in order. Most applications should do everything possible to avoid restarts because of the large overhead they incur.
Correction	If the numbers observed for dirty read restarts are significant, you can want to look into modifying applications that use dirty reads to accomplish data acquisition.

Dirty Read Requests

Sybase allows dirty reads, which are reads of uncommitted data. To accomplish a dirty read, Sybase uses a special lightweight protection mechanism to gain access to an object without using actual page locks. This statistic displays the number of dirty reads that occurred since the last refresh in Embarcadero Performance Analyst.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane
Metrics	Dirty read page requests can incur significant overhead if they are observed with many dirty read restarts.

Procedure Requests

The Procedure Requests statistic reports the number of times that stored procedures were executed since Embarcadero Performance Analyst was last refreshed. Such a request could use either an unused copy of the procedure's query plan in memory or if no such copy exists, the procedure must be read in from disk.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane

Procedure Reads from Disk

The Procedure Reads from Disk statistic reports the number of times since Embarcadero Performance Analyst was last refreshed that stored procedures were read from disk rather than copied in the procedure cache.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane
Metrics	You should examine this number in conjunction with the overall procedure cache hit rate. Observing large numbers in this statistic indicates a lower than ideal procedure cache hit rate. Note that when a database is first started, this statistic is likely larger than desired because all code being used is relatively new and as such, must be read in from disk and placed into the cache. If, however, after a solid hour or two of steady database time, the procedure cache hit rate has not increased to desirable levels and this statistic continues to sport high numbers, you should look into the possibility of increasing the amount of memory allocated to the cache.
Correction	If a problem is found in Sybase servers, versions 11-12, you can increase the amount of the total memory configuration parameter and/or increase the percentage of memory allocated to the procedure cache (by default, the data cache assumes any free memory left over after Sybase has met its kernel and procedure cache needs). Take care when increasing the procedure cache alone, as this could increase query response times due to more physical I/O being performed. For Sybase 12.5, the total memory configuration parameter can again be increased to provide more memory for the Sybase server, but in 12.5, if you want to increase the size of the procedure cache, note that it is now configured in terms of literal size instead of a percentage of the overall configured memory. Once the procedure cache has been adjusted, monitor Sybase to see if the cache hit rate improves. If it does not, another increase may be necessary. Also, keep a careful eye on the actual machine's memory limits and swap activity. Increased swap activity can be indicative of too little memory left for the server machine.

Bottleneck Analysis pane

When Sybase Adaptive Server is running, every connected process is either busy doing work or waiting to work. A process that is waiting might mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. And this is where wait-based or bottleneck analysis comes into play. DBAs use this form of performance analysis to determine if perceived bottlenecks in a database are contributing to a performance problem.

>

Bottleneck analysis is a valid method of measuring performance because it helps a DBA track where a database has been spending its time. If lock contention or heavy table scan activity has been dragging down database performance, a DBA can use bottleneck analysis to confirm the actual root cause. Once one or more wait events or other bottlenecks have been pinpointed as possible performance vampires, the DBA can drill down and oftentimes discover a fair amount of detail about which sessions and objects are causing the problem.

Memory bottlenecks can definitely cause performance degradation in an otherwise well-running database. Typically, these bottlenecks center around the server latches, which are lightweight locks used to protect certain resources in memory. To help you identify such problems, the following statistics are presented on the Performance Analyst Memory home page:

Log Semaphore Contention	Latch Contention	Waits on Execution Plans
Large I/Os Denied	Buffers Grabbed Dirty	Procedure Removals

Log Semaphore Contention

The Log Semaphore Contention statistic measures the total number of times tasks requested a log semaphore. This includes those granted immediately and those for which the task had to wait.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	In high throughput environments with a large number of concurrent users committing transactions, you should expect a certain amount of contention for the log semaphore.

Latch Contention

The Latch Contention statistic reports the number of times a task was switched out because it needed to wait for a latch. If your user tables use only allpages-locking, this latch contention is taking place either on a data-only-locked system table or on allocation pages. If your applications use data-only-locking, the contention reported here includes all waits for latches, including those on index pages and OAM pages as well as allocation pages.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	In SMP environments where inserts and expanding updates are extremely high, so that page allocations take place very frequently, contention for the allocation page latch can reduce performance. Normally, Adaptive Server allocates new pages for an object on an allocation unit that is already in use by the object and known to have free space. For each object, Adaptive Server tracks this allocation page number as a hint for any tasks that need to allocate a page for that object. When more than one task at a time needs to allocate a page on the same allocation unit, the second and subsequent tasks block on the latch on the allocation page.
Correction	You can specify a "greedy allocation" scheme, so that Adaptive Server keeps a list of eight allocation hints for page allocations for a table. The effect of dbcc tune(des_greedyalloc) are not persistent, so you need to reissue the commands after a reboot. You should use this command only if all of the following are true: *You have multiple engines. It is rarely useful with fewer than four engines.* *A large number of pages are being allocated for the object. You can use sp_spaceused or optdiag to track the number of pages.* *The latch contention counter shows contention.*

Waits on Execution Plans

The Waits on Execution Plans value represents the number of times that a process attempting to use sp_showplan had to wait to acquire read access to the query plan.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	Query plans may be unavailable if you run sp_showplan before the compiled plan is completed or after the query plan finished executing. In these cases, Sybase tries to access the plan three times and then returns a message to the user.

Large I/Os Denied

The Large I/Os Denied statistic represents the number of times that large I/O could not be performed.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	Adaptive Server cannot perform large I/O for the following reasons: *If any page in a buffer already resides in another pool.* *When there are no buffers available in the requested pool.* *On the first extent of an allocation unit, since it contains the allocation page, which is always read into the 2K pool.*
Correction	If a high percentage of large I/Os were denied, it indicates that using larger pools might not be as effective as it could be. If a cache contains a large I/O pool, and queries perform both 2K and 16K I/O on the same objects, there will always be some percentage of large I/Os that cannot be performed because pages are in the 2K pool. If more than half of the large I/Os were denied, and you are using 16K I/O, try moving all of the space from the 16K pool to the 8K pool. Re-run the test to see if total I/O is reduced. Note that when a 16K I/O is denied, Adaptive Server does not check for 8K or 4K pools, but uses the 2K pool.

Buffers Grabbed Dirty

The Buffers Grabbed Dirty statistic represents the number of times Sybase found dirty buffers since the last refresh in Embarcadero Performance Analyst.

As users request information, buffers are moved into and out of the Sybase data cache. Pages are also modified in the cache (termed dirty buffers) and need to be written out to disk. If Sybase has to wait for a dirty buffer to be written out to disk before a requested buffer is placed into the cache, performance can suffer.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	Ideally, the dirty buffer grab statistic should stay close to zero.
Correction	Seeing high numbers for this statistic could indicate that the cache size is too small. You should look into carefully adjusting the total memory configuration parameter to be higher. However, keep a careful eye on the actual machine's memory limits and swap activity. Increased swap activity can indicate too little memory left for the server machine.

Procedure Removals

The Procedure Removals statistic reports the number of times that stored procedures were aged out of the procedure cache since Embarcadero Performance Analyst was last refreshed.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane

Metrics	High numbers, along with a lower than desired procedure cache hit rate, could indicate a procedure cache that is too small.
Correction	If a problem is found in Sybase servers, versions 11-12, you can increase the amount of the total memory configuration parameter and/or increase the percentage of memory allocated to the procedure cache (by default, the data cache assumes any free memory left over after Sybase has met its kernel and procedure cache needs). Take care when increasing the procedure cache only, as this could increase query response times due to more physical I/O being performed. For Sybase 12.5, the total memory configuration parameter can again be increased to provide more memory for the Sybase server, but in 12.5, if you wish to increase the size of the procedure cache, note that it is now configured in terms of literal size instead of a percentage of the overall configured memory. Once the procedure cache has been adjusted, monitor Sybase to see if the cache hit rate improves. If it does not, another increase may be necessary. Also, keep a careful eye on the actual machine's memory limits and swap activity. Increased swap activity can indicate too little memory left for the server machine.

Metadata Cache Ratio Analysis pane

The following statistics are used on the Performance Analyst for Sybase Memory Home Page to succinctly communicate the general overall performance levels of the memory structures:

Open Objects	Open Indexes	Open Databases
Object Spinlock Contention	Index Spinlock Contention	Hash Spinlock Contention

Open Objects

The Open Objects statistic represents the use of the Adaptive Server metadata cache for open objects. The statistic includes number of active objects, free objects, and the maximum number of objects used.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Metadata Cache Ratio Analysis pane
Metrics	When Adaptive Server accesses an object, it needs to read information about it in the corresponding system table: sysobjects. The metadata cache for objects lets Adaptive Server access the information that describes it in the sysobjects row directly in its in-memory structure. This improves performance because Adaptive Server bypasses expensive calls that require disk access. It also reduces synchronization and spinlock contention when Adaptive Server has to retrieve object information at runtime. Managing individual metadata caches for databases, indexes, or objects is beneficial for a database that contains a large number of indexes and objects, and where there is high concurrency among users.
Correction	You can use this information to set the configuration parameter’s number of open objects.

Open Indexes

The Open Indexes statistic represents the use of the Adaptive Server metadata cache for open indexes. The statistic includes the number of active indexes, free indexes, and the maximum number of indexes.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Metadata Cache Ratio Analysis pane
Metrics	When Adaptive Server accesses an index, it has to access information about it in the respective system table: sysindexes. The metadata cache for indexes lets Adaptive Server access the information that describes it in the sysindexes row directly in its in-memory structure. This improves performance because it allows the Adaptive Server to bypass expensive calls that require disk access. It also reduces synchronization and spinlock contention when Adaptive Server has to retrieve index information at runtime. Managing individual metadata caches for databases, indexes, or objects is beneficial for a database that contains a large number of indexes and objects, and where there is high concurrency among users.
Correction	You can use this information to set the configuration parameter’s number of open indexes.

Open Databases

The Open Databases statistic shows the use of the Adaptive Server metadata cache for open databases. This statistic charts the number of active databases, free databases, and the maximum number of databases used.

When Adaptive Server opens a database, it has to access information about it in the respective system table: sysdatabases. The metadata cache for databases lets Adaptive Server access the information that describes it in the sysdatabases row directly in its in-memory structure. This improves performance because it allows the Adaptive Server to bypass expensive calls that require disk access. It also reduces synchronization and spinlock contention when Adaptive Server has to retrieve database information at runtime.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Metadata Cache Ratio Analysis pane
Metrics	Managing individual metadata caches for databases, indexes, or objects is beneficial for a database that contains a large number of indexes and objects, and where there is high concurrency among users.
Correction	You can use this information to set the configuration parameter's number of open databases.

Object Spinlock Contention

The Object Spinlock Contention statistic represents spinlock contention on the object descriptor caches.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Metadata Cache Ratio Analysis pane
Correction	You can use this information to tune the configuration parameter's open object spinlock and open index spinlock ratios. If the reported contention is more than 3%, decrease the value of the corresponding parameter to lower the number of objects or indexes that are protected by a single spinlock.

Index Spinlock Contention

The Index Spinlock Contention value represents spinlock contention on the index descriptor caches.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Metadata Cache Ratio Analysis pane
Correction	You can use this information to tune the configuration parameter's open object spinlock and open index spinlock ratios. If the reported contention is more than 3%, decrease the value of the corresponding parameter to lower the number of objects or indexes that are protected by a single spinlock.

Hash Spinlock Contention

This metric represents the contention for the spinlock on the index metadata cache hash table.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Metadata Cache Ratio Analysis pane
Correction	You can use this information to tune the open index hash spinlock ratio configuration parameter. If the reported contention is greater than 3%, decrease the value of the parameter.

Memory Analysis pane

The following statistics are used on the Performance Analyst for Sybase Memory Home Page to succinctly communicate the general overall performance levels of the memory structures:

Buffer Cache	Procedure Cache	Connection Memory
Open Database Memory	Open Object Memory	Open Index Memory

>

Buffer Cache

The Buffer Cache statistic displays the amount of memory currently allocated to the Buffer Cache.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Memory Analysis pane

Procedure Cache

The Procedure Cache value displays the amount of memory currently allocated to the procedure cache.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Memory Analysis pane
Metrics	The memory allocated for the procedure cache holds the optimized query plans (and occasionally trees) for all batches, including any triggers.

Connection Memory

The Connection Memory value indicates the amount of memory currently allocated to user connections.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Memory Analysis pane
Metrics	Each connection requires approximately 60 KB.

Open Database Memory

The Open Database Memory value represents the current amount of memory allocated in the number of open databases parameter.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Memory Analysis pane

Open Object Memory

This value represents the current amount of memory allocated in the number of open objects parameter.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Memory Analysis pane

Open Index Memory

This value represents the current amount of memory allocated in the number of open indexes parameter.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Memory Analysis pane

Workload Analysis pane

The following statistics are used on the Performance Analyst for Sybase Memory Home Page to succinctly communicate the general overall performance levels of the memory structures:

- Top Memory Hogs

>

Top Memory Hogs

This metric identifies the process that currently is using the highest percentage of memory in the database.

Location	Sybase ASE Performance Analyst statistics > Memory > Home > Workload Analysis pane
Metrics	It is not uncommon for one or two users to cause the majority of runtime problems that plague a server. The problem could be a runaway process, an untuned batch procedure, or other user-initiated operation. If your database server does not have an overabundance of memory, then you should periodically check to see who your heavy memory users are along with the total percentage of memory each takes up. If you see one or two users who have more than 25-50% of the total memory usage, then you should further investigate the sessions to see what activities they are performing.

Data Cache Tab

The Data Cache tab of the Memory Detail View includes the following sections:

Cache Configurations	Cache Efficiency	Cached Objects
Cache I/O Activity		

The following statistic, available on this tab, duplicates a statistic available on the Key Ratio Analysis pane of the Home page:

- Data Cache Hit Ratio

Cache Configurations

The Cache Configuration view represents the default and named data caches and the memory configuration in MB.

Location	Sybase ASE Performance Analyst statistics > Memory > Data Cache Tab

Cache Efficiency

The Cache Efficiency section summarizes behavior for the default data cache and all named data caches combined. The table below describes the information available on the Cache Efficiency detail view:

Column	Description
Cache Name	The name of the specified cache
Hit Rate	The percentage of cache hits compared to the total number of cache searches.
Percent Used	The percentage of searches using this cache as a percentage of searches across all caches.
Spinlocks	The number of spinlocks.
LRU Buffers	The number of buffers that were retrieved and replaced in the Least Recently Used section of the pool.
MRU Buffers	The number of buffers that were retrieved and replaced in the Most Recently Used section of the pool.
Large I/Os	The size (bytes) of the I/O buffer for the pool.
Dirty Reads	The number of dirty reads.

Location	Sybase ASE Performance Analyst statistics > Memory > Data Cache Tab

>

Metrics	Interpreting cache hit data requires an understanding of how the application uses each cache. In caches that are created to hold specific objects such as indexes or look up tables, cache hit ratios can reach 100%. In caches used for random point queries on huge tables, cache hit ratios can be quite low but still represent effective cache use. This data can also help you determine if adding more memory would improve performance. For example, if "Cache Hits" is high, adding memory probably would not help much.
Correction	You can compare the "percent used" value for each cache to determine if there are caches that are over- or under-utilized. If you decide that a cache is not well utilized, you can: *Change the cache bindings to balance utilization.* *Resize the cache to correspond more appropriately to its utilization.*

Cached Objects

The Cached Objects section shows the objects that are currently cached. The table below describes the information available on the Cached Objects detail view:

Column	Description
Cache Name	The name of the specified cache.
Database Name	The name of the database where the object resides.
Object Owner	The owner of the object.
Object Name	The name of the object.
Object Type	The type of the object.
Cached (KB)	The amount, in Kilobytes, of the cached object.
Processes Accessing	The number of processes accessing the specified object.

Location	Sybase ASE Performance Analyst statistics > Memory > Data Cache Tab

Cache I/O Activity

The Cache I/O Activity view shows how many searches were directed to each cache. The table below describes the information available on the Cache I/O Activity view:

Column	Description
Cache Name	The name of the specified cache.
Physical Reads	The number of physical reads associated with the specified cache.
Pages Touched	The number of pages touched by the cache.
Pages Read	The number of pages read by the cache.
Buffers to MRU	The number of buffers that were retrieved and replaced in the Most Recently Used section of the pool.
Buffers to LRU	The number of buffers that were retrieved and replaced in the Least Recently Used section of the pool.
I/O Buffer Size	The size (bytes) of the I/O buffer for the pool.
Allocated (KB)	The number of KB that have been allocated for the pool.
Stalls	The number of times a process had to wait for a free procedure cache buffer when installing a stored procedure into cache.

>

Location	Sybase ASE Performance Analyst statistics > Memory > Data Cache Tab
Metrics	You can compare the values in this view for each cache to determine if there are caches that are over- or under-utilized.
Correction	If you determine that a cache is not used as well as you would like, you can: *Change which objects are bound to each cache to balance use.* *Resize the cache to correspond more appropriately to its use.*

Procedure Cache Tab

The Procedure Cache tab of the Memory Detail View includes the following sections:

Procedure Cache Activity	Procedure Cache Detail	Current Procedure Usage

The following statistic, available on this tab, duplicates astatistic available on the Key Ratio Analysis pane of the Home page:

- Procedure Plan Hit Ratio

Procedure Cache Activity

The Procedure Cache Activity view summarizes the key metrics concerning the Procedure Cache. The table below describes the information available on the Procedure Cache Activity view:

Column	Description
Requests	The number of stored procedures requested from the cache.
Loads	The number of stored procedures loaded into cache.
Writes	The number of times a stored procedure was normalized and the tree was written back to sysprocedures. This happens when a stored procedure (or trigger) is created during the interval between polls. A degradation in performance can occur if an application program generates stored procedures.
Stalls	The number of times a process had to wait for a free procedure cache buffer when installing a stored procedure into cache.

Location	Sybase ASE Performance Analyst statistics > Memory > Procedure Cache Tab

Procedure Cache Detail

The Procedure Cache Detail view summarizes the key metrics concerning the Procedure Cache. The table below describes the information available on the Procedure Cache Detail view:

Column	Description
Object Name	The name of the object.
Object Owner	The object's owner.
Object Type	The type of object.

>

Column	Description
Database	The name of the database where the object resides.
Compile Date	The last month/day/year object data was compiled.
Memory Used (KB)	The amount of memory used by the object in kilobytes.

Location	Sybase ASE Performance Analyst statistics > Memory > Procedure Cache Tab

Current Procedure Usage

The Current Procedure Usage view displays the number of procedures currently in use. The table below describes the information available in the Current Procedure Usage view:

Column	Description
SPID	The process name.
User Name	The name of the user attached to the process.
Database	The database attached to the process.
Object Owner	The object's owner.
Object Name	The object's name.
Object Type	The object's type.
Memory Used (KB)	The amount of memory the procedure used in kilobytes.
Program Name	The name of the program where the procedure is used.

Location	Sybase ASE Performance Analyst statistics > Memory > Procedure Cache Tab

I/O

In addition to a Home page, the I/O category of Sybase Performance Analyst includes the following tabbed pages:

Databases tab	Devices tab	Engines Tab
Network tab	Objects tab	

Home

The I/O performance category page displays the following vital Sybase I/O statistics:

Bottleneck Analysis pane	Key System Analysis pane	Network Analysis pane
Transaction Log Analysis pane	Workload Analysis pane	

Key System Analysis pane

The following statistics are used on the Performance Analyst for Sybase I/O Home Page to succinctly communicate the general overall performance levels of I/O:

Physical Reads	APF Reads	Physical Writes
Procedure Requests from Disk	Data Cache Searches	Checkpoints
Large I/Os Performed		

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Key System Analysis pane

Physical Reads

Physical Reads represents the total number of physical reads performed by the database server since the last refresh.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	Large numbers of physical reads could reflect a too small data or procedure cache. You should examine the data and procedure cache hit rates to determine the overall effectiveness of logical vs. physical I/O.

APF Reads

This metric represents the number of pages that were read by asynchronous prefetch since the last refresh.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	Asynchronous prefetch tries to find a page it needs to read in the data cache using a quick scan without holding the cache spinlock. If that does not succeed, it then performs a thorough scan holding the spinlock.

Physical Writes

This value reflects total number of physical writes performed by the database server since the last refresh.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Key System Analysis pane

Procedure Requests from Disk

The Procedure Requests from Disk value represents the number of times that stored procedures were read from disk rather than found and copied in the procedure cache.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	If you see a relatively high number, it could indicate that the procedure cache is too small.
Correction	You should increase the size of the memory allocated to the procedure cache memory structure.

Data Cache Searches

The Data Cache Searches value represents the number of times that a needed page was found in any cache.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Key System Analysis pane
Correction	If you see a high number of cache misses, you should investigate statistics for each cache.

Checkpoints

The Checkpoints value represents the combined number of normal and free checkpoints that occurred during the sample interval.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	Checkpoints write dirty pages (pages that have been modified in memory, but not written to disk) to the database device. Adaptive Server's automatic (normal) checkpoint mechanism works to maintain a minimum recovery interval. By tracking the number of log records in the transaction log since the last checkpoint was performed, it estimates whether the time required to recover the transactions exceeds the recovery interval. If so, the checkpoint process scans all data caches and writes out all changed data pages. When Adaptive Server has no user tasks to process, a housekeeper task begins writing dirty buffers to disk. These writes are done during the server's idle cycles, so they are known as "free writes." They result in improved CPU use and a decreased need for buffer washing during transaction processing. If the housekeeper process finishes writing all dirty pages in all caches to disk, it checks the number of rows in the transaction log since the last checkpoint. If there are more than 100 log records, it issues a checkpoint. This is called a "free checkpoint" because it requires very little overhead. In addition, it reduces future overhead for normal checkpoints.

Large I/Os Performed

The Large I/Os Performed statistic measures the number of times that a requested large I/O was performed.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	Adaptive Server cannot perform large I/O for the following situations: *If any page in a buffer already resides in another pool.* *When there are no buffers available in the requested pool.* *On the first extent of an allocation unit, since it contains the allocation page, which is always read into the 2K pool.*
Correction	If a high percentage of large I/Os were denied, it indicates that the use of the larger pools might not be as effective as it could be. If a cache contains a large I/O pool, and queries perform both 2K and 16K I/O on the same objects, there will always be some percentage of large I/Os that cannot be performed because pages are in the 2K pool. If more than half of the large I/Os were denied, and you are using 16K I/O, try moving all of the space from the 16K pool to the 8K pool. Re-run the test to see if total I/O is reduced. Note that when a 16K I/O is denied, Adaptive Server does not check for 8K or 4K pools, but uses the 2K pool.

Bottleneck Analysis pane

The following statistics are used on the Performance Analyst for Sybase I/O Home Page to succinctly communicate the general overall performance levels of I/O:

Delayed Disk IOs	Outstanding Disk I/Os	Disk I/O Structure Contention
Disk I/O Structure Contention	Total I/O	I/O Errors
Hottest Devices		

The following statistic, available on this pane, duplicates a statistic available on the Key Ratio Analysis pane of the Home page:

- I/O Busy

>

The following statistics, available on this pane, are duplicates of statistics available on the Bottleneck Analysis pane of the Home page:

Device I/O Contention	Group Commit Sleeps	

Delayed Disk IOs

The Delayed Disk IOs value is the number of I/Os delayed when the limit on disk I/O structures is reached.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Metrics	When Adaptive Server exceeds the number of available disk I/O control blocks, I/O is delayed because Adaptive Server requires that tasks get a disk I/O control block before initiating an I/O request.
Correction	If you see a nonzero value for delayed disk I/Os, try to add to the number of available disk I/O control blocks by increasing the configuration parameter.

Outstanding Disk I/Os

The Outstanding Disk I/Os statistic reflects the maximum number of I/Os pending for Sybase since the last refresh.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Metrics	If you see nonzero numbers for both outstanding and delayed disk I/Os, there could be a problem in many areas.

Disk I/O Structure Contention

When a Sybase server begins to experience I/O delays, the result can be a very dissatisfied user community. When such problems begin to occur, you should investigate Sybase or operating system limits. It could be that I/O operations are being blocked by one or both.

The Disk I/O Structures statistic represents the number of I/O delays caused by Sybase reaching the limit on disk I/O structures. When Sybase exceeds the number of available disk I/O control blocks, I/O is deferred because Sybase requires that any task get a disk I/O control block before beginning an I/O request.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Metrics	Consistent numbers above zero should be a cause for concern.
Correction	If you continue to see nonzero numbers for this statistic, you can try increasing the number of available disk I/O control blocks by increasing the configuration parameter disk I/O structures.

Total I/O

The Total I/O statistic represents the total number of physical reads and writes.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane

I/O Errors

I/O error rate reflects total number of I/O errors (errors during read and write operations) encountered by the server since the last refresh inside Performance Analyst. The I/O errors rate is a percentage based on Total I/O (the sum of the physical reads and writes).

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Metrics	You should observe few, if any errors.
Correction	If you notice any errors, you should check the Sybase error log for details.

I/O Busy

This metric represents the number of clock ticks in the sample interval during which the user task performed I/O operations.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Metrics	High numbers indicate an I/O-intensive process. If idle time is also high, the application could be I/O bound.
Correction	The application might achieve better throughput if you assign it a higher priority, bind it to a lightly loaded engine or engine group, or partition the application's data onto multiple devices.

Hottest Devices

The Hottest Devices view identifies the configured devices engaged in the most I/O activity. In Sybase, when a task or process needs to perform physical I/O, the server fills out the actual block I/O structure and links it to a per-engine I/O queue. If two Sybase engines request the same I/O structure from the exact same device at the identical time, one of them is put to sleep by the server and it waits for the semaphore it needs.

Device I/O Contention reflects the number of times a task or process was put to sleep while waiting for a semaphore for a particular database device.

The table below describes the information available on the Hottest Devices view:

Column	Description
Device Name	The name of the device.
APF Reads	The number of asynchronous pre-fetch buffers read.
Physical Reads	The total number of physical reads performed by the database server since the last refresh
Physical Writes	The total number of physical writes performed by the database server since the last refresh
Total I/O	The combined total of read and write operations.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Correction	If you consistently see high numbers for Device I/O Contention, you can try redistributing the tables across devices (that are located on separate hard disks) or by adding devices (again, that are located on separate hard disks) and moving tables and indexes to them.

Network Analysis pane

The following statistics are used on the Performance Analyst for Sybase I/O page to succinctly communicate the general overall performance levels of I/O:

Packets Sent	Packets Received	Network Delays
Network Errors		

Packets Sent

Packets Sent displays the number of times Sybase sent a packet to a client application.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Network Analysis pane

Packets Received

Total Packets Received reflects the number of times Sybase received a packet from a client application.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Network Analysis pane

Network Delays

The Sybase server normally sends and receives network packets at a regular rate. If the network begins to be saturated, Sybase can experience delays in network I/O. The Network Delays statistic reflects the number of times network I/O activity was delayed.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Network Analysis pane
Metrics	If you see a network contention rate greater than 1%, it could indicate a challenged network structure.
Correction	Beyond ensuring that the existing network is sufficient for handling the current load (database and non-database), other items to look at from a Sybase-specific standpoint include validating that only the necessary amount of data is being sent and returned to requesting users. You can also examine the default packet size and see if it is too small for the average packet size being sent/received by the Sybase server.

Network Errors

The Network Errors statistic reflects the number of times that network errors were detected by Sybase while reading and writing packets.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Network Analysis pane
Metrics	If you consistently see a value much above zero, it could indicate the network structure is challenged.
Correction	Beyond ensuring that the existing network is sufficient for handling the current load (database and non-database), other items to look at from a Sybase-specific standpoint include validating that only the necessary amount of data is being sent and returned to requesting users. You can also examine the default packet size and see if it is too small for the average packet size being sent/received by the Sybase server.

>

Transaction Log Analysis pane

The following statistics are used on the Performance Analyst for Sybase I/O page to succinctly communicate the general overall performance levels of I/O:

Log Writes	Log Allocations	ULC Flushes
Average Number of Writes per Log Page		

Log Writes

Transaction Log Writes refers to the number of times Sybase wrote a transaction log page to disk since the last refresh. When the current log page becomes full, Sybase writes it out to disk. Sybase also writes transaction log pages to disk after a transaction commits.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Transaction Log Analysis pane

Log Allocations

Transaction Log Allocations refers to the number of times since the last refresh that additional pages were allocated to the transaction log. This statistic gives you a feel for the rate of transaction log growth.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Transaction Log Analysis pane

ULC Flushes

ULC Flushes is the total number of times that ULCs (user log caches) were flushed to a transaction log during the sample interval.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Transaction Log Analysis pane

Average Number of Writes per Log Page

The Average Number of Writes is the average number of times each log page was written to disk.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Transaction Log Analysis pane
Metrics	In high throughput applications, this number should be as low as possible. If the transaction log uses 2K I/O, the lowest possible value is 1; with 4K log I/O, the lowest possible value is .5, since one log I/O can write 2 log pages. In low throughput applications, the number will be significantly higher. In very low throughput environments, it may be as high as one write per completed transaction.

Workload Analysis pane

The following statistics are used on the Performance Analyst for Sybase I/O page to succinctly communicate the general overall performance levels of I/O:

- Top I/O Hogs

Top I/O Hogs

The Top I/O Process statistic identifies the processes that has currently caused the most I/O usage on the database. The table below describes the information available from the Top I/O Hogs view of the I/O page:

Column	Description
SPID	The process ID.
Login	The login name the session is using.
Physical I/O	The physical amount of I/O the process is using.
% Used	Percent of total I/O on the server this process is consuming.

Location	Sybase ASE Performance Analyst statistics > I/O > Home > Workload Analysis pane
Correction	If any one session uses more than 50% of a total resource (CPU, memory, etc.) you should drill down into that particular session and investigate the cause.

Databases tab

The database I/O statistic displays the major I/O statistics as they relate to each associated database. The table below describes the information available from the I/O Databases detail tab:

Column	Description
Database Name	The name of the database.
Logical Reads	The number of logical reads associated with the database.
Physical Reads	The number of physical reads associated with the database.
APF Reads	The number of asynchronous pre-fetch buffer reads associated with the database.
Pages Read	The number of pages read associated with the database.
Physical Writes	The number of physical writes associated with the database.
Pages Written	The number of pages written to the associated database.
Rows Inserted	The number of rows inserted in the associated database.
Rows Updated	The number of rows updated in the associated database.
Rows Deleted	The number of rows deleted from the associated database.
Lock Requests	The number of requests for a lock on database objects.
Lock Waits	The number of times a task waited for a lock on a database object.

Correction

The major guidelines for improving I/O performance in Adaptive Server are as follows:

- Spread data across disks to avoid I/O contention.
- Isolate server-wide I/O from database I/O.
- Separate data storage and log storage for frequently updated databases.
- Keep random disk I/O away from sequential disk I/O.
- Mirror devices on separate physical disks.

- Partition tables to match the number of physical devices in a segment.

Objects tab

The following views are used on the Performance Analyst for Sybase I/O Drilldown Objects to succinctly communicate the general overall performance levels of I/O:

Index Scans	Open Object I/O	

Index Scans

The following views are used on the Index Scans view of I/O Drilldown Objects to succinctly communicate the general overall performance levels of I/O:

Ascending Scans	DOL Ascending Scans	DOL Descending Scans
Descending Scans		

Ascending Scans

The Ascending Scans value represents the number of forward scans on allpages-locked tables.

Location	Sybase ASE Performance Analyst statistics > I/O > Objects tab

DOL Ascending Scans

The DOL value represents the number of forward scans on data-only-locked tables.

Location	Sybase ASE Performance Analyst statistics > I/O > Objects tab

DOL Descending Scans

The DOL Descending Scans value represents the number of backward scans on allpages-locked tables.

Location	Sybase ASE Performance Analyst statistics > I/O > Objects tab

Descending Scans

The Descending Scans value represents the number of backward scans on data-only-locked tables.

Location	Sybase ASE Performance Analyst statistics > I/O > Objects tab

Open Object I/O

The Open Object I/O view displays a grid that shows a list of each object on the selected database that is either open, or has been open, and related I/O activity.

>

The table below describes the information available on the Open Object I/O grid:

Column	Description
Database Name	The name of the database.
Object Name	The name of the object.
Index ID	The ID number of the associated index.
Logical Reads	Total number of buffers read.
Physical Reads	Total number of buffers read from disk
APF Reads	The number of asynchronous pre-fetch buffer reads associated with the database.
Pages Read	The number of pages read associated with the database.
Physical Writes	The number of physical writes associated with the database.
Pages Written	The number of pages written to the associated database.
Rows Inserted	The number of rows inserted in the associated database.
Rows Updated	The number of rows updated in the associated database.
Rows Deleted	The number of rows deleted from the associated database.
Used Count	The number of rows used.
Lock Requests	The number of requests for a lock on database objects.
Lock Waits	The number of times a task waited for a lock on a database object.

Location	Sybase ASE Performance Analyst statistics > I/O > Objects tab
Metrics	Objects are only removed from the list if the configured number of "open objects" is exceeded. Then they are aged out to make room for newly opened objects.

Network tab

The Network tab of the I/O Detail view presents the following statistics:

Network I/O Summary	Network I/O by Engine	Network I/O by Session

Network I/O Summary

Here you see a summary of data that includes the following statistics:

Packets Sent: Total Packets Sent reflects the number of times Sybase sent a packet to a client application.

Packets Received: Total Packets Received reflects the number of times Sybase received a packet from a client application.

Bytes Sent: Total Bytes Sent reflects the number of bytes sent to Sybase since the last refresh.

Bytes Received: Total Bytes Received reflects the number of bytes received by Sybase since the last refresh.

Location	Sybase ASE Performance Analyst statistics > I/O > Network tab

>

Network I/O by Engine

The statistics, presented in tabular format, include:

Engine: The name of the engine under scrutiny.

Packets Sent: The number of packets sent by the engine.

Packets Received: The number of packets received by the engine.

Bytes Sent: The number of bytes sent by the engine.

Bytes Received: The number of bytes received by the engine.

Location	Sybase ASE Performance Analyst statistics > I/O > Network tab

Network I/O by Session

This view shows the amount of network I/O committed by registered sessions on the Sybase Server.

Presented in tabular format, you see:

Login: The log in name used by the session.

SPID: The process ID.

KPID: The kernel process ID.

Packets Sent: The number of packets sent by the session.

Packets Received: The number of packets received by the session.

Bytes Sent: The number of bytes sent by the session.

Bytes Received: The number of bytes received by the session.

Network Packet Size: Network packet size used by the current session.

Location	Sybase ASE Performance Analyst statistics > I/O > Network tab
Metrics	This metric can be used to see if a user or batch job is generating more (or less) network traffic than expected. The process' network traffic, measured in conjunction with overall server network traffic, could help distinguish between process level or network-wide loads being the source of contention.

Devices tab

The following statistics are used on the Performance Analyst for Sybase I/O Drilldown Devices to succinctly communicate the general overall performance levels of I/O:

Device I/O Summary	Device I/O Type	Real Time Device I/O

Device I/O Summary

Devices are accessed repeatedly in a dynamic Sybase environment to satisfy user requests for data, to handle write activity that records transactions in a database's transaction log, and to manage other I/O operations. Viewing the I/O activity for each device is a good way to see what the "hot" devices are in a Sybase server with respect to I/O usage. The same information can be used to spot heavy database usage in systems where the device-to-database mappings are one-to-one.

>

The table below describes the information available in the Device I/O Summary view:

Column	Description
Device	The device's name.
APF Reads	The number of asynchronous pre-fetch buffer reads associated with the database.
Physical Reads	The number of physical reads associated with the database.
Physical Writes	The number of physical writes associated with the database.
Total I/O	The combined total of read and write operations.

Location	Sybase ASE Performance Analyst statistics > I/O > Devices tab
Metrics	If devices have one-to-one relationships with user databases, you can quickly tell which databases are experiencing the most demand on a server. If the device-to-physical drive/file system is a one-to-one relationship, you can also spot which physical drives on a server are getting the heaviest workouts. For servers that have many drives, it is normally desirable to spread devices across different physical drives and controllers so contention does not occur at the disk level. In addition, separating databases and their corresponding logs is normally recommended so that each is located on a distinct physical drive. If possible, write-intensive devices, like log devices, are best suited for non-RAID5 scenarios.
Correction	If device hit rates are low, you can add more devices or redistribute objects among different devices. Typically, segments can be used to redistribute objects among different devices or physical hard disks. Common techniques include placing tables and indexes on different segments and partitioning large tables.

Device I/O Type

The Device I/O Type view shows the type of I/O registering against each configured device.

Devices are accessed repeatedly in a dynamic Sybase environment to satisfy user requests for data, to handle write activity that records transactions in a database's transaction log, and to manage other I/O operations. Viewing the I/O activity for each device is a good way to see what the "hot" devices are in a Sybase server with respect to I/O usage. The same information can be used to spot heavy database usage in systems where the device-to-database mappings are one-to-one.

The table below describes the information available from the Device I/O Type detail view:

Column	Description
Device	The name of the device.
I/O Type	The type of I/O.
I/O Time (sec)	The amount of time, in seconds, spent performing I/O operations.
Total I/O	The combined total of read and write operations.

Location	Sybase ASE Performance Analyst statistics > I/O > Devices tab
Metrics	If devices have one-to-one relationships with user databases, you can quickly tell which databases are experiencing the most demand on a server. If the device-to-physical drive/file system is a one-to-one relationship, you can also spot which physical drives on a server are getting the heaviest workouts. For servers that have many drives, it is normally desirable to spread devices across different physical drives and controllers so contention does not occur at the disk level. In addition, separating databases and their corresponding logs is normally recommended so that each is located on a distinct physical drive. If possible, write-intensive devices, like log devices, are best suited for non-RAID5 scenarios.

Correction	If device hit rates are low, you can add more devices or redistribute objects among different devices. Typically, segments can be used to redistribute objects among different devices or physical hard disks. Common techniques include placing tables and indexes on different segments and partitioning large tables. If device loads appear skewed (one device has much more activity than others), you should focus attention on that device. Again, redistributing objects can lessen the device's workload.

Real Time Device I/O

Devices are accessed repeatedly in a dynamic Sybase environment to satisfy user requests for data, to handle write activity that records transactions in a database's transaction log, and to manage other I/O operations. Viewing the current, real-time I/O activity for each device is a good way to see what the "hot" devices are in a Sybase server with respect to I/O usage. The same information can be used to spot heavy database usage in systems where the device-to-database mappings are one-to-one.

The table below describes the information available on the Real Time Device I/O detail view:

Column	Description
Device	The logical name of the device.
Physical Name	The physical name of the device.
APF Reads	The number of asynchronous pre-fetch buffer reads associated with the database.
Physical Reads	Number of buffers read from disk.
Physical Writes	Number of buffers written to disk.
I/O Time (Secs)	The amount of time, in seconds, spent performing I/O operations.
Semaphore Requests	The number of device semaphore requests immediately granted on the device.
Semaphore Waits	The number of tasks forced to wait for synchronization of an I/O request.

Location	Sybase ASE Performance Analyst statistics > I/O > Devices tab
Metrics	The amount of Reads, APF Reads, and Writes can help you spot overworked devices. You should also keep your eye on the semaphore waits and I/O time for each device to help you determine devices that are experiencing I/O related delays.
Correction	If device hit rates are low, you can add more devices or redistribute objects among different devices. Typically, segments can be used to redistribute objects among different devices or physical hard disks. Common techniques include placing tables and indexes on different segments and partitioning large tables. If device loads appear skewed (one device has much more activity than others), you should focus attention on that device. Again, redistributing objects can lessen the device's workload.

Engines Tab

In symmetric multiprocessing (SMP) environments, a DBA can configure the Sybase server to use more than one "engine," which represents a certain amount of CPU power. By default, Sybase configures one engine for use. If you have a server machine with multiple CPUs, you can enable more engines to take advantage of the machine's additional processing ability.

>

The Engines tab of the I/O Detail view displays information with respect to how each engine is handling I/O. The table below describes the information available on the Engines tab of the I/O Detail view:

Column	Description
Engine	The name of the configured engine.
Completed I/Os	The number of I/Os completed during the sample interval.
Outstanding I/Os	The number of I/Os left outstanding during the sample interval.

TIP: To configure the grid to show/hide row numbers, use the Options Editor.

Location	Sybase ASE Performance Analyst statistics > I/O > Devices tab
Metrics	If the number of outstanding I/Os remains high or increases during periods of heavy activity, there may not be enough engines configured for the system.
Correction	If the server machine has multiple CPUs, you can configure more engines for Sybase to use by following this process: *Use the sp_configure procedure to change the current engine configuration. For example, to change the number of engines from one to two, you would run: 'sp_configure "max online engines",2'* *Stop and restart the Sybase server.*

Index Scans Tab

Indexes are accessed frequently in dynamic Sybase server environments. The type of index access often determines the response time an end user experiences. Single row index accesses are the quickest, and complete index scans are the most time consuming (for large indexes at least).

The Index Scans tab of the I/O Detail view presents information with respect to index scans. The table below describes the information available on the Index Scans tab of the I/O Detail view:

Column	Description
Scan Type	The type of index scan.
Count	The number of scans per type for the sample interval.

Location	Sybase ASE Performance Analyst statistics > I/O > Devices tab
Metrics	There are two basic scan types - ascending and descending. For ascending scans, Sybase moves up the index page chain from beginning to end. Sybase follows the page chain in reverse order for descending scans. Descending scans are normally the result of requests made for data by descending column order. Within ascending and descending scans, a data-only lock, or DOL, styled access can also occur.

Space

In addition to a Home page, the Space category of Sybase Performance Analyst includes the following tabbed pages:

Databases tab	Devices tab	

>

Home

The Space performance category view displays the following vital Sybase space statistics:

Space Analysis pane	Transaction Log Analysis	Bottleneck Analysis pane
Device Analysis pane		

Space Analysis pane

A Sybase Server is broken up into many different databases, each of which has its own level of activity. The Space Analysis section displays a summary of activities that have occurred within each database. It also displays the total number of databases defined on the Sybase server.

The table below describes the database information available on the Space Analysis view:

Column	Description
Database Name	The name of the Sybase database.
Total Space (MB)	The total amount of space allocated to the database expressed in MB.
Free Space (MB)	The amount of space currently free expressed in MB.
Used Space (MB)	The amount of space currently in use, expressed in MB.
Percent Free	The available free space on the database, expressed as a percentage.
Log Total Space (MB)	
Log Free Space (MB)	
Log Used Space (MB)	
Log Percent Free	
Log Percent Used	

Transaction Log Analysis

The following statistics are used on the Performance Analyst for Sybase Space page to succinctly communicate the general overall performance levels of space related metrics:

- Transaction Log Analysis

Transaction Log Analysis

The Transaction Log Analysis view indicates if any transaction log for a database in the Sybase server is approaching low levels of available free space. The table below describes the information available on the Transaction Log Analysis grid:

Axis	Description
Percent Space Used	The amount of space currently used in the log, expressed as a percentage.
Databases	The name of the Sybase database.

Location	Sybase ASE Performance Analyst statistics > Space > Home > Space Analysis pane

Metrics	If the transaction log runs out of space, no transactions can take place in the database until you free space in the log.
Correction	If a database log consistently approaches its limit on used space, there are a few actions a DBA can take to prevent a database from freezing. A backup (dump) of the log can be taken: *If log backups are not required for disaster recovery, the truncate log on checkpoint option can be set for the database. Setting this option deletes any space devoted to transactions in the log that have already completed when a database checkpoint occurs.* *You can extend the log onto a new device. Or, if space exists on the device on which the database log currently resides, you can allocate more space for the log on the current device.*

Bottleneck Analysis pane

The following statistics are used on the Performance Analyst for Sybase Space page to succinctly communicate the general overall performance levels of space related metrics:

- Devices low on space

The following statistics, available on this pane, are duplicates of statistics available on the Bottleneck Analysis pane of he Home page:

- Databases Low on Space
- Logs Low on Space

Devices low on space

This metric indicates if any device in the Sybase server is approaching low levels of available free space.

Location	Sybase ASE Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	The environment of the particular Sybase server, as well as your work style dictate the metrics you should use to evaluating a device that is running into trouble with space. Many DBAs create devices that parallel a single, corresponding database in size, and therefore, such devices show 100% utilization. Other DBAs create large devices that are not completely utilized.
Correction	If a device has become too full, a DBA can begin the process of manually relocating databases from it onto other devices. The process of moving logs is easier and can be accomplished via singular commands (sp_logdevice).

Device Analysis pane

The Device Analysis view presents space related details about a selected device located on the Sybase server. The table below describes the information available on the Device Analysis view:

Column	Description
Devices	Total number of devices configured on the server.
Device Name	The name of the device.
Physical Name	The name used by the operating system to identify the device.
Total Space (MB)	Total space allocated on this device in MB.
Used Space (MB)	Total space, in MB, currently in use on the device.
Free Space (MB)	Amount of space that has not been allocated on this device.

>

Column	Description
Percent Free (MB)	The available free space on the database, expressed as a percentage.
Percent Used (MB)	The amount of space used on the database, expressed as a percentage.

Location	Sybase ASE Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	The environment of the particular Sybase server as well as your work style dictate the metrics you should use to evaluate a device that is running into trouble with space. Many DBAs create devices that parallel a single, corresponding database in size, and therefore, such devices show 100% utilization. Other DBAs create large devices that are not completely utilized.
Correction	If a device has become too full, a DBA can begin the process of manually relocating databases from it onto other devices. The process of moving logs is easier and can be accomplished via singular commands (sp_logdevice).

Databases tab

The following views are used on the Performance Analyst for Sybase Space drill-down page to succinctly communicate the general overall performance levels of space related metrics:

Database Object Summary	Database Object Detail	

Database Object Summary

A Sybase server contains many databases, some of which are devoted to system-level activities (the master and tempdb databases, for example) and others that hold user data. The database overview displays details about the space situation for each database in a Sybase server, including the total, used, and free space. The percentage used amount for each database is also shown.

The table below describes the information available on the Database Object Summary detail view:

Column	Description
Database Name	The name of the database.
Total DB Space (KB)	Total amount of space allocated to the database.
Free DB Space (KB)	Total amount of space allocated but free on the database.
Used DB Space (KB)	Total amount of space currently used by the database.
Percent DB Free	The amount of free space expressed as a percentage.
User Table Count	The number of user tables in the database.
Table Reserved (KB)	The amount of total space reserved (used and free) for tables.
Table Data (KB)	The amount of total space used by tables.
Table Unused (KB)	The amount of total unused space reserved for tables.
User Index Count	The number of user indexes.
Index Reserved (KB)	The amount of index space that is reserved.
Index Data (KB)	The amount of data stored in indexes.
Index Unused (KB)	The amount of index space that is unused.
Total Log Space (KB)	The total amount of log space.

>

Column	Description
Free Log Space (KB)	The amount of log space that is allocated but unused.
Used Log Space (KB)	The amount of log space that is currently in use.
Percent Log Free (KB)	The amount of free log space expressed as a percent.

Location	Sybase ASE Performance Analyst statistics > Space > Databases tab
Metrics	If a database's used space percent amount goes above 90%, and the database is dynamic in nature (meaning that users are constantly adding and modifying data), then you should take action to ensure that the database does not run out of available free space.
Correction	If the percent used amount of a database is approaching problematic levels, there are two ways you can rectify the situation: *If the database device that the database currently resides on contains additional free space, you can ALTER the database to consume more available space on that device.* *If the database device that the database currently resides on does not contain additional free space, you can do one of the following:* *Create a new device and issue an ALTER for the database to use that device for space in addition to those currently used.* *Choose another existing device that has free space and ALTER the database to use that device in addition to those currently in use.*

Database Object Detail

Tables and indexes comprise every Sybase database. The Database Object Detail view presents space-related information about tables and indexes. You can choose to view object details by table or index. Or, view objects matching a pattern you specify in the text box.

The table below describes the information available in the Database Object Detail view:

Column	Description
Owner	The owner of the table or index.
Table Name	The name of the table.
Index Name	The name of the index.
Index Type	The type of index.
Key Count	The number of keys.
Fill Factor	The percentage value of the how full each index page can be.
Reserved (KB)	The amount of space reserved for the object, in KB.
Index Size (KB)	The size of the index, in KB.
Unused (KB)	The amount of unused space (free space) that the object contains in KB.
Segment Name	The name of the segment.
Extent Estimate	An estimation as to the number of extents for the object.

Location	Sybase ASE Performance Analyst statistics > Space > Databases tab

Devices tab

The following statistics are used on the Performance Analyst for Sybase Space Drill-down Devices to succinctly communicate the general overall performance levels of space related metrics:

Device Summary	Device Fragment Detail	

Device Summary

The Device Summary view presents space related details about a selected device located on the Sybase server. The table below describes the information available in the Device Summary detail view:

Column	Description
Device Name	The logical name of the device.
Physical Name	The physical name of the device.
Total Space (MB)	Amount of space configured when the device was created.
Used Space (MB)	Amount of space available to any database (via the create or alter database commands).
Free Space (MB)	The amount of free space left in the device.
Percent Free	The amount of free space for the device expressed as a percentage of total space.
Percent Used	The amount of used space for the device expressed as a percentage of total space.

Location	Sybase ASE Performance Analyst statistics > Space > Devices tab
Metrics	If a device's used space percent amount goes above 90%, you should take action to ensure that the device does not run out of available free space.
Correction	If the percent used amount of a device is approaching problem levels, there are two ways you can rectify the situation: *Create a new device and issue an ALTER for the database(s) to use that device for space in addition to those currently used.* *Choose another existing device that has free space and ALTER the database(s) to use that device in addition to those currently in use.*

Device Fragment Detail

Space from a device is allocated among one or more databases on a Sybase server. The Device Fragment Detail section presents all databases and their space fragments that reside on a particular device.

The table below describes the information available on the Device Fragment Detail view:

Column	Description
Database Name	The name of the database.
Type	The type of database fragment (data only, log only, data and log).
Size (MB)	The size of the fragment in MB.

Location	Sybase ASE Performance Analyst statistics > Space > Devices tab

Databases

In addition to a Home page, the Databases category of Sybase Performance Analyst includes the following tabbed pages:

Overview Tab	Objects Tab	Configuration Tab
Engines Tab		

Home

The Database performance category view displays the following vital Sybase database statistics:

Server Analysis pane	Bottleneck Analysis pane	Backup Analysis pane
Engine Analysis pane	Workload Analysis pane	

Server Analysis pane

The following statistic is used on the Performance Analyst for Sybase Databases Page to succinctly communicate the general overall performance levels of database specific metrics:

- Sample Query Response Time

The following statistics, available on this pane, are duplicates of statistics available on the Key Ratio Analysis pane of the Home page:

CPU Busy	I/O Busy	

Sample Query Response Time

Performance Analyst allows you to define a query that is periodically run and timed by the product. The sample query helps you gauge the response time effectiveness of the Sybase server.

Location	Sybase ASE Performance Analyst statistics > Space > Home > Server Analysis pane
Metrics	Because you control the complexity of the sample query, you must determine what is a “valid” response time for the query. Once you have determined that, you can create an alarm in Performance Analyst that will alert you whenever the sample query falls outside of its acceptable baseline.

Bottleneck Analysis pane

The following statistics are used on the Performance Analyst for Sybase Databases page to succinctly communicate the general overall performance levels of database specific metrics:

Databases with Suspect Pages	Engine Configuration Limit	Index Maintenance from Deletes

>

Index Maintenance from Insert/Updates	Operating System Limit	Page Splits
Server Configuration Limit		

The following statistics, available on this pane, are duplicates of statistics available on the Bottleneck Analysis pane of the Home page:

Suspect Databases		

Databases with Suspect Pages

Suspect objects normally indicate an internal problem in a server. They can also indicate that physical damage has occurred to part of a database.

Location	Sybase ASE Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	Suspect objects have no place in a production database. If you identify any suspect objects, you should take immediate action.
Correction	If the suspect object is an index, you could try dropping and recreating it, or use the DBCC REINDEX command. Other damaged objects can complete rebuilding the database. A suspect database can be a difficult thing to recover from quickly. There are times when the cache of the database is suspect and not the database itself. Stopping and starting the Sybase server can verify if this is the case. If the database itself is actually damaged, there could be a true recovery situation. The suspect database can be dropped using the DBCC DBREPAIR DROPDB command. You would then need to recreate the database and perform a recovery operation using the most recent database dump.

Engine Configuration Limit

An engine can exceed its limit for outstanding asynchronous disk I/O requests based on the maximum number of asynchronous I/Os per engine parameter.

Location	Sybase ASE Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	You can change this limit with the maximum asynchronous I/Os per engine configuration parameter.

Index Maintenance from Deletes

Indexes can considerably speed data access operations. The trade-off, however, is that indexes require maintenance whenever indexed columns are added to, modified, or removed in the parent table. The index maintenance from insert/updates metric provides a count of how many index maintenance operations have occurred from insert or update operations.

Location	Sybase ASE Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	Some index maintenance activity is to be expected (unless you are operating in a read-only environment), but excessive index maintenance can quickly become a bottleneck in heavy OLTP environments.
Correction	If you observe a lot of index maintenance activity, you should begin a periodic review of your indexing schemes to ensure that all indexes currently maintained on the system are used and necessary. If you are using Sybase 12.5.03 or later (and have installed the monitoring tables), you can use Performance Analysts object I/O drill-down view to diagnose which tables and indexes are the targets of heavy maintenance activity.

Index Maintenance from Insert/Updates

Indexes can considerably speed data access operations. The trade-off, however, is that indexes require maintenance whenever indexed columns are added to, modified, or removed in the parent table. The index maintenance from insert/updates metric provides a count of how many index maintenance operations have occurred from insert or update operations.

Location	Sybase ASE Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	Some index maintenance activity is to be expected (unless you are operating in a read-only environment), however excessive index maintenance can quickly become a bottleneck in heavy OLTP environments.
Correction	If you observe a lot of index maintenance activity, you should begin a periodic review of your indexing schemes to ensure that all indexes currently maintained on the system are used and necessary. If you are using Sybase 12.5.03 or higher (and have installed the monitoring tables), you can use Performance Analysts object I/O drill down view to diagnose which tables and indexes are the targets of heavy maintenance activity.

Operating System Limit

When a Sybase server begins to experience I/O delays, the result can be a very dissatisfied user community. When such problems occur, you should investigate the Sybase or operating system limits. It could be that I/O operations are being blocked by one or both.

The Operating System Limit statistic shows nonzero numbers if Sybase detects that the limit for asynchronous I/Os has been exceeded.

Location	Sybase ASE Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	You should be concerned if you consistently see numbers above zero.
Correction	In most UNIX operating systems, there is a kernel parameter that limits the number of asynchronous I/Os that can occur at one time. If you continue to see nonzero numbers for this statistic, you should look into raising this limit.

Page Splits

When data is inserted or updated in a table, Sybase may reorganize the data storage in the table's index pages. When an index page becomes full, but a DML operation demands room on that page, Sybase moves half the rows to a new page to accommodate the request. This reorganization is known as a page split.

Location	Sybase ASE Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	Performance for DML actions can be impaired from page split operations. In addition, more index pages can make for longer index scan times.

Server Configuration Limit

When a Sybase server begins to experience I/O delays, the result can be a very dissatisfied user community. When such problems begin to occur, you should investigate the Sybase or operating system limits. It could be that I/O operations are being blocked by one or both limits.

The Server Configuration Limit statistic shows nonzero numbers if Sybase has exceeded its limit for the number of asynchronous disk I/O requests that can be outstanding for the server at one time.

Location	Sybase ASE Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	You should be concerned if you consistently see numbers above zero.

Correction	If you continue to see nonzero numbers for this statistic, you can raise this limit using sp_configure with the max async I/Os per server parameter.

Backup Analysis pane

The following statistics are used on the Performance Analyst for Sybase Databases page to succinctly communicate the general overall performance levels of database specific metrics:

Backup Server Status	Last Transaction Log Dump	

Backup Server Status

The Backup Server Status value indicates whether or not the Backup Server is up and running.

Location	Sybase ASE Performance Analyst statistics > Space > Home > Backup Analysis pane

Last Transaction Log Dump

The Last Transaction Log Dump value represents the last time the transaction was dump for the associated database.

Location	Sybase ASE Performance Analyst statistics > Space > Home > Backup Analysis pane

Engine Analysis pane

The following statistics are used on the Performance Analyst for Sybase Databases page to succinctly communicate the general overall performance levels of database specific metrics:

Engine Busy %	Maximum Engines	

Engine Busy %

The Engine Busy % value represents the percentage of time the Adaptive Server Kernel is busy executing tasks on each Adaptive Server engine (rather than time spent idle).

Location	Sybase ASE Performance Analyst statistics > Space > Home > Engine Analysis pane
Metrics	When Adaptive Server has no tasks to process, it enters a loop that regularly checks for network and completed disk I/Os, and tasks in the run queue. Operating system commands to check CPU activity may show high usage for an Adaptive Server engine because they are measuring the looping activity, while Engine Busy % does not include time spent looping it is considered idle time.
Correction	This category can help you decide whether there are too many or too few Adaptive Server engines. Adaptive Server's high scalability is due to tunable mechanisms that avoid resource contention. If values are consistently very high (more than 90%), it is likely that response time and throughput could benefit from an additional engine.

Maximum Engines

The Maximum Engines value represents the maximum number of engines currently connected.

Location	Sybase ASE Performance Analyst statistics > Space > Home > Engine Analysis pane

>

Workload Analysis pane

The following statistic is used on the Performance Analyst for Sybase Databases page to succinctly communicate the general overall performance levels of database specific metrics:

- Workload Analysis

Workload Analysis

The Sybase Server is broken up into many different databases, each of which has its own levels of activity. The Workload Analysis section displays a summary of resource consumption across all databases.

The table below describes the information available in the Workload Analysis section:

Column	Description
DB Name	The name of the database.
Connections	The number of connections to the database.
Blocked Users	The number of users currently blocked on the associated database.
Total Memory	The total amount of memory currently being used by the database.
Total I/O	The total amount of I/O currently being registered by the database.
Total CPU	The total amount of cumulative CPU for all processes connected to the database.

Location	Sybase ASE Performance Analyst statistics > Space > Home > Engine Analysis pane

Overview Tab

The following statistics are used on the Performance Analyst for Sybase Databases Detail Overview Page to succinctly communicate the general overall performance levels of database specific metrics:

Database Size Analysis	Database Object Analysis	Database User Analysis
Database Summary		

Database Size Analysis

The Database Size Analysis chart displays each database in terms of its total percentage of size to the total size of all databases on the Sybase Server.

Location	Sybase ASE Performance Analyst statistics > Space > Overview Tab

Database Object Analysis

The Database Object Analysis chart displays each database in terms of its terms of its total percentage of objects to the total objects for all databases on the Sybase Server.

Location	Sybase ASE Performance Analyst statistics > Space > Overview Tab

>

Database User Analysis

The Database User Analysis chart displays each database in terms of its total percentage of user accounts to the total user accounts of all databases on the Sybase Server.

Location	Sybase ASE Performance Analyst statistics > Space > Overview Tab

Database Summary

This section displays summarized metadata for all defined databases on the Sybase Server.

NOTE: Right-click any column heading to sort the column(s) in ascending or descending order.

The information available in the Database Summary table is described below:

Column	Description
Database	The name of the database.
Created	Gives the month/day/year and hour/minutes/second on which the database was first generated.
Owner	Lets you see who owns the database.
Database Size (MB)	Displays the size of the database in megabytes.
Log Size (MB)	Displays the log size in megabytes.
Tables	Gives the number of tables in the database.
Indexes	Displays the number of indexes in the table.
Users	Lets you see the number of database users.
Last Trans Log Backup	Let you see the last month/day/year and hour/minutes/second on which the transaction logs were most recently backed-up.
Suspect	Lets you see in YES or NO format whether the database is suspect.
Suspect Pages	Lets you see in YES or NO format whether there are any suspect pages in the database.

Location	Sybase ASE Performance Analyst statistics > Space > Overview Tab
Metrics	The following items merit attention: *Any critical database that shows a last transaction log backup date that is older than the database's required backup needs.* *Any database that shows a status of offline or suspect.*

Objects Tab

he following statistics are used on the Performance Analyst for Sybase Databases Detail Objects Page to succinctly communicate the general overall performance levels of database specific metrics:

Database Object Detail	Suspect Objects	Cache Bindings

>

Database Object Detail

This view gives you detailed information for each database. Select the database you want to scrutinize and choose between reviewing tables or indexes. You can also enter a preference to see objects matching at pattern you specify in the text box. The table below describes the information available in the Database Object Detail table:

Column	Description
Owner	Displays the name of the table or index owner, depending on the view you selected.
Table/Index Name	Displays the name of the table or index, depending on the view you selected.
Segment Name	Lets you see associated table/index segment.
Locking Scheme	**TABLE ONLY:** Lets you see the type of locking scheme the table uses.
Last Statistic Date	**TABLE ONLY:** Lets you see the most recent date and time table statistics were gathered.
Row Count	**TABLE ONLY:** Lets you see the number of rows in the target table.
Reserved (KB)	Lets you see the amount of space reserved for the target object in kilobytes.
Avg Row Size	**TABLE ONLY**: Lets you see the average row size for the target table.
Tree Depth	**INDEX ONLY:** Lets you see the number of layers, or depth, of the index.
Leaf Count	INDEX ONLY: Lets you see the number of leaf-level pages in the index.
Leaf Row Size	**INDEX ONLY:** Lets you see the row size for each index leaf.

Location	Sybase ASE Performance Analyst statistics > Space > Objects Tab

Suspect Objects

Suspect objects normally indicate an internal problem in a server or that physical damage has occurred to part of a database. The table below describes the information available in the Suspect Objects grid:

Column	Description
Owner	Displays the name of the table owner.
Table Name	The affected table's name.
Index Name	The affected index's name.

Location	Sybase ASE Performance Analyst statistics > Space > Objects Tab
Metrics	Suspect objects have no place in a production database. If any suspect objects are identified, you should immediately take action.
Correction	If the suspect object is an index, you could try dropping and re-creating it, or use the DBCC REINDEX command. Other damaged objects can complete rebuilding the database. A suspect database can be a difficult thing to recover from quickly. There are times when the cache of the database is suspect and not the database itself. Stopping and starting the Sybase server can verify if this is the case. If the database itself is actually damaged, there could be a true recovery situation. The suspect database can be dropped using the DBCC DBREPAIR DROPDB command. You would then need to re-create the database and perform a recovery operation using the most recent database dump.

>

Cache Bindings

The Sybase cache areas are devoted to facilitating the transfer of data and information between clients and the Sybase database. Tables and/or indexes can be assigned to various caches in hopes of speeding up access to table or index information. The Cache Bindings Detail section provides information about the current cache bindings of tables and/or indexes. The table below describes the information available in the Cache Bindings chart:

Column	Description
Owner	The object owner's name.
Table Name	The table's name.
Index Name	The index's name.
Cache Binding	The name of the cache associated with the table or index.Databa

Location	Sybase ASE Performance Analyst statistics > Space > Objects Tab
Metrics	The objects best suited for being bound to a cache include frequently accessed objects that are oftentimes small in nature. Large tables that are scanned are normally not suitable for being bound to a cache and are best left to being indexed properly.

Configuration Tab

The following statistics are used on the Performance Analyst for Sybase Databases Detail Configuration Page to succinctly communicate the general overall performance levels of database specific metrics:

Database Configuration	Server Configuration	

Database Configuration

The Database Configuration view allows you to see all the settings currently configured parameters and their settings on the associated database. Select a database from the drop-down list to see the parameter/setting information.

Location	Sybase ASE Performance Analyst statistics > Space > Configuration Tab

Server Configuration

The Server Configuration view allows you to see all the settings of currently configured parameters on the Sybase Server. Select a database from the drop-down list to see the following information in the Server Configuration grid:

Column	Description
Parameter	Lets you see the name of the configured parameter.
Configured Value	Displays the configured value of the parameter.
Run Value	Displays the run value.
Minimum	Displays the minimum configuration value.
Maximum	Displays the maximum configuration value.
Dynamic	Indicates whether or not the parameter is dynamic or not.

Location	Sybase ASE Performance Analyst statistics > Space > Configuration Tab

Engines Tab

The following statistic is used on the Performance Analyst for Sybase Databases Detail Engines Page to succinctly communicate the general overall performance levels of database specific metrics:

- Database I/O

Database I/O

The Database I/O view on the Engines tab of the Databases Detail view allows you to see statistics about database objects. The table below describes the information available to you about Database I/O:

Column	Description
Name	The name of the engine.
Start Date/Time	The date and time when the Sybase engine started.
Stop Date/Time	The date and time when the Sybase engine stopped
OSPID	The operating system process ID.
KPID	The kernel process ID.
Status	The engine's current status and whether it is connected or not.
Context Switches	The current number of context switches registering on the associated engine.
Current Connections	The current number of connections registering on the associated engine.
CPU Time	The current amount of CPU time consumed by the associated engine.
CPU Time (System)	The current amount of system CPU time.
CPU Time (User)	The current amount of user CPU time.

Location	Sybase ASE Performance Analyst statistics > Space > Configuration Tab

Users

In addition to a Home page, the Users category of Sybase Performance Analyst includes the following tabbed pages:

Top Sessions tab	Locks tab	Transactions tab
System Waits tab	Session Waits tab	

Home

The Users performance category view displays the following vital Sybase user statistics:

Key User Analysis pane	Bottleneck Analysis pane	Index Scan Analysis pane

Transaction Analysis - Users	Workload Analysis pane	

Key User Analysis pane

The following statistics are used on the Performance Analyst for Sybase User page to succinctly communicate the general overall performance levels of user specific metrics:

Active User Processes	Extended Procedure Requests	Inactive User Processes
Parallel Queries	Procedure Requests	System Processes

Active User Processes

The Active User Processes statistic is the total number of active and open threads reported on the Server. Active Sessions displays the number of processes actively performing work.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Key User Analysis pane

Extended Procedure Requests

The Extended Procedure Requests value represents the number of times extended procedures were executed in the sample interval.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Key User Analysis pane

Inactive User Processes

The Inactive Sessions statistic represents the total number of threads logged onto the server that are idle at the current time.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Key User Analysis pane
Metrics	A large number of inactive users could indicate user sessions that have mistakenly been left logged on. Since each user thread consumes a portion of memory on the server, to reduce resource usage, you should sever any sessions not needing a connection.
Correction	Drill down into the Session Waits page and check sessions that have many seconds idle or that have been logged on for very long periods of time, as indicated by the logon time column. After verifying that a session is no longer necessary, you can KILL it.

Parallel Queries

The Parallel Queries value represents the total number of queries eligible to be run in parallel.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Key User Analysis pane

Metrics

The optimizer determines the best plan, and decides whether a query should be run serially or in parallel and how many worker processes should be used for parallel queries.

>

Procedure Requests

The Procedure Requests value represents the number of times stored procedures were executed in the sample interval.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Key User Analysis pane

System Processes

The System processes value represents the total number of threads logged onto the server that are Sybase internal processes.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Key User Analysis pane

Bottleneck Analysis pane

The following statistics are used on the Performance Analyst for Sybase User Home Page to succinctly communicate the general overall performance levels of user specific metrics:

Average Hash Chain Length	Average Lock Contention %	Current Locks
Deadlock %	Lock Promotions	Lock Timeouts
Total Priority Changes		

The following statistics, available on this pane, are duplicates of statistics available on the Bottleneck Analysis pane of the Home page:

Deadlocks	Blocked Connections	

Average Hash Chain Length

This value represents the average number of locks per hash bucket during the sample interval.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Correction	If the average number of locks per hash chain is more than four, consider increasing the size of the hash table. Large inserts with bulk copy are an exception to this guideline. Lock hash chain lengths may be longer during large bulk copies.

Average Lock Contention %

The Average Lock Contention % represents the average number of times there was lock contention as a percentage of the total number of lock requests.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Correction	If the lock contention average is high, investigate the locks to see if the application is causing a lock contention or deadlock-related problem.

Current Locks

There are varieties of operations in Sybase that require the use of locks. The Current Locks statistic represents the number of total locks currently active in Sybase.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	You should watch to see if the if the number approaches the Sybase limit for available locks.
Correction	If the number of current locks in a Sybase server approaches the Sybase limit for available locks, you can look into increasing the Number of Locks configuration parameter.

Deadlock %

A deadlock occurs when two processes have a lock on a separate page or object and each wants to acquire a lock on the other process' page or object. Each waits for the other to release the necessary lock. Sybase constantly checks for deadlocks and, when found, chooses the transaction that has accumulated the least amount of CPU time and terminates the transaction. The server then rolls back that transaction and issues a notification of the event. The other process gets to move forward.

The deadlock % rate displays the percentage of times deadlocks occurred vs. the percentage of locks requested and immediately granted.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	Immediately investigate a percentage much above zero to prevent the situation from mushrooming. You can easily drill down and discover the exact process(es) holding locks and deadlocks that are blocking out other user activity.
Correction	Well-designed applications can minimize deadlocks by always acquiring locks in the same order. You should always do updates to multiple tables in the same order. Once Sybase discovers a deadlock, it takes action and remedies the situation. Embarcadero Performance Analyst makes it easier to discover how prevalent deadlock activity is on a system; preventing deadlocks from occurring in the first place is more difficult. Those responsible for writing systems can minimize deadlocks by ensuring that applications acquire locks in the same order. Likewise, you should always do updates and other DML that act on multiple tables in the same order. You can also shrink the amount of time that Sybase waits to check for deadlocks by modifying the deadlock checking period configuration parameter.

Lock Promotions

The Lock Promotions value represents the average number of lock promotion types combined per second and per transaction.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	Lock promotions occur when the following escalations take place: *"Sh-Page to Sh-Table" - Shared page to shared table.* *"Ex-Row to Ex-Table" - Exclusive row to exclusive table.* *"Sh-R to Sh-Table - Shared row to shared table.* *"Sh-Next-Key to Sh-Table" - Shared next-key to shared table.* *"Ex-Page to Ex-Table" - Exclusive page to exclusive table.*
Correction	If lock contention is high and lock promotion is frequent, you should consider changing the lock promotion thresholds for the tables involved. You can configure the lock promotion threshold either server-wide or for individual tables.

Lock Timeouts

The Lock Timeouts value represents the total number of locks timing out during the sample period.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	You can specify the time that a task waits for a lock: *At the server level, with the lock wait period configuration parameter* *For a session or in a stored procedure, with the set lock wait command* *For a lock table command*
Correction	Except for lock table, a task that attempts to acquire a lock and fails to acquire it within the time period returns an error message and the transaction is rolled back. Using lock timeouts can be useful for removing tasks that acquire some locks, and then wait for long periods of time blocking other users. However, since transactions are rolled back, and users may simply resubmit their queries, timing out a transaction means that the work needs to be repeated.

Total Priority Changes

Total Priority Changes represent the total number of priority changes during the sample period.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	It is normal to see some priority switching due to system-related activity. Such priority switching occurs, for example, when: *A task sleeps while waiting on a lock - Adaptive Server temporarily raises the task's priority.* *The housekeeper task sleeps - Adaptive Server raises the priority to medium while the housekeeper sleeps, and changes it back to low when it wakes up.* *A task executes a stored procedure - the task assumes the priority of the stored procedure and resumes its previous priority level after executing the procedure.*
Correction	If you are using logical process management and there are a high number of priority changes compared to steady state values, it may indicate that an application, or a user task related to that application, is changing priorities frequently. Check priority change data for individual applications. You should verify that applications and logins are behaving expected. If you determine that a high-priority change rate is not due to an application or to related tasks, then it is likely due to system activity.

Index Scan Analysis pane

The following statistics are used on the Performance Analyst for Sybase User page to succinctly communicate the general overall performance levels of user specific metrics:

Ascending	Descending	DOL Ascending
DOL Descending	Total Index Scans	

Ascending

The Ascending value represents the number of forward scans on allpages-locked tables.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Index Scan Analysis pane

>

Descending

The Descending statistic represents the number of backward scans on allpages-locked tables.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Index Scan Analysis pane

DOL Ascending

The DOL Ascending value represents the number of forward scans on data-only-locked tables.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Index Scan Analysis pane

DOL Descending

The DOL Descending value represents the number of backward scans on data-only-locked tables.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Index Scan Analysis pane

Total Index Scans

The Total Index Scans statistic represents the total number of index scans (forward and backward).

Location	Sybase ASE Performance Analyst statistics > Users > Home > Index Scan Analysis pane

Transaction Analysis - Users

The following statistics are used on the Performance Analyst for Sybase User page to succinctly communicate the general overall performance levels of user-specific metrics:

Active Transactions	Committed Transactions	

Active Transactions

The Active Transaction statistic represents a count of the number on in-process transactions for SQL Server.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Index Scan Analysis pane

Committed Transactions

Committed Transactions gives the number of transactions committed since the last refresh inside Performance Analyst. This includes transactions that meet explicit, implicit, and ANSI definitions for committed transactions. Note that multi-database transactions are counted.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Index Scan Analysis pane
Metrics	Multi-database transactions generally incur more overhead than single database transactions (for example, a transaction that modifies two databases is counted as two transactions). They usually involve more log activity and two-phase commits between the different databases, as well as cause more ULC flushes. You can improve performance by reducing the number of multi-database transactions.

>

Workload Analysis pane

The following statistics are used on the Performance Analyst for Sybase User page to succinctly communicate the general overall performance levels of user specific metrics:

- Workload Analysis

Workload Analysis

When your phone starts ringing with complaints of performance slowdowns, some of the first questions you should answer are:

- What users are currently connected to the database?
- What resources are they using?
- What are they currently executing?

The Workload Analysis section of the Performance Analyst page provides insight into the leading resource hogs of a server. The table below describes the information available in the Workload Analysis section of the Users page:

Column	Description
Activity	Statistics are presented for the Top CPU Process, Top I/O Process, or Top Memory Process.
SPID	The process ID.
Login	The logon name the session is using.
% Used	The percentage of CPU, I/O, or Memory consumed by the process.

Location	Sybase ASE Performance Analyst statistics > Users > Home > Workload Analysis pane
Metrics	If any one session appears to be using more than 50% of a total resource (CPU, memory, etc.), then you should drill down into the session to find out what they are currently executing.

Top Sessions tab

The following statistics are used on the Performance Analyst for Sybase Users Drill-down Top Sessions to succinctly communicate the general overall performance levels of user specific metrics:

Top I/O Processes	Top Memory Processes	Top CPU Processes
Top Sessions		

Top I/O Processes

The Top I/O Processes statistic identifies the current Sybase process with the highest % of I/O activity.

Location	Sybase ASE Performance Analyst statistics > Users > Top Sessions tab
Correction	If any one session uses more than 50% of a total resource (CPU, memory, etc.), you should drill down into that particular session and investigate the cause.

>

Top Memory Processes

The Top Memory Process statistic identifies the Sybase process that currently is using the highest percentage of memory in the database.

Location	Sybase ASE Performance Analyst statistics > Users > Top Sessions tab

Top CPU Processes

The Top CPU Process statistic identifies the Sybase process that currently has the highest percentage of CPU usage in the database.

Location	Sybase ASE Performance Analyst statistics > Users > Top Sessions tab

Top Sessions

When a system experiences heavy activity, sometimes you will find that all the user connections are contributing somewhat equally to the overall load. More often than not, however, one or two user connections are causing most of the activity. It might be that a large batch load or other typical process is running that is perfectly okay for your system. Or it might be a runaway process or other rogue connection that needs to be tracked down and possibly eliminated.

The table below describes the information available on the Top Sessions detail view:

Column	Description
SPID	The process ID.
User	The name of the session user.
FID	The process ID of the worker process' parent.
Status	The current status of the session.
Database	The database the session is attached to.
Host	The client machine name the session is using.
Program	This identifies the program being run by the session.
Mem Usage	The amount of memory being used by the session.
CPU Time	The cumulative CPU time for the process.
Physical Reads	The number of physical reads produced by the session.
Logical Reads	The number of logical reads produced by the session.
Physical Writes	The number of physical writes generated by the session.

Location	Sybase ASE Performance Analyst statistics > Users > Top Sessions tab

Locks tab

The following statistics are used on the Performance Analyst for Sybase User Drill-down Locks tab to succinctly communicate the general overall performance levels of user specific metrics:

Locks	Locks by Lock Type	Locks by Database

Locks

To modify database information or structures, a user session must obtain a lock on the object to perform its task. In addition to user locks, Sybase itself issues lock requests to carry out its internal duties. The Locks section gives information about the locks currently on the system and also indicates if any blocking situations are occurring. You can elect to show only blocking locks in the Lock detail view.

The table below describes the information available on the Locks detail view:

Column	Description
SPID	The process ID.
User	The name of the current user
Database	The name of the database where the process is running.
Lock Type	The type of lock (database, table, row id, etc.).
Object Name	The name of the object involved in the lock.
Status	The lock's status (waiting or granted).
Lock Page	The number of pages, if any, that are experiencing locks.
Lock Class	Lets you see the lock class.
Host	The client machine name the session is using.
Program	The executable the process is using against the server.
Command	The command currently being issued by the process.
CPU Time	The amount of time the process took to parse, execute, and fetch the data needed to satisfy the process' execution.
I/O	The amount of I/O activity for the target process.
Mem Usage	The amount of memory being used by the target process.
FID	The process ID of the worker process' parent.
Transaction	The current transactions generated by processes on the Sybase Server

Location	Sybase ASE Performance Analyst statistics > Users > Locks tab Sybase ASE Performance Analyst statistics > Users > Locks/Waits
Metrics	A single blocking user has the potential to stop work for nearly all other processes on a small system, and can cause major headaches even on large systems. Blocks are most often caused by user processes holding exclusive locks and not releasing them via a proper COMMIT frequency. Unless a process times out via an application timeout mechanism, or the process has specified a timeout period, a process waiting for a lock will wait indefinitely.
Correction	Once discovered, a blocking lock situation can normally be quickly remedied - the DBA issues a KILL against the offending process, which eliminates the user's stranglehold on the objects the user was accessing. Other user processes then nearly almost always complete in an instant. Discovering the blocked lock situation is made easier by using tools like Performance Analyst, but preventing the blocking lock situation in the first place is where it gets tricky. The DBA can drill down into user detail and view all current blocking locks to see exactly which sessions are holding the currently restrictive locks. The culprit of blocking lock scenarios is usually the application design, or the SQL being used within the application itself. Properly coding an application to reference database objects in an efficient order, and then using the right SQL to get the job done, is an art. The key to avoiding lock contention is to process user transactions in the quickest and most efficient manner possible - something not always easy to do.

Locks by Database

The Locks by Database view represents the total locks occurring on the Sybase Server stratified by which database they are occurring on.

Location	Sybase ASE Performance Analyst statistics > Users > Locks tab

Locks by Lock Type

To modify database information or structures, a user session must obtain a lock on the object to perform its task. In addition to user locks, Sybase itself issues lock requests to carry out its internal duties. The Locks by Lock Type section provides a breakdown of the different locks types being issued across all databases.

Location	Sybase ASE Performance Analyst statistics > Users > Locks tab

Transactions tab

The Transactions view displays current transactions generated by processes on the Sybase Server.

The table below describes the information available on the Transactions detail view:

Column	Description
SPID	The process ID.
User	The name of the user generating the session.
Status	The current status of the session.
Start Time	The time the session started.
Type	The type of session.
State	The current state of the session, for example, whether or not it is active.
Command	The command currently being issued by the session.
Connection	The amount of memory allocated to the session's connections
CPU Time	The amount of time the process took to parse, execute, and fetch the data needed to satisfy the process' execution.
Mem Usage	The amount of memory being used by the session.
Physical I/O	The amount of physical I/O for the session.

Location	Sybase ASE Performance Analyst statistics > Users > Transactions tab
Metrics	Transaction activity is a good measurement of general performance by a user. If the total or committed numbers do not increase over a period of time, either the process is inactive or it may be hung or blocked.
Correction	Look at CPU and I/O statistics to see if any activity for the process is being noted there.

>

System Waits tab

The following statistics are used on the Performance Analyst for Sybase User Drill-down System Waits tab to succinctly communicate the general overall performance levels of user specific metrics:

System Waits	Wait Percent by Total Waits (By Wait Class)	Wait Percent by Time Waited (By Wait Class)

System Waits

When the Sybase Server is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting may mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. The System Waits section identifies the waits that have occurred on the Sybase Server.

The table below describes the information available on the System Waits detail view. You can elect to display the internal system and scheduler events.

Column	Description
Wait Event	The name of the wait event.
Wait Class	The type of wait, often called a wait category or class.
Wait Count	The number of waits that have occurred.
Percent Waits	The percent of waits as compared to the total number of waits.
Wait Time	The amount of time waited in seconds.
Percent Wait Time	The percent of wait time as compared to the total wait time.

Location	Sybase ASE Performance Analyst statistics > Users > System Waits tab

Wait Percent by Total Waits (By Wait Class)

When the Sybase Server is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting might mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. The Wait Percent by Total Waits section graphically depicts the top waits by their percentage to all waits.

Location	Sybase ASE Performance Analyst statistics > Users > System Waits tab

Wait Percent by Time Waited (By Wait Class)

When the Sybase Server is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting might mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. The Wait Percent by Time Waited section graphically depicts the top waits by their percentage of time waited to total time waited.

Location	Sybase ASE Performance Analyst statistics > Users > System Waits tab

Session Waits tab

The following statistics are used on the Performance Analyst for Sybase User Drill-down Locks tab to succinctly communicate the general overall performance levels of user specific metrics:

Session Waits	Wait Percent by Total Waits (By Wait Class)	Wait Percent by Time Waited (By Wait Class)

Session Waits

Session waits is merely a subset of contention that is viewable at the global database level. Often it takes analysis at the session level to pinpoint the exact source of the wait that is occurring globally. Therefore, you need to become accustomed to viewing contention statistics and waits at the user process level.

When monitoring waits with respect to user sessions, there are two areas of interest:

- What HAS the user session been waiting on?
- What IS the user session waiting on?

The table below describes the information available on the Session Waits detail view. You can elect to include background process in the detail display.

Column	Description
SPID	The unique identifier for the process.
User	The name of the process' user.
FID	The process ID of the worker process' parent.
Status	The current status of the process.
Database	The name of the database where the process is running.
Hostname	The client machine name where the session is running.
Program Name	The name of the program being run by the process.
Wait Event	The name of the wait event.
Wait Class	The type of wait, often called a wait category or class.
Waits	The number of waits that have occurred.
Process Waits	The percent of waits as compared to the total number of waits.
Wait Time (s)	The number of seconds the process spent waiting for this wait.
Percent Waits	The percent of wait time as compared to the total wait time.

Location	Sybase ASE Performance Analyst statistics > Users > Session Waits tab
Metrics	To determine the actual wait causes currently experienced by user connections, you should drill down from the global count of users waiting, into the actual system and user wait details.
Correction	If you find a problem, drill down into wait details to determine whether the waits are resource-related.

>

Wait Percent by Total Waits (By Wait Class)

The Wait Percent by Total Waits section graphically depicts the waits that have experienced the highest percentage of occurrences on the system.

Location	Sybase ASE Performance Analyst statistics > Users > Session Waits tab

Wait Percent by Time Waited (By Wait Class)

When the Sybase Server is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting might mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. The Session Wait Percent by Time Waited section graphically depicts the top session waits by their percentage of time waited to total time waited.

Location	Sybase ASE Performance Analyst statistics > Users > Session Waits tab

Session Details

The Sessions Detail view includes the following tabbed pages:

SQL tab	I/O tab	Object Access tab
Locks/Waits		

SQL tab

The SQL tab of the Sessions Detail view offers the following information:

Current SQL	Current Procedures	

Current SQL

This view shows a snippet of the SQL associated with the user in question.

Location	Sybase ASE Performance Analyst statistics > Users > Session Waits tab

Current Procedures

Here, in tabular format, you find the following information about the user's current procedures:

Owner

Object Name

Object Type

Database: The name of the database on which the procedure resides.

Compile Date: The last date on which the object was compiled.

Memory (KB): The number of pages in the procedure cache that is currently allocated to the process. A negative number indicates that pages are being released (freed) from the process.

Location	Sybase ASE Performance Analyst statistics > Users > Session Waits tab

>

I/O tab

The I/O tab of the Session Detail view displays the following statistical information about the current user's activities in a graphical display:

Physical Reads

Logical Reads

Pages Read

Pages Written

Physical Writes

Table Accesses (scans): The number of table accesses (scans) by the user.

Index Accesses: The number of index accesses by the user.

Tempdb Objects: The number of times the Tempdb was accessed by the user.

Object Access tab

The Object Access tab of the Session Detail view displays information about the current objects being accessed. Here you'll see, in tabular format:

Database: The name of the database

Owner ID: The ID number assigned to the process accessing the object.

Object Name

Object Type: The type of object being accessed.

Index ID

Physical Reads: The number of physical disk reads issued against the object.

Physical APF Reads: Number of physical Asynchronous Pre-Fetch buffers read.

Logical Reads: The number logical reads issued against the object.

Locks/Waits

The Lock/Waits tab of the Session Details page offers statistics on:

- Waits

The following statistic, available on this tab, duplicates a statistic on the Locks tab of the Users page:

- Locks

Waits

When the Sybase Server is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting may mean nothing in the overall scheme of things or it can be an indicator that a database bottleneck exists. The Waits section shows processes that are currently waiting on the Sybase Server.

Here you find, in tabular format, details for:

Login: The logon name used by the session.

Database: The name of the database the process is attached to.

>

Wait Event: The name of the wait event.

Wait Class: The type of wait, often called a wait category or class.

Waits: The number of waits that have occurred.

Wait Time (sec): The amount of time waited in seconds.

Location	Sybase ASE Performance Analyst statistics > Users > Session Waits tab

OS

In many scenarios, an optimally tuned database may not perform well because there are constraints imposed by the system where the database is running. These constraints may include processes competing with the database sever for resources (CPU, I/O, or Memory), a slow CPU, insufficient or slow I/O devices, and insufficient memory. The OS Statistics page of Performance Analyst lets you examine operating system metrics for the following platforms:

- AIX
- HP-UX
- Linux
- Solaris
- Unix
- Windows XP and 2000

NOTE: The statistics available on the OS page depend on the operating system platform.

TIP: If you magnify a graph to change back to the original size or minimize the size, close the OS Page and then reopen it again. Performance Analyst displays the graph in its original form and size.

In addition to a Home page, the OS category of Sybase Performance Analyst includes the following tabbed pages:

CPU Tab	I/O Tab	Memory Tab
Network Tab	Processes Tab	Space Tab

Home

The OS home page includes the following sections:

Bottleneck Analysis pane	Disk Analysis pane	Key Resource Usage pane
Memory Analysis pane	Workload Analysis pane	

Key Resource Usage pane

The following ratios are used on the Performance Analyst OS home page to communicate the general overall performance levels of the operating system:

Disk Time	Load Average	Processor Time
Paged Memory Used (Windows)	Swap Memory Used (AIX, HP-UX, Linux, Solaris, Unix)	

TIP: To open the CPU tab, right-click Processor Time. To open the I/O tab, right-click Disk Time. To open the Memory tab, right-click Paged Memory Used.

Disk Time

The Disk Time statistic is the percentage of elapsed time that the selected disk drive/device was busy servicing read or write requests.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	You should avoid consistently seeing values for this statistic greater then 90%.
Correction	Add more disk drives and partition the files among all of the drives.

Load Average

The Load Average statistic represents the system load averages over the last 1, 5, and 15 minutes.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	High load averages usually mean the system is being used heavily and the response time is correspondingly slow.

Processor Time

The Processor Time statistic indicates the percentage of time the processor is working. This counter is a primary indicator of processor activity.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	If your computer seems to be running sluggishly, this statistic could be displaying a high percentage.
Correction	Upgrade to a processor with a larger L2 cache, a faster processor, or install an additional processor.

Paged Memory Used

The Paged Memory Used statistic is the ratio of Commit Memory Bytes to the Commit Limit. Committed memory is where memory space has been reserved in the paging file if it needs to be written to disk. The commit limit is determined by the size of the paging file. As the paging file increases, so does the commit limit.

NOTE: This statistic is available for the Windows platform.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	This value displays the current percentage value only and not an average. If the percentage of paged memory used is above 90%, you may be running out of memory.
Correction	Increase the size of page file.

Swap Memory Used

The Swap Memory Used statistic is the percentage of swap space currently in use.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	If the percentage of swap memory used is above 90%, you may be running out of memory.
Correction	Increase the size of your swap files.

Bottleneck Analysis pane

The following ratios are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Average Disk Queue Length	Interrupts/Sec	Network Output Queue Length/Network Queue Length
Page Faults/Sec	Processor Queue Length	Processor Speed
Processor		

TIP: To open the I/O tab, right-click any Details menu item. To open the CPU tab, right-click the Item Processor Queues Length Details menu item. To open the Network tab, right-click Network Output Queues Length.

NOTE: The statistics available in this section depend on the platform of operating system.

Average Disk Queue Length

The Average Disk Queue Length statistic is the average number of both read and write requests that were queued for the selected disk during the sample interval.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Bottleneck Analysis pane
Metrics	This metric is useful in identifying I/O related bottlenecks. If the disk queue lengths for certain disks are consistently much higher than others, you may need to redistribute the load among available disks. If the disk queues lengths for all disks are consistently large, and you see a high amount of I/O activity, your disks may be inefficient.
Correction	Some things you can do if you have problems with this statistic include: *Redistribute the data on the disk with the large average disk queue to other disks.* *Upgrade to faster disk(s).*

Interrupts/Sec

Interrupts/Sec is the average rate, in incidents per second, at which the processor received and serviced hardware interrupts. This value is an indirect indicator of the activity of devices that generate interrupts, such as the system clock, the mouse, disk drivers, data communication lines, network interface cards, and other peripheral devices. These devices normally interrupt the processor when they have completed a task or require attention. Normal thread execution is suspended. The system clock typically interrupts the processor every 10 milliseconds, creating a background of interrupt activity. This statistic shows the difference between the values observed in the last two samples, divided by the duration of the sample interval.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Bottleneck Analysis pane Sybase ASE Performance Analyst statistics > OS > CPU Tab

Metrics	A high value indicates possible excessive hardware interrupts; justification is dependent on device activity.

Network Output Queue Length/Network Queue Length

The Network Output Queue Length statistic is the number of threads in the processor queue.

NOTE: The name of this statistic depends on the platform of the operating system.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Bottleneck Analysis pane
Metrics	Unlike the disk counters, this counter shows ready threads only, not threads that are running. There is a single queue for processor time even on computers with multiple processors. Therefore, if a computer has multiple processors, you need to divide this value by the number of processors servicing the workload. A sustained processor queue of less than 10 threads per processor is normally acceptable, dependent of the workload.
Correction	A sustained high value in the Processor Queue Length could indicate that a processor bottleneck has developed due to threads of a process requiring more process cycles than are available. If this is the case, you should look at installing a faster (or an additional) processor.

Page Faults/Sec

The Page Faults/Sec statistic is the overall rate faulted pages are handled by the processor. It is measured in numbers of pages faulted per second. A page fault occurs when a process requires code or data that is not in its working set. This counter includes both hard faults and soft faults.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Bottleneck Analysis pane Sybase ASE Performance Analyst statistics > OS > Memory Tab
Metrics	This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
Correction	If the number of page faults remains consistently high, you can check with your Windows System Administrator for further investigation. Often, large numbers of page faults are not a problem so long as they are soft faults. However, hard faults, that require disk access, can cause delays.

Processor Queue Length

The Processor Queue Length statistic is the number of threads in the processor queue.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Bottleneck Analysis pane Sybase ASE Performance Analyst statistics > OS > CPU Tab
Metrics	Unlike the disk counters, this counter shows ready threads only, not threads that are running. There is a single queue for processor time even on computers with multiple processors. Therefore, if a computer has multiple processors, you need to divide this value by the number of processors servicing the workload. A sustained processor queue of less than 10 threads per processor is normally acceptable, dependent of the workload.
Correction	A sustained high value in the Processor Queue Length could indicate that a processor bottleneck has developed due to threads of a process requiring more process cycles than are available. If this is the case, you should look at installing a faster (or an additional) processor.

Processor Speed

The Processor Speed statistic displays the speed of the active processor in MHz. The speed is approximate.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Bottleneck Analysis pane

Processor

The Processor Statistic displays the type of processor currently in use, for example, GenuineIntel.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Bottleneck Analysis pane

Memory Analysis pane

The following metrics are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Available Paged Memory (Windows)	Available Swap Memory	Available Swap Memory(AIX, HP-UX, Linux, Solaris, Unix)
Used Disk Space	Total Physical Memory	

TIP: To open the Memory tab, right-click any Details menu item.

NOTE: The statistics available in this section depend on the platform of operating system.

Available Paged Memory

The Available Paged Memory statistic shows the amount of virtual memory available for the processes.

NOTE: This statistic is available for the Windows platform.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	If the available virtual memory is less than 10% of the total virtual memory, your system may run out of memory.
Correction	Increase the size of page file.

Available Physical Memory

The Available Physical Memory statistic represents the amount of RAM available to all processes.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	This counter displays the last observed value only and not an average. Use this value with the Total physical memory and paging metrics (Memory details page). If the available physical memory is very small compared to this value, and the paging activity is high, your system may be running low on memory.
Correction	Some things you can do if you have problems with this statistic include: *Check the running processes to see if there are any memory leaks.* *Stop any services that are not required.* *Install additional RAM.*

Available Swap Memory

The Available Swap Memory statistic represents the amount of virtual memory available for the processes.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	If the available Available Swap Memory is less than 10% of the total Swap Memory, your system may run out of memory.
Correction	Increase the size of swap files.

Total Paged Memory/Total Swap Memory

The Total Paged Memory statistic shows the maximum amount of virtual memory available to all processes.

NOTE: The name of this statistic depends on the platform of the operating system.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	It is recommended that this be 1½ - 3 times the amount of RAM on the system.

Total Physical Memory

The Total Physical Memory statistic shows the amount of physical memory installed on your computer.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	This is an informational metric and displays the total amount installed on the machine. Use this value with the available physical memory and paging metrics (Memory details page). If the available physical memory is very small compared to this value, and the paging activity is high, your system may be running low on memory.

Disk Analysis pane

The following ratios are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Free Disk Space	Total Disk Space	Used Disk Space

TIP: To open the Space tab, right-click any Details menu item.

NOTE: The statistics available in this section depend on the platform of operating system.

Free Disk Space

The Free Disk Space statistic shows the unallocated space, in megabytes on all logical disk drives.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Disk Analysis pane

Correction	There are many things a DBA can do to ensure that a database does not encounter a space problem due to physical space limitations, here are two: *If a database currently resides on a disk that has little free space, you can add more files to the database. Of course, you should add the new files to other physical hard disks that can accommodate a growing database.* *You should examine hard disks with shrinking disk space to see if you can relocate or delete files to allow more free space.*

Used Disk Space

The Used Disk Space statistic shows the amount of allocated space, in megabytes on all logical disk drives.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Disk Analysis pane
Correction	There are many things a DBA can do to ensure that a database does not encounter a space problem due to physical space limitations, here are two: *If a database currently resides on a disk that has little free space, you can add more files to the database. Of course, you should add the new files to other physical hard disks that can accommodate a growing database.* *You should examine hard disks with shrinking disk space to see if you can relocate or delete files to allow more free space.*

Total Disk Space

Total Disk Space displays the total allocated and unallocated space, in megabytes on all logical disk drives.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Disk Analysis pane
Correction	There are many things a DBA can do to ensure that a database does not encounter a space problem due to physical space limitations, here are two: *If a database currently resides on a disk that has little free space, you can add more files to the database. Of course, you should add the new files to other physical hard disks that can accommodate a growing database.* *You should examine hard disks with shrinking disk space to see if you can relocate or delete files to allow more free space.*

Workload Analysis pane

The following ratios are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Number of Logins	Number of Processes	Top CPU Process
Top I/O Process	Top Memory Process	

TIP: To open the Processes tab, right-click any Details menu item.

NOTE: The statistics available in this section depend on the platform of operating system.

>

Top Memory Process

Top Memory Process shows the current process that is consuming the most amount of memory. The information displayed is dependent on the platform of the operating system. Information displayed includes the name of the process, process ID, amount of memory consumed expressed in KB, amount of CPU expressed as a percentage, the amount of Major Page Faults, and the amount of I/O expressed in KB/sec.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Workload Analysis pane
Metrics	If you are running out of memory on the system, this is a quick way to identify the top memory user. If the displayed process is using a significant portion of the total memory, it could be causing the memory issues.

Top CPU Process

Top CPU Process shows the current process that is consuming the most amount of CPU. The information displayed is dependent on the platform of the operating system. Information displayed includes the name of the process, process ID, amount of memory consumed expressed in KB, amount of CPU expressed as a percentage, the amount of Major Page Faults, and the amount of I/O expressed in KB/sec.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Workload Analysis pane
Metrics	If the amount of CPU time used by this process is close to 100% and the CPU usage is very high, this process may be the bottleneck on the server.
Correction	Investigate the process further to see if it is in an inconsistent state. Also, look at minimum requirements for CPU speed for the process. You may need to upgrade your CPU.

Top I/O Process

The Top I/O Process statistic shows the current process that is consuming the most amount of CPU. The information displayed is dependent on the platform of the operating system. Information displayed includes the name of the process, process ID, amount of memory consumed expressed in KB, amount of CPU expressed as a percentage, the amount of Major Page Faults, and the amount of I/O expressed in KB/sec.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Workload Analysis pane

Number of Logins

This statistic displays the total number of logins on the server.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Workload Analysis pane

Number of Processes

This statistic displays the total number of processes on the server.

Location	Sybase ASE Performance Analyst statistics > OS > Home > Workload Analysis pane

CPU Tab

The CPU tab of the OS Detail includes the following sections:

Context Switches/Sec	CPU Events	CPU Utilization

>

Load Averages		

The following statistics, available on this tab, are duplicates of statistics available on the Bottleneck Analysis pane of the OS Home page:

Interrupts/Sec	Processor Queue Length	

NOTE: The sections available on this tab depend on the platform of operating system.

CPU Utilization

The CPU Utilization section includes the following information:

% Privileged Time	% User Time	

% Privileged Time

The % Privileged Time statistic is the percentage of elapsed time that the process threads spent executing code in privileged mode.

NOTE: For Windows systems, when a Windows system service is called, the service will often run in privileged mode to gain access to system-private data. Such data is protected from access by threads executing in user mode. Calls to the system can be explicit or implicit, such as page faults or interrupts. These kernel commands, are considered privileged to keep the low-level commands executing and prevent a system freeze. Unlike some early operating systems, Windows uses process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. Some work done by Windows on behalf of the application might appear in other subsystem processes in addition to the privileged time in the process.

Location	Sybase ASE Performance Analyst statistics > OS > CPU Tab
Metrics	The ideal range should be 0-40% (less than 40% indicates excessive system activity).
Correction	If your CPU consistently runs at less than 40% you may need to upgrade your system to include a faster processor(s).

% User Time

The % User Time statistic is the percentage of elapsed time the processor spends in the user mode. User mode is a restricted processing mode designed for applications, environment subsystems, and integral subsystems. The alternative, privileged mode, is designed for operating system components and allows direct access to hardware and all memory. The operating system switches application threads to privileged mode to access operating system services. This counter displays the average busy time as a percentage of the sample time.

Location	Sybase ASE Performance Analyst statistics > OS > CPU Tab
Metrics	If the Privileged Time is high in conjunction with Physical Disk Reads, consider upgrading the disk I/O subsystem.

>

CPU Events

The CPU Events section includes the following information:

Context Switches/Sec	Load Averages	System Calls/Sec

NOTE: The statistics available in this section depend on the platform of operating system.

Context Switches/Sec

The Context Switches/Sec section shows the combined rate at which all processors on the computer are switched from one thread to another. Context switches occur when a running thread voluntarily relinquishes the processor, is preempted by a higher priority ready thread, or switches between user-mode and privileged (kernel) mode to use an Executive or subsystem service.

Location	Sybase ASE Performance Analyst statistics > OS > CPU Tab
Metrics	The ideal range should be between 0-10,000. GA number greater then 10,000 may indicate too many threads contending for resources.

System Calls/Sec

System calls are functions which require the services of the operating system. Usually they involve access to some data that users must not be permitted to corrupt or even change.

Location	Sybase ASE Performance Analyst statistics > OS > CPU Tab
Correction	Try to find a good ratio between the amount of time the operating system runs (fulfilling system calls and doing process scheduling), and the amount of time the application runs. Your goal should be running 60% to 75% of the time in application mode, and 25% to 40% of the time in operating system mode. If you find that the system is spending 50% of its time in each mode, then you should investigate to determine what is wrong.

Load Averages

The Load Average statistic represents the system load averages over the last 1, 5, and 15 minutes.

Location	Sybase ASE Performance Analyst statistics > OS > CPU Tab
Metrics	High load averages usually mean that the system is being used heavily and the response time is correspondingly slow.

Processes Tab

The Processes tab of the OS Detail page succinctly communicates the general overall performance levels of processes. The columns available in this table depend on the platform of operating system. The table below describes the information available in the table on this tab:

Column	Description
Process	The name of the process.
User	The user of the process.
ID	The ID Process is the unique identifier of this process. ID Process numbers are reused, so they only identify a process for the lifetime of that process.

Column	Description
CPU	The CPU is the percentage of elapsed time that all of process threads used the processor to execution instructions.
User Mode	The User Mode is the percentage of elapsed time that the process threads spent executing code in user mode.
Memory **WINDOWS ONLY**	Memory is the current size, in bytes, of the virtual address space the process is using. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is finite, and the process can limit its ability to load libraries.
Memory (MB)	Memory is the current size, in bytes, of the virtual address space the process is using. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is finite, and the process can limit its ability to load libraries.
Memory	Memory is the percentage of the memory used of the total memory.
Active Memory	Active Memory is the amount of committed virtual memory, in bytes for this process. Active memory is the physical memory which has space reserved on the disk paging file(s). There can be one or more paging files on each physical drive. This counter displays the last observed value only; it is not an average.
I/O Data	The rate at which the process is reading and writing bytes in I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/Os.
Elapsed Time	The total elapsed time, in seconds, that this process has been running.
Thread Count	The number of threads currently active in this process. An instruction is the basic unit of execution in a processor, and a thread is the object that executes instructions. Every running process has at least one thread.
Handle Count	The total number of handles currently open by this process. This number is equal to the sum of the handles currently open by each thread in this process.
Priority	The current base priority of this process. Threads within a process can raise and lower their own base priority relative to the process' base priority.
Creating Proc ID	The Creating Process ID value is the Process ID of the process that created the process. The creating process may have terminated, so this value may no longer identify a running process.
Page Faults/Sec	Page Faults/Sec is the rate at which page faults by the threads executing in this process are occurring. A page fault occurs when a thread refers to a virtual memory page that is not in its working set in main memory. This may not cause the page to be fetched from disk if it is on the standby list and hence already in main memory, or if it is in use by another process with whom the page is shared.
Page File	Page File is the current number of kilobytes that this process has used in the paging file(s). Paging files are used to store pages of memory used by the process that are not contained in other files. Paging files are shared by all processes, and the lack of space in paging files can prevent other processes from allocating memory.
Private	Private is the current size, in kilobytes, of memory that this process has allocated that cannot be shared with other processes.

I/O Tab

The I/O tab includes the following sections:

Average Disk Queue/Total Disk Queue	Bytes per I/O Operation	Disk I/O Time
Disk Transfers/Sec	I/O Details	Transfer Rate

NOTE: The sections available on this tab depend on the platform of operating system.

>

I/O Details

The table below describes the information available in this section:

Column	Description
Disk	The disk number assignment.
Reading (KB/s)	The amount of bytes read from the device.
Writing (KB/s)	The amount of bytes written to the device.
Disk Read Time	Disk Read Time is the percentage of elapsed time that the selected disk drive was busy servicing read requests.
Disk Write Time	Disk Write Time is the percentage of elapsed time that the selected disk drive was busy servicing write requests.
Disk Time	Disk Time is the percentage of elapsed time that the selected disk was busy servicing requests.
Avg. Read Queue	Avg. Disk Read Queue Length is the average number of read requests that were queued for the selected disk during the sample interval.
Avg. Write Queue	Avg. Disk Write Queue Length is the average number of write requests that were queued for the selected disk during the sample interval.
Disk Reads/Sec	Disk Reads/Sec is the rate of read operations on the disk.
Disk Writes/Sec	Disk Writes/Sec is the rate of write operations on the disk.

NOTE: The columns available in this table depend on the platform of operating system.

Location	Sybase ASE Performance Analyst statistics > OS > I/O Tab

Transfer Rate

The Transfer Rate section succinctly communicates the general overall performance levels of the server's I/O. The table below describes the statistics in this section:

Statistic	Description	Metrics
Reading	The average number of bytes transferred from the disk during read operations.	It is measured in number of KBs per second.
Writing	The average number of bytes transferred to the disk during write operations.	It is measured in number of KBs per second.
Transferred	The amount of time required to transfer the data between the device's cache and the actual device media.	N/A

NOTE: The statistics available in this section depend on the platform of operating system.

Location	Sybase ASE Performance Analyst statistics > OS > I/O Tab

Bytes per I/O Operation

The Bytes per I/O Operation section of the OS Detail succinctly communicates the general overall performance levels of the server's I/O. The table below describes the statistics in this section:

Statistic	Description	Metrics
Reading	The average number of bytes transferred from the disk during read operations.	It is measured in number of bytes per second.
Writing	The average number of bytes transferred to the disk during write operations.	It is measured in number of bytes per second.

Location	Sybase ASE Performance Analyst statistics > OS > I/O Tab

Average Disk Queue/Total Disk Queue

The Average Disk Queue section of the OS Detail page succinctly communicates the general overall performance levels of the server's I/O. The table below describes the statistics in this section:

Statistic	Description	Information
Average Read Queue	Average Read Queue Length is the average number of read requests that were queued for a particular disk.	Sustained loads on any of the disk queue lengths could represent problems within your system.
Average Write Queue	Average Write Queue is the average number of write requests that were queued for a particular disk.	Sustained loads on any of the disk queue lengths could represent problems within your system. Contact your Windows System Administrator for assistance in Correction these problems.
Average Queue	Average time in ms. to service I/O requests which for disks includes seek, rotational latency and data transfer time).	Sustained loads on any of the disk queue lengths could represent problems within your system. Contact your Unix System Administrator for assistance in Correction these problems.

NOTE: The statistics available in this section depend on the platform of operating system.

Location	Sybase ASE Performance Analyst statistics > OS > I/O Tab
Correction	Contact your Windows System Administrator for assistance in Correction these problems.

>

Disk I/O Time

The Disk I/O Time section of the OS Detail page succinctly communicates the general overall performance levels of the disk's I/O. The table below describes the statistics in this section:

Statistic	Description	Metrics
% Disk Read Time	% Disk Read Time is the percentage of time during the sample interval that the disk is busy servicing read requests.	Use this in conjunction with the % Disk Write Time to identify the disks that may be running at capacity and may require upgrades. This metric may also be used to balance the load between under/over utilized disks.
% Disk Write Time	% Disk Write Time is the percentage of elapsed time that the selected disk drive was busy servicing write requests.	Use this in conjunction with the % Disk Read Time to identify the disks that may be running at capacity and may require upgrades. This metric may also be used to balance the load between under/over utilized disks.
% Disk Time	% Disk Time is the percentage of time during the sample interval that the disk is busy servicing requests.	Use this statistic to identify the disks that may be running at capacity and may require upgrades. This metric may also be used to balance the load between under/over utilized disks.

NOTE: The statistics available in this section depend on the platform of operating system.

Location	Sybase ASE Performance Analyst statistics > OS > I/O Tab

Disk Transfers/Sec

The Disk Transfers/Sec section of the OS Detail page succinctly communicates the general overall performance levels of the disk transfers. The table below describes the statistics in this section

Statistic	Description
Reads/Sec	The the rate of read operations on the disk.
Writes/Sec	The rate of write operations on the disk.
Transfers/Sec	The rate of transfer operations per second.

NOTE: The statistics available in this section depend on the platform of operating system.

Location	Sybase ASE Performance Analyst statistics > OS > I/O Tab

Memory Tab

The Memory tab of the OS Detail page includes the following sections:

Buffer Size	Cache Efficiency	Cache Size
Memory Available	Paging Activity	Page Replacements

The following statistic, available on this tab, duplicates a statistic available on the Bottleneck Analysis pane of the OS Home page:

- Page Faults/Sec

NOTE: The sections available on this tab depend on the platform of operating system.

>

Paging Activity

The Paging Activity section includes the following statistics:

Blocks Input/Sec	Blocks Output/Sec	Paged In
Paged Out	Pages Input/Sec	Pages Output/Sec

NOTE: The statistics available in this section depend on the platform of operating system.

Blocks Input/Sec

The Blocks Input/sec statistic is the number of blocks sent to a block device per second.

Location	Sybase ASE Performance Analyst statistics > OS > Memory Tab
Correction	Although it never hurts to have as much physical memory as your system can handle, here are two things you can check within your system to alleviate the memory bottleneck: *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your swap file. Normally, the bigger the initial size of your swap file, the better, in performance terms.*

Blocks Output/Sec

The Blocks Output/sec statistic is the number of blocks received from a block device per second).

Location	Sybase ASE Performance Analyst statistics > OS > Memory Tab
Correction	Although it never hurts to have as much physical memory as your system can handle, here are two things you can check within your system to alleviate the memory bottleneck. *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your swap file. Normally, the bigger the initial size of your swap file, the better, in performance terms.*

Paged In

The number of disk block paged into core from disk which occurs when the server is experiencing a shortage of RAM. Any nonzero value indicates excessive swap activity. In any virtual memory environment, it is not uncommon to see RAM pages moved out to a swap disk. This is a special disk area in UNIX that's reserved for holding memory pages so that the processor is capable of addressing RAM in excess of its full capability. While page-out of operations are a normal part of any server's operations, page-in operations indicate that the real amount of available RAM has been exceeded and that additional RAM pages are required on the server.

Location	Sybase ASE Performance Analyst statistics > OS > Memory Tab
Correction	High values indicate the possible need to increase the amount of RAM on the system.

Paged Out

The number of disk block paged out of core from disk. In any virtual memory environment, it is not uncommon to see RAM pages moved out to a swap disk. This is a special disk area in UNIX that's reserved for holding memory pages so that the processor is capable of addressing RAM in excess of its full capability. While page-out of operations are a normal part of any server's operations, page-in operations indicate that the real amount of available RAM has been exceeded and that additional RAM pages are required on the server.

Pages Input/Sec

The Pages Input/Sec statistic is the number of pages read from disk to resolve hard page faults. Hard page faults occur when a process requires code or data that is not in its working set or elsewhere in physical memory, and must be retrieved from disk.

Location	Sybase ASE Performance Analyst statistics > OS > Memory Tab
Metrics	This value was designed as a primary indicator of the kinds of faults that cause system-wide delays. It includes pages retrieved to satisfy faults in the file system cache (usually requested by applications) and in non-cached mapped memory files. This counter counts numbers of pages, and can be compared to other counts of pages, such as Memory: Page Faults/sec, without conversion. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
Correction	Although it never hurts to have as much physical memory as your system can handle, there are some things you can check within your system to alleviate the memory bottleneck. *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your page file. Normally, the bigger the initial size of your page file, the better, in performance terms.*

Pages Output/Sec

This metric is the number of pages written to disk to free up space in physical memory. Pages are written back to disk only if they are changed in physical memory. A high rate of pages output might indicate a memory shortage.

Location	Sybase ASE Performance Analyst statistics > OS > Memory Tab
Metrics	Windows NT writes more pages back to disk to free up space when low in physical memory. This counter counts numbers of pages, and can be compared to other counts of pages, without conversion. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
Correction	Although it never hurts to have as much physical memory as your system can handle, there are some things you can check within your system to alleviate the memory bottleneck. *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your page file. Normally, the bigger the initial size of your page file, the better, in performance terms.*

Memory Available

The Memory Available section of the OS Detail page succinctly communicates the general overall performance levels of the server's memory. The statistics available in this section depend on the platform of operating system. The table below describes these statistics:

Statistic	Description	Metrics	Correction
Physical Memory	The Physical Memory statistic is the amount of committed virtual memory, in megabytes. Committed memory is where memory space has been reserved in the paging file if it needs to be written to disk.	This counter displays the last observed value only and not an average.	N/A
Physical	The Physical statistic is the total physical memory available.	N/A	N/A
Paged Memory	The Paged Memory statistic is the amount of committed virtual memory, in bytes. Used Memory is the physical memory which has space reserved on the disk paging file(s). There can be one or more paging files on each physical drive.	This counter displays the last observed value only; it is not an average.	N/A
Swap Memory	The Swap Memory statistic is the free, unreserved swap space.	N/A	Increase the size of your swap file or add additional files if you consistently see a shortage of free, unreserved swap space.
Swap	The Swap statistic is the total swap memory being used.	N/A	N/A
Total Physical	The Total Physical statistic is the total physical memory available.	N/A	N/A
Total Swap	The Total Swap statistic is the total swap memory available.	N/A	N/A
Swap-ins	The number of memory pages paged in from swapspace.	N/A	If the system is consistently low on free memory (over a 30 second average), the memory scheduler will start to swap out processes. Increase the size of your swap file or add additional files.
Swap-outs	The number of memory pages paged in from swapspace.	N/A	If the system is consistently low on free memory (over a 30 second average), the memory scheduler will start to swap out processes. Increase the size of your swap file or add additional files.
Buffers	The Buffers statistic represents the amount of buffers used by the kernel in kilobytes.	N/A	N/A

>

Statistic	Description	Metrics	Correction
Page Faults/Sec	The Page Faults/Sec statistic is the overall rate faulted pages are handled by the processor. It is measured in numbers of pages faulted per second. A page fault occurs when a process requires code or data that is not in its working set. This counter includes both hard faults and soft faults.	This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.	If the number of page faults remains consistently high, you can check with your Windows System Administrator for further investigation. Often, large numbers of page faults are not a problem so long as they are soft faults. However, hard faults, that require disk access, can cause delays.

Buffer Size

The Buffer Size statistic represents the amount of cache used by the kernel in kilobytes.

Location	Sybase ASE Performance Analyst statistics > OS > Memory Tab

Cache Size

The Cache Size statistic represents the amount of cache used by the kernel in kilobytes.

Location	Sybase ASE Performance Analyst statistics > OS > Memory Tab

Page Replacements

The following statistics are available in this section:

Memory Freed (Pages/sec)	Clock Algorithm Scans (Pages/sec)	

Memory Freed (Pages/sec)

Freed or destroyed (Kb/s).

Location	Sybase ASE Performance Analyst statistics > OS > Memory Tab

Clock Algorithm Scans (Pages/sec)

The VMM uses a technique known as the clock algorithm to select pages to be replaced. This technique takes advantage of a referenced bit for each page as an indication of what pages have been recently used (referenced). When the page-stealer routine is called, it cycles through the PFT, examining each page's referenced bit.

Location	Sybase ASE Performance Analyst statistics > OS > Memory Tab

Cache Efficiency

The Cache Efficiency section of the Memory tab succinctly communicates the general overall performance levels of the server's memory. The following statistics are available in this section:

Copy Read Hits %	Data Map Hits %	MDL Read Hits %

Pin Read Hits %		

Copy Read Hits %

The Copy Read Hits % statistic is the percentage of cache copy read requests that hit the cache and does not require a disk read to provide access to the page in the cache.

Location	Sybase ASE Performance Analyst statistics > OS > Memory Tab
Metrics	When the page is pinned in the memory, the page's physical address in the file system cache will not be altered. A copy read is a file read operation where a page in the cache is copied to the application's buffer. Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.
Correction	Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate

Data Map Hits %

The Data Map Hits % statistic is the percentage of data maps in the file system cache that could be resolved without having to retrieve a page from the disk.

Location	Sybase ASE Performance Analyst statistics > OS > Memory Tab
Metrics	Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.
Correction	Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate.

MDL Read Hits %

The MDL Read Hits % statistic is the percentage of Memory Descriptor List Read requests to the file system cache that hit the cache and does not require disk access to provide memory access to the pages in the cache.

Location	Sybase ASE Performance Analyst statistics > OS > Memory Tab
Metrics	Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.
Correction	Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate.

Pin Read Hits %

The Pin Read Hits % statistic is the percentage of pin read requests that hit the file system cache and does not require a disk read in order to provide access to the page in the file system cache.

Location	Sybase ASE Performance Analyst statistics > OS > Memory Tab
Metrics	Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.

>

Correction

Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate.

Space Tab

The Space tab of the OS Detail page includes the following sections:

Device Summary	Device Detail	

Disk Space Free

The Disk Space Free metric displays the amount of free disk space in megabytes.

Disk Space Detail

The Disk Space Detail section of the Space tab succinctly communicates the general overall performance levels of the server's disks and space allotment. The table below describes the statistics in this section:

Statistic	Description
Partition	The drive letter of the disk.
Local Filesystem	The name of the file system.
Type	The type of file system.
Total Space	Total size of the disk/device's capacity expressed in MBs.
Used Space	Amount of MBs currently allocated on the particular disk/device.
Free Space	Amount of MBs currently unallocated and free on the particular disk/device.
Capacity	The percentage of space used on the device.
Mounted On	The mount point of the device.

Location	Sybase ASE Performance Analyst statistics > OS > Space Tab

Device Summary

The Device Summary section of the Space tab displays a bar chart for all devices. The Device Summary section succinctly communicates the general overall performance levels of the server's disks and space allotment. The table below describes the statistics in this section:

Statistic	Description
Used	Amount of MBs currently allocated on the particular disk/device.
Free	Amount of MBs currently unallocated and free on the particular disk/device.

Location	Sybase ASE Performance Analyst statistics > OS > Space Tab

>

Device Detail

The Device Detail section of the Space tab succinctly communicates the general overall performance levels of the server's disks and space allotment. The table below describes the statistics in this section:

Statistic	Description
Partition	The drive letter of the disk.
Local Filesystem	The name of the file system.
Type	The type of file system.
Total Space	Total size of the disk/device's capacity expressed in MBs.
Used Space	Amount of MBs currently allocated on the particular disk/device.
Free Space	Amount of MBs currently unallocated and free on the particular disk/device.
Capacity	The percentage of space used on the device.
Mounted On	The mount point of the device.

NOTE: The statistics available in this section depend on the platform of operating system.

Location	Sybase ASE Performance Analyst statistics > OS > Space Tab

Network Tab

The Network tab of the OS Detail page succinctly communicates the general overall performance levels of the server's networking. The Network tab of the OS Detail page includes the following sections:

Network Details	Packet Collisions	Packet Discards
Packet Errors	Transmission Rate	Transmission Rate (Bytes)
Transmission Queue Length		

NOTE: The sections available on this tab depend on the platform of operating system.

Network Details

The statistics available in this section depend on the platform of operating system. The table below describes the information available in this section:

Column	Description
Network Interface	The name of network interface.
INET Address/Address	The IP address assigned to the network interface.
Pkts Sent/Sec	The number of packets sent per second.
Pkts Received/Sec	The number of packets received per second.
Sent (KB/Sec)	The number of bytes sent per second.
Received (KB/Sec)	The number of bytes received per second.
Out Pkts Discarded	The number of outbound packets discarded.
In Pkts Discarded	The number of inbound packets discarded.
Out Pkt Errors	The number of outbound packet errors.

>

Column	Description
In Pkt Errors	The number of inbound packet errors.
Queue Length	The queue length.
Collisions	The number of collisions.
Packets Discarded	The number of packets discarded.

Location	Sybase ASE Performance Analyst statistics > OS > Network Tab

Transmission Rate

The Transmission Rate section of the Network tab succinctly communicates the packet transmission rate. The following statistics are available in this section:

Packets Received/Sec	Packets Sent/Sec	

Packets Sent/Sec

The Packets Sent/Sec statistic is the number of packets sent over each network adapter per second.

Location	Sybase ASE Performance Analyst statistics > OS > Network Tab

Packets Received/Sec

The Packets Received/Sec statistic is the number of packets received over each network adapter per second.

Location	Sybase ASE Performance Analyst statistics > OS > Network Tab

Transmission Rate (Bytes)

The Transmission Rate (Bytes) section of the Network tab succinctly communicates the packet transmission rate. The following statistics are available in this section:

Received (KB/Sec)	Sent (KB/Sec)	

Sent (KB/Sec)

The Sent (KB/Sec) statistic is the rate at which bytes are sent over each network adapter including framing characters.

Location	Sybase ASE Performance Analyst statistics > OS > Network Tab

Received (KB/Sec)

The Received (KB/Sec) statistic is the rate at which bytes are received over each network adapter, including framing characters.

Location	Sybase ASE Performance Analyst statistics > OS > Network Tab

Transmission Queue Length

The number of pending outgoing packets on either the software transmit queue or the hardware transmit queue

Location	Sybase ASE Performance Analyst statistics > OS > Network Tab

Packet Collisions

A collision occurs when two devices on the network sense that the network is idle and end up trying to send data at the same time. Some collisions are normal on a repeated network, but excess collisions can cause serious performance problems.

Location	Sybase ASE Performance Analyst statistics > OS > Network Tab
Correction	Contact your network administrator if there is high rate of persistent Packet Collisions.

Packet Discards

The Packet Discards section of the Network tab succinctly communicates the packet discard rate. The following statistics are available in this section:

Inbound Packets Discarded	Outbound Packets Discarded	Packet Discards

Outbound Packets Discarded

The Outbound Packets Discarded statistic is the number of outbound packets that were discarded even though no errors had been detected to prevent from being transmitted.

Location	Sybase ASE Performance Analyst statistics > OS > Network Tab
Correction	A possible reason for discarding such a packet could be to free up buffer space.

Inbound Packets Discarded

Received Packets Discarded is the number of received packets that were discarded even though no errors had been detected to prevent from being transmitted.

Location	Sybase ASE Performance Analyst statistics > OS > Network Tab
Correction	A possible reason for discarding such a packet could be to free up buffer space.

Packet Discards

The Packet Discards statistic represents the number of network packets discarded.

Location	Sybase ASE Performance Analyst statistics > OS > Network Tab
Correction	Contact your network administrator if there is high rate of persistent packet discards.

Packet Errors

The Packet Errors section of the Network tab succinctly communicates the packet error rate. The following statistics are available in this section:

Inbound Packet Errors	Outbound Packet Errors	

Outbound Packet Errors

The Outbound Packet Errors statistic is the outbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.

Location	Sybase ASE Performance Analyst statistics > OS > Network Tab

Inbound Packet Errors

The Packets Received Errors statistic is the received packets that contained errors preventing them from being deliverable to a higher-layer protocol.

Location	Sybase ASE Performance Analyst statistics > OS > Network Tab

Top SQL Statistics

A lot of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements or stored procedures can drag an otherwise well-configured database down in terms of user response times.

Before you can identify problem SQL in your database, you have to ask the question of what 'bad SQL' is. What criteria do you use when you begin the hunt for problem SQL in your critical systems? Understand that even the seasoned experts disagree on what constitutes efficient and inefficient SQL; so there is no way to sufficiently answer this question to every SQL Server professional's satisfaction. The Top SQL results grid displays the top SQL statements across various performance categories.

The SQL Details tab of the Top SQL view presents information as follows:

Active SQL Results	Selected SQL Text	

Metrics

When you begin to look for inefficient SQL in a database, there are two primary questions you need to answer:

- What has been the worst SQL that’s historically been run in my database?
- What is the worst SQL that’s running right now in my database?

When Correction a slow system, you should be on the lookout for any query that shows an execution count that is significantly larger than any other query on the system. It may be that the query is in an inefficient loop, or other problematic programming construct. Only by bringing the query to the attention of the application developers will you know if the query is being mishandled from a programming standpoint.

Active SQL Results

The Active SQL results section displays the following information in grid format:

SQL Text: A snippet of the full SQL statement. Clicking on the statement will cause the full statement to be presented in the Selected SQL Text grid.

Database: The name of the database.

SPID: The process ID.

CpuTime: The accumulated CPU time for the statement.

Wait Time: The time the process has been waiting to be serviced, in milliseconds. A value of zero indicates there was no waiting.

Memory Usage (KB): The number of pages in the procedure cache that is currently allocated to the process. A negative number indicates that pages are being released (freed) from the process.

Physical Reads: The current cumulative number of physical disk reads issued by the process.

Logical Reads: The current cumulative number logical reads issued by the process.

Pages Modified: The number of pages that were modified by the process.

Packets Received

Packets Sent

Start Time: The date/time when the statement was first executed.

Batch ID

Selected SQL Text

In this section of SQL details, you see the fully expanded SQL statement that you highlighted in the Active SQL Results grid. By clicking Explain SQL you open the Explain Plan with its ISQL window.

Performance Analyst for IBM DB2 for Linux, Unix and Windows Tutorial

Embarcadero Performance Analyst provides key insight into the complex world of database performance. Running inside the Embarcadero DBArtisan console, or standalone in a Web browser, Performance Analyst helps you pinpoint performance problems that threaten a database's availability or response times. The powerful analysis facilities will satisfy advanced experts, yet novices will find Performance Analyst's interface easy to use and the expert help a welcome aid to troubleshooting difficult performance issues.

About This Performance Analyst for IBM DB2 Tutorial

This tutorial is intended to help you get started with Embarcadero Performance Analyst.

After completion of this tutorial, you will have the foundation you need to explore the many features and benefits of Performance Analyst. You will have learned how to accurately monitor your database using the most popular performance methodologies in use today. You will know how to drill down into the many layers of database performance and also be able to locate performance issues in your database in a fast and effective manner. Finally, you'll understand how to fix performance problems using the integration that exists between Performance Analyst and Embarcadero DBArtisan.

This tutorial is divided into 8 sections.

You can use this basic tutorial as a roadmap of product highlights.

After starting this tutorial, you can select **Help** from the menu to find many additional resources that complement and build on many of the activities shown in this brief tutorial.

For more information, see Additional Evaluation Resources.

Session 1: Getting Started with Performance Analyst for DB2

Download and Install

You can obtain the latest version of Performance Analyst by downloading the DBArtisan software from the Embarcadero Web site at http://www.embarcadero.com/downloads/downloaddbartisan.asp. Performance Analyst runs fully contained within DBArtisan so you will need the DBArtisan product to use Performance Analyst (although Performance Analyst can also run in standalone fashion within a Microsoft Internet Explorer Web browser).

Click **Download**, and follow the steps indicated. Save the file on your computer and then double-click to launch the self-extracting file that will guide you through the installation process.

When you first install an evaluation copy of DBArtisan, you can use the tool for 14 days. After that time, a permanent license is needed.

>

Proceed to Overview for Performance Analyst for DB2

Overview for Performance Analyst for DB2

When DBArtisan is connected to an DB2 database, you can open Performance Analyst two ways:

1. On the toolbar, click the **Performance Analyst** icon.
2. On **Utilities** menu, select **Performance Analyst**.

   DBArtisan opens Performance Analyst.

DB2 requires that certain monitor switches be turned on for collecting various performance metrics. When Performance Analyst is first invoked, you will be prompted to turn on the monitor switches. Also, you can increase or decrease the size of the monitor heap depending on the activity of your database.

If you do not want to see this dialog each time you start Performance Analyst, you can check the "Save these settings as default" check box.

Proceed to Product Navigation

Product Navigation

Navigating through Performance Analyst is easily accomplished through two main mechanisms:

1. The main Performance Analyst toolbar contains buttons and drop-down menus that invoke the performance analysis main views and drill-downs within the product as well as frequently used actions like Refresh and Performance Threshold Customizations.
2. Positioning the mouse pointer on any area within Performance Analyst and right-clicking brings up a navigation menu that allows you to easily move between any of the main Performance Analyst home pages and drill-down views. The right mouse menu also contains context-sensitive options like drilling down further into a particular statistic, copying the display to a text file, and more.

Proceed to Session 2: Using Ratio-Based Performance

Session 2: Using Ratio-Based Performance

Ratio-based analysis has been around for many years, and used to be the only technique database administrators used when they were called on to diagnose the cause of a database slowdown. Performance ratios are very good roll-up mechanisms for busy DBAs, making possible the analysis-at-a-glance approach, however ratios must be computed correctly for them to be meaningful. Key ratios in Performance Analyst are calculated using *delta* statistics, which is superior to computing ratios with cumulative values. Delta statistics, which are values derived by using the before and after measures in each statistical refresh done in Performance Analyst, lets you see changes in performance metrics that would otherwise be invisible if cumulative statistics were used. Performance Analyst makes it easy to see the ratios that are computed with delta statistics vs. cumulative. A triangle symbol denotes a delta-based statistic, while square symbols represent a cumulative-based measure.

Ratio-based analysis is available for all performance categories (memory, etc.) where ratios can assist in determining the performance level of a database. For example, if you would like to see memory-specific ratios, on the Performance Analyst toolbar click **Memory**. Performance Analyst displays the Memory home page that showcases all important memory-related ratios.

Proceed to Obtaining Detail Information for a Ratio Using Performance Analyst for DB2

Obtaining Detail Information for a Ratio Using Performance Analyst for DB2

A ratio gives you a global picture for a particular performance measure. However, you often need more information regarding a statistical measure so you can make an informed decision. Performance Analyst has easy-to-use drill down features that make obtaining detailed information for various ratios a breeze.

For example, to view more information for the Buffer Cache Hit Ratio statistic, do the following:

1. On either the main **Home** page or the **I/O** home page, locate the **Buffer Pool Hit Ratio** statistic.
2. Position the pointer over the statistic.

   The pointer changes to a hand symbol.
3. Double-click the left mouse button.

   Performance Analyst opens the Buffer Cache Tab of the memory drill down view.
4. Or right-click the statistic, select **Details** or **Navigate to… > I/O > Buffer Pool**.

Once either of these actions has been taken, you will be taken to a drill-down view that provides more information on the DB2 buffer pools.

Proceed to Session 3: Using Bottleneck/Wait-based Performance Analysis

Session 3: Using Bottleneck/Wait-Based Performance Analysis in Performance Analyst for DB2

When a DB2 database is up and running, every connected process is either busy doing work or waiting to perform work. A process that is waiting may mean nothing in the overall scheme of things, or it can be an indicator that a database bottleneck exists. Bottleneck analysis is a valid method of measuring performance because it helps you track where a database and user sessions have been spending their time. If lock contentions or heavy table-scan activity has been dragging a database's performance down, you can use bottleneck analysis to confirm the actual root cause.

Performance Analyst uses several key areas of DB2 metrics to form its bottleneck analysis. These areas are usually the main sources of performance drag on the databases.

Performance Analyst prominently displays bottleneck analysis on every performance home page. On the main home page, you will see a summary of all key bottlenecks threatening the database. You can also see specific bottlenecks for memory, I/O, space, objects, and users by navigating to each home page and viewing the Bottleneck Analysis section for that home page.

For example, to see bottlenecks that specifically relate to user applications, follow these steps:

1. Click the Users Performance Analyst toolbar button
2. Locate the Bottleneck Analysis section at the lower-left portion of the screen

The Users Bottleneck Analysis section shows the top bottlenecked applications, concurrency issues, and bottleneck activity count.

Proceed to Obtaining Detail on System Bottlenecks

Obtaining Detail on System Bottlenecks

You can easily obtain more information regarding global bottlenecks that are plaguing a database by drilling down into a particular bottleneck metric. For example, if you would like to see more data regarding the current top concurrency bottlenecks, do the following:

1. Open the Performance Analyst main **Home** page.
2. Double-click the **Concurrency** area in the Bottlenecks section to be taken to the Lock Waits drill-down section.

   Performance Analyst opens the System Waits view.
3. Or right-click anywhere in Performance Analyst, point to **Users**, and then click **Lock Waits**.

   Performance Analyst opens the System Waits view.

The Lock Waits view displays critical details regarding all the lock wait events that DB2 has recorded. You can easily see what applications are waiting on locks and the average lock wait times for all applications.

> **NOTE:** If you would like to see an expanded view of the Lock Waits grid shown on the top of the Lock Waits view, click the Expand button located at the top right corner of the section. All section views in Performance Analyst can be maximized/minimized using this control.

Proceed to Viewing Detail for Application Bottlenecks

Viewing Detail for Application Bottlenecks

Not only can you obtain detail information for system bottlenecks, but you can also use Performance Analyst to get drill-down information regarding the locks held by and waited on by applications.

To see more information regarding application locks, do the following:

1. Open the Performance Analyst main **Home** page.
2. Double-click the **Concurrency Bottlenecks** area to be taken to the Lock Waits drill-down view.

   Performance Analyst opens the Lock Waits tab.
3. Or right-click anywhere in Performance Analyst, point to **Users**, and then click **Locks**.

Once on the Locks drill-down view, you can view the details for locks held and waited on by an application details by following these steps:

1. Examine the Application Details grid
2. Select an application that appears to be experiencing lock waits problem

You will see details on the locks held by, and the locks waited on, by the selected application. You can get further details on the application that is holding the locks by navigating to the Attributes drill-down and selecting the appropriate Application ID.

Proceed to Session 4: Using Workload Performance Analysis

Session 4: Using Workload Performance Analysis

Key ratios give you a global perspective on database activity. Bottleneck analysis gives you insight into things that are holding up user activity and throughput. But another technique is necessary if a database professional is to really get a handle on what's occurring inside a badly performing database.

Workload analysis involves the investigation of two critical areas of a database's performance:

- Application resource consumption and activity
- SQL execution

Without looking at these two key performance categories, you can miss a lot of what could be responsible for perceived performance problems. When performance on a database takes a sudden nosedive, it is not uncommon to find one or two sessions that are causing the bulk of the workload. Finding these processes can be easily accomplished by viewing session metadata coupled with resource consumption and statistical execution statistics.

Proceed to Pinpointing Top Resource Consumers

Pinpointing Top Resource Consumers

Performance Analyst identifies top resource consumers in a couple of different ways. On the main home page, Performance Analyst highlights the top resource consumers across lock usage, CPU usage, and I/O usage. The percentage used across all statistical categories are displayed so a DBA can immediately pinpoint a session that is using all or most of a particular resource. The top resource sessions are also displayed on the memory, I/O, and other home pages as well with a listing for the top sessions for that category (for example, the top IO consumers appear on the I/O home page, etc.).

Proceed to Getting Workload Details Using Performance Analyst for DB2

Getting Workload Details Using Performance Analyst for DB2

Performance Analyst offers several details on top resource applications in addition to the summaries provided on each home page. A detailed and flexible top applications view lets you view a summary of all the major resource applications on the database as well as every critical detail regarding each application connected to the database.

Further drill-downs are available that highlights specific details regarding each application's activities and submitted SQL. To see detailed information regarding an application identified in one of Performance Analyst's Workload Analysis sections, do the following:

1. In a **Workload Analysis** display, position the pointer over one of the processes.
2. Double-click the process or right-click, and then select **Details**.

If you want like to see all key workload analysis metadata for sessions, do the following:

1. On the toolbar, click the **Top Sessions** button.
2. Or right-click anywhere in Performance Analyst, point to **Users**, and then click **Top Sessions**.

Proceed to Session 5: Using SQL Performance Analysis

Session 5: Using SQL Performance Analysis

SQL-based performance analysis is really a subset of Workload Analysis. Understanding current and historical SQL execution patterns lets you to have the second set of data points necessary to properly perform workload analysis. You may find that optimizing SQL code produces some of the best performance-enhancing boosts available for a database.

Proceed to Understanding SQL Activity

Understanding SQL Activity

Performance Analyst offers plenty of insight into SQL performance and execution patterns. Each one of the home pages depicts the current levels of SQL activity as determined by a variety of different metrics for that performance category. For example, Performance Analyst might show that rows reads account for 75% of all SQL statement I/O on the database.

Proceed to Getting Details on Top SQL

Getting Details on Top SQL

Getting details on identified SQL statements is easy in Performance Analyst. For example, if you would like to see further information on a SQL statement identified on the Performance Analyst Home page, do the following:

1. In the **SQL Analysis** section, position the pointer over **Hash Values**.
2. Double-click or right-click, and then select **Details**.

   Performance Analyst opens the Top SQL summary page that shows a breakdown of all top SQL statements along with their execution statistics

If you would like to see the full SQL text of a SQL statement, do the following

1. In the **Top SQL Details** section, double-click a statement.

   Performance Analyst opens the SQL Details Tab of the Top SQL view to display the entire SQL statement and the execution metrics. While on this pane, you can also perform an EXPLAIN of the SQL statement and examine its access path information.

   NOTE: To automatically tune a problem SQL statement if you have Embarcadero SQL Tuner installed on your workstation, do the following:

2. On the **SQL Details Tab**, select a problem statement, and then click the **Tune SQL** button.

   SQL Tuner opens and you can begin a tuning session.

Proceed to Session 6: Using Operating System Performance Analysis

Session 6: Using Operating System Performance Analysis

In many scenarios, an optimally tuned database may not perform well because there are constraints imposed by the system where the database is running. These constraints may include processes competing with the database sever for resources (CPU, I/O or Memory), a slow CPU, insufficient I/O devices, and insufficient memory. Performance Analyst allows you to examine the relevant operating system metrics for DB2.

To view the operating system metrics of the machine hosting the DB2 server you are monitoring, do the following:

1. Click the **OS** Performance Analyst toolbar button.

   OR

2. Navigate to the OS Home Page by going to **OS**, and then **Home**.

   NOTE: The first time you invoke the OS feature you will be prompted for an operating system login. Please specify the correct operating system, host name, username and password, and protocol.

The Operating System home page has been designed in a similar fashion as the datasource home page. On this page you can see the pertinent operating system metrics displayed in five quadrants broken down into the following categories:

- Resource Usage
- Memory Analysis
- Disk Analysis
- Workload Analysis
- Bottleneck Analysis

Proceed to Obtaining Details on the Operating System

Obtaining Details on the Operating System

You can easily obtain more information regarding operating system problems that are degrading database performance by drilling down into a particular analysis quadrant. For example if you'd like to see more information regarding the processes running on the host machine, do the following:

1. Navigate to the Performance Analyst **Operating System** home page.
2. Double-click the **Top Memory Process** in the **Workload Analysis** quadrant.

   Performance Analyst opens the Processes detail view.

   OR
3. Right-click anywhere in Performance Analyst and select the **Navigate to**, **OS**, and then **Processes**.

The Processes detail tab shows all the processes on the host machine. From this view you can easily determine which processes are showing high marks in I/O or CPU, for example, by simply double clicking on the column headers and sorting the processes in ascending or descending order.

Proceed to Session 7: Identifying and Getting Help with Performance Problems

Session 7: Identifying and Getting Help with Performance Problems

You do not need to be an advanced expert in performance tuning to spot trouble in a database if you use Performance Analyst. There are a number of mechanisms in the tool that help you quickly locate areas of concern within a database.

Proceed to Using the Alarm Log

Using the Alarm Log

Performance Analyst has a central and flexible alarm log that lists every performance violation found in all monitored databases. Unlike some products that only show you alarms for a particular database, Performance Analyst lists all alarms for each database you are monitoring in one central log so you can see every problem in one place. To view the Performance Analyst Alarm log, do the following:

1. On the toolbar, click the **Alarms** button.

Once in the Alarm log, you can perform a number of actions including:

- Selecting an alarm and viewing the application page that shows details regarding the metric in violation.
- Filtering the types of alarms shown through the Options dialog.
- Clearing all current alarms.
- Copying or saving the Alarm log into a different format.

To select any of these options, do the following:

1 Position the pointer over any area of the **Alarm Log**.

2 In the shortcut menu, select an option.

Proceed to Recognizing Performance Violations

Recognizing Performance Violations

Performance Analyst makes it easy to visually pinpoint performance violations in a database. Any statistic that does not conform to your predefined performance standards is colorized in any view that it appears so you can easily see potential troublespots.

The colors used to highlight a particular performance violation depend on the severity of the problem. Performance Analyst uses several layers of performance standards/threshold violations with each layer being represented by a different color. You can control everything regarding threshold violations, including how thresholds are determined, their color representation in Performance Analyst's display, and more by using the Threshold Template feature of Performance Analyst.

How can Performance Analyst inform you of performance violations in your database if you minimize the product to your desktop and are working on other things? Even if Performance Analyst isn't right in front of you, it can still let you know that a problem has been found in your database through the systray alarm icon. When Performance Analyst detects a problem in one of your databases, one of the things the product does is enable an alarm icon in your workstation's systray area that communicates the presence of alarms in your database. So even if the Performance Analyst interface isn't in sight, the product can let you know that something in your database requires attention.

Proceed to Creating Threshold Templates

Creating Threshold Templates

How does Performance Analyst know what to signal as a problem in a database? It uses predefined performance standards that you define and enforce. All of this is accomplished by using Performance Analyst's Threshold templates. A template is a set of performance metrics and their thresholds that can be reused to enforce performance standards across any number of databases. Embarcadero provides a default performance template that contains nearly all of the most-watched performance statistics so you can begin monitoring with confidence right out of the box. However, you can create your own Threshold templates so you only are alerted to performance violations that are important to you and your particular environment.

To view or create Threshold templates, do the following:

1 On the toolbar, click the **Thresholds** button.

Performance Analyst opens the Threshold maintenance dialog box.

The Threshold Template Properties dialog box lets you select the options to add or edit templates, as well as other maintenance options. For example, you can copy the original Embarcadero Threshold template and modify it to fit your own needs. You can control what statistics are sampled for performance violations as well as the levels of severity you would like reported. Once you are satisfied with your template, you can save it and then apply it to one or more datasources in one action to immediately begin enforcing your performance standards.

Proceed to Getting Help with Performance Problems

Getting Help with Performance Problems

As databases become more complex, so do the problems that plague them. Knowing what to do for a particular performance problem can be hard in some cases, but Performance Analyst helps by providing a complete set of expert help that you can use as a guide when performance problems occur. To obtain expert help for any statistic shown in Performance Analyst, do the following:

1. Position the pointer over a statistic that you need help with.
2. Right-click and select **Help**.

The Expert Help gives you a general description of the statistic, the range of values that are considered optimal, and some suggestions on how to fix a problem situation. Expert help can also be invoked for any of the Ratio, Bottleneck, SQL, or Workload Analysis sections shown on every home page by clicking on the help icon for that section, which is located at the top right corner of the section.

Proceed to Session 8: Fixing Identified Problems

Session 8: Fixing Performance Problems Using Performance Analyst for DB2

Finding performance problems is one thing, but fixing them is another. Many performance monitors point out a particular problem in a database, but then offer no way to actually correct the situation. Such is not the case with Performance Analyst. Because of its tight integration with DBArtisan, certain problems found in Performance Analyst can be visually corrected by DBArtisan.

If a DB2 tablespace is about to run out of free space, Performance Analyst can issue an alarm that a space shortage exists. Once the space problem has been identified, do the following:

1. On the **Space** home page toolbar, click the **Space** button.
2. In the **Space Analyst** grid, right-click a tablespace, and then select **Manage with DBArtisan**.

   You will then be taken to DBArtisan's Tablespace Editor where you can perform corrective actions, such as adding new containers to the tablespace, or resizing the tablespace's existing container(s)

>

IBM DB2 (W/U/L) Performance Analyst statistics

This section includes expert help for all DB2 categories and statistics in the Embarcadero Performance Analyst views and pages.This help is divided into the following sections:

Home	Memory	I/O
Space	Users	Cfg (Database Configuration)
Instance	OS	Top SQL

Home

The Embarcadero Performance Analyst Home view lets you review availability and overall performance of all monitored databases from a single window. Statistics on the Home view are organized into the following categories:

Key Ratio Analysis pane	Bottleneck Analysis pane	SQL Analysis pane
Storage Analysis pane	Workload Analysis pane	

Key Ratio Analysis pane

Database performance analysts typically use ratio-based or wait/bottleneck-based methods for examining the performance levels. Ratio-based analysis involves examining a number of key database ratios that indicate how well a database is running. Performance ratios serve as very good roll-up mechanisms for at-a-glance performance analysis. Many DBAs have large database farms to contend with and cannot spend time checking detailed wait-based analysis outputs for each and every database they oversee. Succinctly presented performance ratios can assist in such situations by giving solid indicators that can be quickly scanned to see if any database needs immediate attention.

While there are certainly many opinions as to what rules to follow, there are some standards that should always be adhered to. To start with, many of the formulas that make up ratio-based analysis must be derived from delta measurements instead of cumulative statistics. Many of the global ratios that a DBA will examine come from the DB2 Snapshot Monitor. The snapshot monitor returns metrics for the Instance and particular databases since the start of monitoring. Some metrics are cumulative (e.g., counters) and others are instantaneous (e.g., gauges).

A final thing to remember about using ratio-based analysis is that, while there are a number of rules of thumb that can be used as starting points to begin the evaluation of database performance, DBAs must determine each SQL Server's individual 'personality' with respect to the key performance ratios. Some hard and fast rules simply do not apply to every database. The danger in using blanket ratio standards is that they can lead the DBA to haphazardly take action, which can at times contribute nothing to the situation, and sometimes even degrade performance.

The following ratios are used on the Performance Analyst Home page to succinctly communicate the general overall performance levels of the monitored database:

Efficiency	Buffer Pool Hit Ratio	Catalog Cache Hit Ratio	Package Cache Hit Ratio
	Shared Workspace Hit Ratio		
Resource Utilization	Database Heap Utilization	Log Space Utilization	Lock List Utilization
	Shared Sort Memory Utilization		

Buffer Pool Hit Ratio

The DB2 database server reads and updates all data from a bufferpool because memory access is much faster than disk access. Data is copied from disk to a bufferpool as needed by the applications using the database. When the server needs to read/write data and the data is already in the bufferpool, no disk access is required. However, if the data is not in the bufferpool, it needs to be read from the disk, which is significantly slower process.

The buffer pool hit ratio indicates the percentage of time that the database server did not need to load a page to service the request for a page. The calculation takes into account all the index and data pages that were requested.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Key Ratio Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane
Metrics	Avoiding disk I/O is the main issue when you try to improve the performance tuning. A high buffer pool hit ratio is desirable because it indicates a lower frequency of synchronous disk I/O. A database where data is accessed uniformly from very large tables will have a poor hit ratio. There is little you can do to improve the performance in such cases.
Correction	The buffer pool hit ratio on the database Home page is the overall hit ratio of the database. First, drill down to the bufferpool level and check the individual buffer pool hit ratios to identify the elements that may require tuning. Increasing the bufferpool size generally improves the buffer pool hit ratio. You can use one of the following methods to increase the bufferpool size depending on the layout of the data: *If the tablespaces using the bufferpool have tables and indexes on them, increase the bufferpool size incrementally until the index page hit ratio stops increasing for that bufferpool. You will not usually get any benefit increasing the bufferpool size after you reach that point.* *If the tablespaces associated with the bufferpool have only indexes on them, increase the bufferpool size until the index page hit ratio stops increasing.* *If the tablespaces associated with the bufferpool have only data on them, increase the bufferpool size until the data page hit ratio stops increasing.* One general design guideline is to try and use different tablespaces for indexes and tables and associate different bufferpools with them. This generally improves performance and makes bufferpool tuning much easier.

Catalog Cache Hit Ratio

Catalog cache is used to cache the following types of catalog information:

- Metadata for tables, views, and aliases.
- Database authorization information that is used to check authorization when performing operations like CONNECT, LOAD, CREATE, BIND, and so on.
- Execute privileges for user-defined functions and stored procedures.

When a database operation accesses catalog information, it inserts this information into the catalog cache so operations accessing the same information can read it from the cache and avoid disk reads. The catalog cache hit ratio indicates how well the catalog cache avoids accesses to the catalog.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Key Ratio Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory >Caches Tab
Metrics	A high catalog cache hit ratio (> 80%) indicates that the catalog cache is working well. A smaller ratio can indicate that this parameter needs tuning. You may see a smaller ratio immediately following the first connection to the database and execution of DDL/DCL statements since these require heavy catalog accesses.

Correction	If the catalog cache hit ratio is consistently small, increase the database configuration parameter *catalogcache_sz*. When you increase this value, pause to consider whether it would be more effective to allocate the extra memory you are reserving to another purpose such as package cache or bufferpools. In a partitioned database environment, make the *catalogcache_sz* larger on the catalog partition than on other partitions because the catalog information required for all partitions will go through this partition. When tuning this parameter, it is advisable to monitor changes to the database catalog via DDL statements. During the execution of DDL statements, there may be a large drop in the catalog cache hit ratio due to invalidation of the cached data.

Package Cache Hit Ratio

Embarcadero DBArtisan's Database Manager caches the sections for frequently used dynamic and static SQL statements in the package cache. The package cache hit ratio indicates how well the package cache is avoiding catalog accesses to packages and recompilations.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Key Ratio Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Caches Tab
Metrics	A high package cache hit ratio (> 80%) indicates that the package cache is working well. In the case of static SQL statements, package caching allows the Database Manager to reduce the internal overhead by eliminating the need to access system catalogs when reloading a package. For dynamic SQL, the benefit of package caching is even greater since it a cache hit eliminates the need for recompilation. The package cache hit ratio metric is particularly important for transaction-processing applications since a typical workload involves repeatedly executing the same SQL statements.
Correction	Executing DDL statements can invalidate sections of SQL statements in the cache, causing the hit ratio to decrease dramatically. Before attempting to tune this parameter, you should check the DDL activity to see if that is causing a skew. No amount of tuning will improve the package cache performance if the DDL activity is causing package invalidations in the cache. If the DDL activity is minimal and package cache hit ratio is consistently small, consider increasing the package cache size (pckcachesz) in the database configuration. When increasing the value of this parameter, consider whether it might be more effective to allocate the extra memory being reserved to another purpose such as catalog cache or bufferpools.

Shared Workspace Hit Ratio

When sections are required by an application for executing dynamic or static SQL statements, they are placed in the shared workspace. The shared workspace exists at the application level and is shared among applications using the database. The hit ratio is a percentage indicating how well the shared SQL workspace is helping to avoid initialization of sections for SQL statements that are about to be executed. A high ratio indicates the shared workspace is successful in avoiding this action.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Key Ratio Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory >Caches Tab
Metrics	A shared workspace is shared by many applications. If applications have similar database access patterns, they can benefit greatly if they find required sections in the shared workspace. If an application finds a section in the shared workspace (e.g., a hit), that application avoids setup and initialization cost for that section. A high hit ratio is desirable for this metric. Typically, you should expect to see a high ratio (>80%) in transaction processing environments.
Correction	Shared workspace memory is allocated from the application control heap (app_ctl_heap_sz database configuration parameter) and increasing this may improve the hit ratio.

Database Heap Utilization

There is one database heap per database, and the database manager uses it on behalf of all applications connected to the database. The database heap utilization is the percentage of database heap that is currently being used.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Key Ratio Analysis pane
Metrics	The database heap contains control block information for tables, indexes, table spaces, and bufferpools. It also contains space for event monitor buffers, the log buffer, and temporary memory used by utilities. Tthe heap utilization can increase or decrease whenever any of the constituent elements change. If the utilization goes above 85% several times or stays above 85% for a sustained period, you may need to increase the maximum database heap size.
Correction	The *dbheap* database configuration parameter determines the maximum amount of heap memory that can be allocated. Examine the individual memory pools and how they are using the memory before increasing this parameter. An excessively large use of one of the memory pools may be the cause of this problem.

NOTE: On the main toolbar, click **Memory** and then **Memory Pools** to see the Memory Pools Details.

Log Space Utilization

Log space utilization is the percentage of total available log space on the database that is currently being used.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Key Ratio Analysis pane
Metrics	Very high log space utilization (>75%) can indicate that the database might run out of log space soon. Usually, consistently high log space utilization has one of the following causes: *Large transaction size. Applications are performing large numbers of updates on the database but are not committing the transactions frequently.* *The log space configured for the database is insufficient for the level of database activity* Very low log space utilization (< 5) %may indicate an excessively large log file and that you may be able to free some of the space occupied by log files on disk.
Correction	If the log utilization is mostly low but increases to near 100% for short periods, then you can increase the number of secondary log files (*logsecond* database configuration parameter). This allocates and deallocates additional log files as needed. If however, the log size is consistently high and changes frequently, you can increase the primary log size and/or the log file size (*logfilsiz* and logprimary database configuration parameters). Also, check the SQL activity on the SQL Activity tab of the Users Detail section to see if any applications are performing infrequent commits. If this is the case, evaluate whether such applications can be changed to have shorter transactions and hence less need for log space.

Lock List Utilization

Lock list utilization is the percentage of total database memory allocated for locks that is currently being used.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Key Ratio Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Lock Memory Tab
Metrics	There is only one lock list for each database and it contains the locks held by all applications connected to the database. Once the lock list is full, the database manager starts escalating row locks to table locks to free up space. This escalation may result in serious performance degradation because of reduced concurrency. Additionally, the number of deadlocks and transaction rollbacks may go up. If this metric reaches the 75% mark, you should consider bringing this percentage down with tuning.

Correction	Depending on the database's activity level, you may be able to reduce the lock utilization by following these recommendations: **Increase size of lock list:** If there is not enough lock list space available, lock escalations will occur, increasing contention and reducing concurrency. Update the *locklist* database configuration parameter to increase this value. **Tune applications to reduce locking:** On the Locks tab of the Users detail section, identify the applications that are holding many locks and then consider the following steps for controlling the size of the lock list: *Make the transactions shorter by increasing the COMMIT frequency. This ensures that locks are released frequently, thus freeing up lock space.* *Before you update many rows in a single table, lock the entire table (using the SQL LOCK TABLE statement). This prevents many row-level locks from being obtained (although this decreases concurrency)* *To control how locking is done for a specific table, use the LOCKSIZE parameter of the ALTER TABLE.* *To decrease the number of share locks held, use the Cursor Stability isolation level when possible. If the applications' integrity requirements are not compromised, use Uncommitted Read instead of Cursor Stability to further decrease the amount of locking.* **Decrease percentage of lock list:** If a small number of applications are consuming most of the lock space, decrease the percentage of lock list for each application. You can throttle back those applications by decreasing the maxlocks database configuration parameter. This reduces the amount of lock list memory available to each application thereby allowing for better distribution of lock list memory. NOTE: Decreasing the percentage of lock list should be a last resort, used only if you cannot decrease utilization with the other recommendations. It can cause a large number of lock escalations.

Shared Sort Memory Utilization

The amount of memory allocated for each sort may be some or all of the available sort heap size. Sort heap size is the amount of memory available for each sort as defined in the database configuration parameter *sortheap*. Shared sort memory utilization gives the percentage of the sort heap being used. It is possible for a single application to have concurrent sorts active. For example, in some cases a SELECT statement with a subquery can cause concurrent sorts.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Key Ratio Analysis pane
Metrics	Memory estimates do not usually include sort heap space. If excessive sorting is occurring, the extra memory used for the sort heap should be added to the base memory requirements for running the database manager. Generally, the larger the sort heap, the more efficient the sort. Typically the shared sort utilization should be less than or equal to 70%. You should consider tuning the database if you see a utilization value greater than this.
Correction	To bring the sort heap utilization to an acceptable level, use the following guidelines: *Examine the queries you are running on the database to see if you can add indexes to columns frequently accessed in the WHERE clause. This minimizes the use of the sort heap.* *If you are using dynamic bitmaps or hash join buffers in your queries, or your queries frequently require large sorts, increase the sortheap size.* *If you adjust the sortheap size, also look at the sheapthres database manager configuration parameter to see if it too needs to be adjusted* *If you are adjusting the sortheap size, you may also benefit from rebinding your packages since the optimizer takes this parameter into account when determining the access paths.*

Bottleneck Analysis pane

When a DB2 database is active, applications accessing the database can either successfully access the database resources, or wait for resources to become available. Waiting applications can indicate normal database activity, or can indicate that a database bottleneck exists. A database administrator can use wait-based or bottleneck performance analysis to determine if perceived bottlenecks in a database are contributing to a performance problem.

Performance Analyst Bottleneck analysis section helps a DBA isolate potential problem areas where there are waits, resource overloads or potential failures. If concurrency issues or heavy table scan activity is impacting database performance, you can use bottleneck analysis to isolate the root cause. Once one or more bottlenecks have been isolated as possible areas of interest, you can drill down and examine details on why the applications or resources are causing the problem. The Performance Analyst home page identifies the top system and application bottlenecks that may be currently a cause of lessened performance.

The following bottleneck indicators are used on the Performance Analyst Home page to succinctly communicate the general overall performance level of the monitored database:

Concurrency	Apps Waiting on Locks	Avg. Lock Wait Time	Deadlocks Detected
	Lock Escalations	Lock Timeouts	
I/O Activity	Log Pages I/O	Synchronous I/O	
Object Bottlenecks	Accesses to Overflowed Records	Page Reorganizations	
Sorts and Joins	Hash Join Overflows	Sort Overflows	
Storage Bottlenecks	DMS Tablespaces Low on Space	Inaccessible Containers	Node with Smallest Log
	Tablespaces in Abnormal State		

Apps Waiting on Locks

This metric gives the percentage of all currently connected applications that are waiting for locks.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Bottleneck Analysis pane
Metrics	If this number is high, you should investigate whether the applications are having concurrency problems.
Correction	Compare this metric with the lock escalations metric to identify if the lock list memory is too small. Go to the Locks Waiting Details section of the Locks tab of the Users detail view and examine the lock activity at application level to identify the applications that are holding a large number of row locks and table-level locks. You may be able to tune the applications with high number of locks.

Avg. Lock Wait Time

The average lock wait time metric gives the average amount of time applications waited for a lock in the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Bottleneck Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Users > Lock Waits tab
Metrics	This metric gives an overview of the concurrency performance of the database. If this value is too high, it may mean that either too many applications are vying for the same resources or that applications are executing long transactions that make other applications wait longer. You should examine this metric over a period of time to determine the normal operational state of the database.

Correction	As a first step, examine the lock activity at the application level to identify if any particular application is causing excessive lock contentions. If you see a high value, consider tuning the application to provide better concurrency. If lock timeouts are excessive, and average lock wait times are very short, you can increase the *locktimeout* database configuration parameter to make the applications wait longer before timing out. NOTE: You can find the lock activity in the Applications Waiting on Locks section of the Lock Waits tab of the User detail page.

Deadlocks Detected

Deadlocks detected is the total number of deadlocks that have occurred since this instance of Performance Analyst started monitoring the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Bottleneck Analysis pane
Metrics	If a large number of deadlocks are detected, it can indicate that applications are experiencing lock contention problems. Deadlocks are usually caused by one of the following situations: *Lock escalations on the database* *Catalog tables locked for Repeatable Read* *Applications are using inappropriate isolation levels at bind time* *Applications are obtaining the same locks in a different order* *Applications are locking tables explicitly where row level locks are sufficient.*
Correction	You may be able to modify the applications causing lock contentions for better concurrency. To identify the applications that may be causing contentions, go to the Lock Waits tab of the Users detail page and review the Application Details section.

Lock Escalations

The lock escalations metric indicates the number of times that locks have been escalated from row locks to table locks, since this instance of Performance Analyst started monitoring the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Bottleneck Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Users > Bottleneck Analysis pane
Metrics	A lock is escalated when the total number of locks held by an application reaches the maximum amount it is allowed in the lock list memory. There are several possible causes of lock escalations: *The database lock list size is too small for the concurrent workload* *The maximum percentage of lock list memory allowed for each application is too small* *One or more applications are obtaining an excessive number of locks* Monitor the lock escalations over a period of time to determine what levels are acceptable in your environment. If the escalations are excessive, or are accompanied by deadlocks or long lock waits, consider tuning the database.
Correction	Examine the *locklist* database configuration parameter together with the Lock List Utilization metric to see if the *locklist* size needs to be increased. If the lock list utilization is within acceptable range, try raising the *maxlocks* database configuration parameter to increase the percentage of lock list space available to each application. You should also examine the detailed lock usage for each application on the Locks tab of the Users detail page to see if applications are holding a high number of locks. In this case, tune your application for better concurrency.

Lock Timeouts

This metric identifies the number of times that a request to lock an object timed out without being granted.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Bottleneck Analysis pane
Metrics	If the number of lock timeouts becomes excessive when compared to the acceptable range for your database, it can indicate that an application is holding locks for long durations. It can also indicate that the amount of time an application waits for a lock before timing out is too short. If you have too few lock timeouts and the average lock wait time is too high, it can indicate that the lock timeout configuration parameter is set to an excessively high value.
Correction	First you should examine the lock activity at the application level to identify any particular application that is causing excessive lock contentions. If so, you can tune the application to provide better concurrency. If lock timeouts are excessive, and average lock wait times are very short, you can increase the *locktimeout* database configuration parameter to make the applications wait longer before timing out.

Log Pages I/O

The log pages I/O statistic is the total number of pages that are read and written to the disk by the database logger.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Bottleneck Analysis pane
Metrics	You can use this element with an operating system monitor to get a bird's eye view of the amount of I/O on a device that is attributable to database activity.

Synchronous I/O

This metric presents the number of I/O operations that had to be performed synchronously by the database agents.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Bottleneck Analysis pane
Metrics	DB2 read and write operations are always performed through bufferpools. If the page being read or written is not present in the bufferpool, it is first read into the bufferpool and then the I/O operation is performed on it. DB2 pre-fetchers fetch the data page into the bufferpool asynchronously in anticipation of a read/write operation. If the page is not available at the time of I/O, the database manager agent has to perform synchronous I/O. Synchronous I/O can decrease responsiveness of a database, and if the database tablespaces are spread across multiple physical devices you may be able to reduce the number of synchronous operations by tuning the database configuration.
Correction	You can set the *num_ioserver* database configuration parameter to 1-2 more than the number of physical storage devices used by the database to ensure that parallel I/O is occurring. To minimize synchronous I/O, you may also need to layout your tablespaces such that the containers are defined on different physical devices.

Accesses to Overflowed Records

This metric is the number of accesses (reads and writes) to overflowed rows in the database tables.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Bottleneck Analysis pane
Metrics	A row overflows if it is updated and no longer fits in the data page where it was originally written. This usually happens as a result of an update of a VARCHAR or an ALTER TABLE statement. Overflowed rows indicate that data fragmentation has occurred. If this number is high, you can try to improve performance by reorganizing the tables using the Space Analyst.

Page Reorganizations

The page reorganization statistic is the number of page reorganizations executed in the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Bottleneck Analysis pane
Metrics	Too many page reorganizations can result in less than optimal insert performance. You can use the Space Analyst to reorganize tables and eliminate fragmentation. You can also use the APPEND parameter for the ALTER TABLE statement to indicate that all inserts are appended at the end of a table and so avoid page reorganizations. In situations where updates to rows cause the row length to increase, the page may have enough space to accommodate the new row, but a page reorganization can be required to defragment that space. Or if the page does not have enough space for the new larger row, an overflow record is created that causes overflow_accesses during reads. You can avoid both situations by using fixed length columns instead of varying length columns.

Hash Join Overflows

Hash join is an option for the DB2 optimizer. A hash join compares the hash codes before comparing the predicates for tables involved in a join. This reduces the number of comparisons. The hash join overflows metric gives the percentage of all hash join operations that ran out of sort heap and may have required disk space for temporary storage since the current instance of Performance Analyst started monitoring the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Bottleneck Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	An overflowing hash join incurs a large overhead because of the amount of disk I/O required to complete the operation. If this value crosses the 30% mark, the DBA should take action.
Correction	Increase the *sortheap* database configuration parameter to reduce the number of overflows.

Sort Overflows

The sort overflows statistic is the total number of sorts that ran out of sort heap and that may have required disk space for temporary storage.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Bottleneck Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Users > Bottleneck Analysis pane
Metrics	When a sort overflows, additional overhead is incurred. The sort requires a merge phase and can potentially require more I/O if data needs to be written to disk.
Correction	Sort overflows can be reduced by increasing the *sortheap* database configuration parameter.

DMS Tablespaces Low on Space

The DMS Tablespaces Low on Space metric gives the number of DMS tablespaces in the database where more than 90% of the total available space is in use.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Bottleneck Analysis pane
Metrics	This metric is an indicator that some of the tablespaces in your database may be running out of space. You should drill down to the Space metrics to identify the tablespaces that may be at the critical utilization level. In tablespaces where there are little on no inserts after the initial data is loaded, little or no free space may have been left on purpose. In such cases, it is normal to see a high utilization.

Correction	If this metric goes beyond the normal operational value for your database, consider creating additional space on the DMS tablespaces that are running out of space. You can do this by performing one or more one of the following tasks: *Clean up unused space by deleting the tables/indexes that are no longer required.* *Resize or extend the existing containers in the tablespaces.* *Add new containers to the tablespaces.*

Inaccessible Containers

The inaccessible containers metric identifies the number of containers on all tablespaces that are currently inaccessible.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Bottleneck Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	Inaccessible containers represent a serious problem on the database. The database is unable to access/write any data on an inaccessible container. Usually this is caused by either media errors or modifications to the container files/directories from outside DB2.
Correction	Drill down to the Space metrics to examine the inaccessible containers and the tablespaces on which they are contained. You may need to perform a redirected restore on the tablespace from a backup to correct this.

Node with Smallest Log

The node with smallest log metric is only returned for global snapshots and indicates the node with the least amount of available log space.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Bottleneck Analysis pane
Metrics	Use this metric in conjunction with the log space utilization metrics to check if the log space or applications on the specified node need to be adjusted.
Correction	If the log utilization is mostly low but increases to near 100% for short periods, then you can increase the number of secondary log files (*logsecond* database configuration parameter). This allocates and deallocates additional log files as needed. If, however, the log size is consistently high and changes frequently, you can increase the primary log size and/or the log file size (*logfilsiz* and *logprimary* database configuration parameters). Also, check the SQL activity, under Users Drilldown > SQL activity to see if any applications are performing infrequent commits. If possible, such applications can be changed to have shorter transactions and less need for log space.

Tablespaces in Abnormal State

This metric is a count of the number of tablespaces in your database that are not in a 'Normal State'.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Bottleneck Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Space > Home > Bottleneck Analysis pane

Metrics	Tablespaces in not normal state are not necessarily a cause for alarm. They may have been transitioned to that state intentionally by the DBA for maintenance related tasks. If this metric shows a value that is higher than you anticipated, you may need to drilldown to the Space metrics to see which tablespaces are not normal. A tablespace can be in one or more of the following states when it is not 'Normal': Quiesced: SHARE, UPDATE, EXCLUSIVE Pending: Load, delete, backup, roll forward, restore, disable, drop In progress: Roll forward, reorg, backup, TBS deletion, TBS creation Storage must be defined Restore in progress Offline and not accessible Storage may be defined Storage definition is in 'final' state Storage definition was changed prior to rollforward DMS rebalancer is active
Correction	If the number of 'Not Normal' tablespaces is higher than you expect, drill down to the space metrics to examine the abnormal tablespaces and take appropriate action.

SQL Analysis pane

A lot of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured server down in terms of user response times.

Before you can identify problem SQL in your database, you need to understand the level of SQL activity that is occurring in the database. A high level overview of the type of SQL activity helps determine not only they types of statements that are being executed but also the type of internal activity they may be causing.

The following indicators are used on the Performance Analyst Home page to succinctly communicate the general overall performance level of the monitored database:

Static SQL	Dynamic SQL	Units of Work
Internal Rebinds	Internal Rollbacks Due to Deadlocks	

Static SQL

The static SQL metric is the number of static SQL statements that were executed on the database since the start of this instance of Performance Analyst.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > SQL Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > SQL Analysis pane
Metrics	Together with dynamic and failed SQL statements, this metric represents the number of SQL statements executed on the database.

Dynamic SQL

This metric is the number of static SQL statements that were executed on the database since the start of this instance of Performance Analyst.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > SQL Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > SQL Analysis pane

Metrics	Together with Static and Failed SQL statements, this metric represents the number of SQL statements executed on the database.

Failed SQL

The failed SQL metric is the number of failed SQL statements that were executed on the database since the start of this instance of Performance Analyst.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > SQL Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	Together with static and dynamic SQL statements, this metric represents the number of SQL statements executed on the database. This metric may also help you determine reasons for poor performance, since failed statements mean time wasted by the database manager. This results in lower throughput for the database.

Units of Work

This is the number of transactions executed on the database since the start of this instance of Performance Analyst.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > SQL Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > SQL Analysis pane
Metrics	A low number of units of work compared to the overall SQL activity (static + dynamic SQL statements) indicate long transactions. This may in turn be an indicator of poor concurrency and heavy log usage.

Internal Rebinds

The internal rebinds metric is the number of automatic rebinds (or recompiles) that have been attempted on the database since the beginning of this instance of the Performance Analyst.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > SQL Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > SQL Analysis pane
Metrics	Automatic rebinds are the internal binds the system performs when a package has been invalidated. The rebind is performed the first time the database manager needs to execute an SQL statement from the package. For example, packages are invalidated when you: *Drop an object, such as a table, view, or index, on which the plan depends* *Add or drop a foreign key* *Revoke object privileges on which the plan depends.* You can use this element to determine the level of database activity at the application or database levels. Since internal rebinds can significantly degrade performance, they should be minimized whenever possible.

Internal Rollbacks Due to Deadlocks

The internal rollbacks due to deadlocks metric is the total number of forced rollbacks initiated by the database manager because of a deadlock for the current instance of Performance Analyst. A rollback is performed on the current unit of work in an application selected by the database manager to resolve the deadlock.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > SQL Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > SQL Analysis pane

Metrics	This metric shows the number of deadlocks that have been broken and can be used as an indicator of concurrency problems. It is important because this metric lowers the throughput of the database.

Storage Analysis pane

While DBAs focus on memory settings and tuning SQL, they frequently forget how dangerous and insidious storage problems can be. This is not a good mindset because storage headaches can play a major role in wrecking an otherwise well-running database.

Storage problems generally take one of two forms:

- The 'hit-the-wall' variety that can bring things to a complete standstill.
- The 'performance vampire' kind that slowly drains the performance of a database over time.

Storage problems have the capability to bring the curtain down on a database very quickly, as in the case of a transaction log running out of free space. But storage problems can also silently work behind the scenes to slowly, but surely rob a database of its performance. For example, a hub table in a busy database might be accessed very quickly when an application is first given life, but over time, if it develops a heavy forwarded row problem, it can cause things to run very differently.

The Storage Analysis section of the Performance Analyst Home page displays the total used and free space of all DMS tablespaces in the database. It also provides the total and available log space for the database.

The following ratios succinctly communicate the general overall performance levels of the datasource:

DMS Space	Total SMS Space Used	Log Space
Last Backup		

DMS Space

The total used and total free metrics show the space details for DMS tablespaces. Specifically, they show the amount of used and free space on the DMS tablespaces in the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Storage Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Space > Home > Storage Analysis pane
Metrics	Unlike SMS tablespaces, the total available space to the DMS tablespaces is defined at the time the database is defined. The DBA needs to explicitly increase the amount of space on DMS tables by adding/extending tablespace containers (using the ALTER TABLESPACE statement). If left unattended, a DMS tablespace can either remain underutilized or fill up completely. Keeping an eye on the DMS space is important because the once it fills up, the applications trying to write to the database will come to a halt suddenly.
Correction	Go to the Space home page or drill-down to see the usage metrics for individual tablespaces and allocate/deallocate space to containers accordingly using ALTER TABLESPACE command.

Total SMS Space Used

This metric shows the total amount of SMS space used by the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Storage Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Space > Home > Storage Analysis pane

Metrics	The maximum size of SMS tablespaces is not set at the time the tablespaces are created. The maximum size that can be reached is the space available on the drives/volumes that the SMS tablespace containers are defined on. You need to periodically examine the available space of the drives/volumes (using OS metrics Space page) to make sure there is enough space available for your database requirements. Since the space is determined by the space available on the drives/volumes, remember that if other applications are reading and writing to the same devices (especially logs, traces, etc), DB2 may be periodically competing for space with these applications.
Correction	Ideally, try and isolate the DB2 SMS drives/volumes from other applications. If the OS level metrics show that you are running out of space on a volume that is used by a tablespace's containers, you can add more containers defined on volumes with more space, to the tablespace using the ALTER TABLESPACE command.

Log Space

The log space metrics show the amount of free and used log space on the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Storage Analysis pane
Metrics	Very high log space utilization (>75%) indicates that the database could run out of log space soon. Usually, consistently high log space utilization has one of the following causes: Large transaction size. Applications are performing a large number of updates on the database but are not committing the transactions frequently. The log space configured for the database is insufficient for the level of database activity Very low log space utilization (< 5) % can indicate an excessively large log file. It's possible you can free some of the space occupied by log files on disk.
Correction	If the log utilization is mostly low but increases to near 100% for short periods, you can increase the number of secondary log files (*logsecond* database configuration parameter). This allocates and deallocates log files as needed. If however, the log size is consistently high, and changes frequently, you can increase the primary log size and/or the log file size (*logfilsiz* and *logprimary* database configuration parameters). Also, check the SQL activity, under Users Drilldown>SQL activity to see if any applications are performing infrequent commits, and if possible such applications can be changed to have shorter transactions and hence less need for log space.

Last Backup

The last backup metric indicates when the database was last backed-up.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Storage Analysis pane

Workload Analysis pane

When experiencing performance slowdowns, some of the first questions you should answer are:

- Who is currently connected to the database?
- What resources are they using?
- What are they currently executing?

The Workload Analysis section of the Performance Analyst Home page provides insight into the leading resource hogs of a server, as well as a count of total active and inactive processes. Drill-down's are available so you can easily get detailed information into what each leading application is currently involved with.

>

Workload Analysis statistics include:

Application Activity Analysis	Applications Currently Executing	Idle Applications

Application Activity Analysis

Ideally, database resource usage should be spread evenly across connected applications. In many cases, however, a couple of renegade application can hog most of the resources and starve the other applications that are competing for the same resources. If you see a slowdown of the database, but cannot identify the root cause from ratio / bottleneck analysis, you need to examine any connected applications that are using more than their fair share of system resources (CPU, I/O, locklist, etc.)

Performance Analyst provides an easy way to identify the top users of the database resources and what percentage of total usage can be attributed to these applications.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Workload Analysis pane
Metrics	If any one application is using more than 50% of the total resource you should drilldown into that application (Users drill-down) to isolate the problem.

Applications Currently Executing

The applications currently executing statistic gives the number of applications for which the database manager is currently executing requests.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Workload Analysis pane
Metrics	You can use this number to understand how many of the database manager agent tokens are being used by applications connected to this database

Idle Applications

The idle applications statistic the number of applications that are currently connected to the database for which the database manager is not executing any requests

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Home > Workload Analysis pane
Metrics	You can use this statistic to help you understand the level of activity within a database and the amount of system resources currently in use.

Memory

The Memory category of DB2 Performance Analyst includes the following tabbed pages:

Caches Tab	Memory Pools Tab	Lock Memory Tab

Home

The Home page of the Memory performance category view has statistics in the following categories:

Key Ratio Analysis pane	Bottleneck Analysis pane	Memory Analysis pane

>

SQL Analysis pane	Workload Analysis pane	

Key Ratio Analysis pane

Database performance analysts typically use one of two methods for examining the performance levels of a database - ratio-based or wait/bottleneck-based. Ratio-based analysis involves examining a number of key database ratios that can be used to indicate how well a database is running. Performance ratios serve as very good roll-up mechanisms for busy DBAs to use for at-a-glance performance analysis. Many DBAs have large database farms to contend with and cannot spend time checking detailed wait-based analysis outputs for each and every database they oversee. Succinctly presented performance ratios can assist in such situations by giving DBAs a few solid indicators that can be quickly scanned to see if any database needs immediate attention.

While there are certainly many opinions as to what rules to follow, there are some standards you should always adhere to. To start with, many of the formulas that make up ratio-based analysis must be derived from delta measurements instead of cumulative statistics. Most of the global ratios that a DBA will examine come from the DB2 Snapshot Monitor. The Snapshot Monitor returns various metrics for the instance and database since the start of monitoring. Some of the metrics are cumulative (i.e., counters) and others are instantaneous (i.e., gauges).

In the case of cumulative metrics, the values can grow quite large over the monitoring period and the ratios calculated on these metrics can hide the skews that may occur during the monitoring. However, if delta statistics are used (taking, for a specified sampling period, the before and after counts of each statistic that make up a ratio), then an accurate and current portrayal of the database activity over the sampling period can be had.

A final thing to remember about using ratio-based analysis is that, while there are a number of rules of thumb that can be used as starting points to begin the evaluation of database performance, DBAs must determine each DB2 databases' individual personality with respect to the key performance ratios. Some hard and fast rules simply do not apply to every database. The danger in using blanket ratio standards is that they can lead the DBA to haphazardly take action, which can at times contribute nothing to the situation, and sometimes even degrade performance.

The following memory ratios are used on the Performance Analyst Key Ratio Analysis section to succinctly communicate the general overall memory performance levels of the monitored database:

Data Page Hit Ratio	Private Workspace Hit Ratio	Index Page Hit Ratio

The following topics, available on this pane, duplicate statistics available on the Home page:

Buffer Pool Hit Ratio	Catalog Cache Hit Ratio	Package Cache Hit Ratio
Shared Workspace Hit Ratio		

Data Page Hit Ratio

This metric is the percentage of all data reads satisfied because the page was already available in a bufferpool.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane
Metrics	The index page hits are used in conjunction with data page hits to calculate the overall buffer pool hit ratio. Use this element, in conjunction with the index page hit ratio when tuning the overall buffer pool hit ratio.

Private Workspace Hit Ratio

Each application executing on the database has access to the private workspace of the agents working for it.

>

The private workspace hit ratio is a percentage indicating how well the private SQL workspace is helping to avoid initialization of sections for SQL statements that are about to be executed. A high ratio indicates it is successful in avoiding this action.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Caches Tab
Correction	Shared workspace memory is allocated from the application heap (*app_ heap_sz* database configuration parameter) and increasing this may improve the private workspace hit ratio.

Index Page Hit Ratio

The index page hit ratio is the percentage of all index reads that were satisfied because the page was already available in a bufferpool.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > Key Ratio Analysis pane
Metrics	The index page hits are used in conjunction with data page hits to calculate the overall buffer pool hit ratio. Use this element, in conjunction with the data page hit ratio when tuning the overall buffer pool hit ratio.

Bottleneck Analysis pane

The Bottleneck Analysis section identifies the number of times when overflows or over-utilization may have occurred in areas of memory. An overflow occurs when the usage of a memory area goes beyond the allocated amount.

Overflows are serious bottlenecks and may cause unnecessary lock escalations, out of memory errors or performance degradations.

The following statistics are used on the Performance Analyst for DB2 Memory home page to succinctly communicate the general overall performance levels of the memory structures:

Catalog Cache Overflows	Hash Join Small Overflows	Package Cache Overflows
Private Workspace Overflows	Shared Workspace Overflows	

The following statistics, available on this pane, dupicte statistics availabe on the Bottleneck Analysis pane of the Home page.

Hash Join Overflows	Lock Escalations	Sort Overflows

Catalog Cache Overflows

The catalog cache overflow statistic is the number of times the catalog cache overflowed the bounds of its allocated memory.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	Catalog cache overflows can cause unnecessary lock escalations. This can result in loss of concurrency, or 'out of memory' errors from other heaps allocated to the database's shared memory. Overflows of the catalog cache can also cause performance degradation. DB2 reclaims the catalog cache space by evicting table descriptor information for tables, views, or aliases and/or authorization information that is not currently being used by any transaction.

Correction	Use this element with the catalog cache high watermark (Memory> Caches Drilldown) to determine whether the size of the catalog cache needs to be increased to avoid overflowing. If the number of overflows is large, the catalog cache may be too small for the workload. Enlarging the catalog cache can improve its performance. If the workload includes transactions that compile a large number of SQL statements referencing many tables, views, aliases, user-defined functions, or stored procedures in a single unit of work, then compiling fewer SQL statements in a single transaction can improve the performance of the catalog cache. Or, if the workload includes binding packages that contain many SQL statements referencing many tables, views, aliases, user-defined functions, or stored procedures, you can try splitting packages so that they include fewer SQL statements to improve performance.

Hash Join Small Overflows

The hash join small overflows statistic is the number of times that hash join data exceeded the available sort heap space by less than 10%.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	If the value of hash join small overflows is greater than 10% of hash join overflows, then you should consider increasing the sort heap size. Values at the application level can be used to evaluate hash join performance for individual applications.
Correction	Increase the *sortheap* database configuration parameter to reduce the number of overflows.

Package Cache Overflows

The package cache overflows metric is the number of times that the package cache overflowed the bounds of its allocated memory.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	Package cache overflows can cause unnecessary lock escalations. This can result in loss of concurrency, or ‘out of memory’ errors from other heaps allocated to the database’s shared memory. Overflows of the package cache can also cause performance degradation.
Correction	Use this element with the package cache high watermark (Memory > Caches Drilldown) to determine whether the size of the package cache needs to be increased to avoid overflowing.

Private Workspace Overflows

The private workspace overflows statistic is the number of times that private workspaces overflowed the bounds of their allocated memory.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	Private workspace overflows can cause performance degradation as well as ‘out of memory’ errors from other heaps allocated to the agent’s private memory.
Correction	Use this element with the private workspace high watermark (Memory -> Caches Drilldown) to determine whether the size of the private workspace needs to be increased to avoid overflowing.

>

Shared Workspace Overflows

The shared workspace overflows metric is the number of times that shared workspaces overflowed the bounds of their allocated memory.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > Bottleneck Analysis pane
Metrics	Overflows of shared workspaces can cause performance degradation. Overflows can also 'out of memory' errors from the other heaps allocated out of application's shared memory.
Correction	Use this element with the shared workspace high watermark (Memory > Caches Drilldown) to determine whether the size of the shared workspaces need to be increased to avoid overflowing.

Memory Analysis pane

This section helps you to track memory usage for database processes through Memory Analysis Heaps.

NOTE: The memory metrics for DB2 are only available if you are using a DB2 version 8 server and a DB2 version 8 client

NOTE: The memory management for DB2 (and hence the metrics) varies between platforms. You only see the memory metrics at the database level for UNIX system. If you are tracking the memory metrics for a Windows system, please refer to Instance -> Memory Pool Drilldown.

Memory Analysis Heaps

Under normal operating conditions you can track the memory usage for the following elements:

Application Heap	Database Heap	Application Control
Lock Manager Heap	Backup/Restore/Utility Heap	Statistics Heap
Package Cache Heap	Catalog Cache Heap	DFM Heap
Query Heap	Database Monitor Heap	Statement Heap
FCMBP Heap	Import Pool	Other Memory
Buffer Pool Heap		

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Home > Memory Analysis pane
Metrics	The various memory heap usages are useful when tuning the corresponding database configuration parameters.
Correction	To see greater details on the current utilizations, maximum available, and high watermarks, go to the Memory Pools tab of the Memory detail view. Based on these you can make decisions about whether you need to increase or decrease the maximum allocations for the heaps (mostly through database configuration parameters)

SQL Analysis pane

A lot of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag an otherwise well-configured server down in terms of user response times.

Before you can identify problem SQL in your database, you need to understand the level of SQL activity that is occurring in the database. A high level overview of the type of SQL activity helps determine not only they types of statements that are being executed but also the type of internal activity they may be causing.

>

The following indicators are used on the Performance Analyst Memory page to succinctly communicate the general overall performance level of the monitored database. They duplicate statistics available on the SQL Analysis pane of the Home page:

Static SQL	Dynamic SQL	Units of Work
Internal Rebinds	Internal Rollbacks Due to Deadlocks	

Workload Analysis pane

When your phone starts ringing with complaints of performance slowdowns, one of the first things you should get a handle on is:

- Who is currently connected to the database?
- What resources are they using?
- What are they currently executing?

The Workload Analysis section of the Performance Analyst Memory page provides insight into the leading applications that caused memory overflows to disk or lock escalations. Drill down's are available so you can easily get detailed information into what each leading application is currently involved with.

Application Activity Analysis pane

Top memory issues are displayed in the following categories:

- Top Locks
- Top Lock Escalations
- Top Cache Overflows
- Top Sort Overflows
- Top Workspace Overflows

Metrics

The list of top memory hogs can help you quickly identify the applications that are causing memory usage problems. You can drill down into the Memory tab views to see how these applications are using memory compared to their peers and the overall memory being used by the database. This is a good place to start, if you think that the memory problems you are noticing on the database may be as a result of poorly written applications rather than database configuration or overall system memory availability.

Caches Tab

The Caches tab of the Memory Detail page offers detailed views of cache usage metrics and includes the following statistics:

Application Details	Catalog Cache Usage Details	Package Cache Usage Details
Private Workspace Usage Details	Shared Workspace Usage Details	

>

The following statistics, available on this page, dupication statistics available on the Key Ratio Analysis pane of the Home page:

Catalog Cache Hit Ratio	Package Cache Hit Ratio	Shared Workspace Hit Ratio

The following statistic, available on this page, dupication statistics available on the Caches Tab of the Memory page:

Private Workspace Hit Ratio		

Application Details

The Application Details section shows the list of applications connected to the database and the various cache usage metrics for each application. It also shows the rolled up metrics at the database level in the same list. You can select the type of cache monitoring you intend to perform for the applications from the following:

- **Catalog Cache:** Shows the details of all catalog cache activity for connected applications
- **Package Cache:** Displays the details of all package cache activity for connected applications
- **Shared Workspace:** Gives the details of all shared workspace activity for connected applications

Private Workspace: This shows the details all of private workspace activity for connected applications

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Caches Tab
Correction	Depending on the category of cache details you select, you can view cache activity associated with the applications. This section is very useful in isolating the applications and the cache areas you need to tune by comparing the cache performance for all the applications and also the overall cache performance at the database level.

Catalog Cache Usage Details

This section shows detailed catalog cache usage metrics for the selected item. The metrics presented here include:

- **Catalog Cache Lookups:** The number of times that the catalog cache was referenced to obtain table descriptor information or authorization information
- **Catalog Cache Inserts:** The number of times that the system tried to insert table descriptor or authorization information into the catalog cache.
- **Catalog Cache Overflows:** The number of times that the catalog cache overflowed the bounds of its allocated memory
- **Catalog Cache High Watermark:** The largest size reached by the catalog cache

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Caches Tab
Metrics	These metrics are used for calculating the catalog cache hit ratio. They also give you an estimate of catalog access activity occurring in the database. You can use the catalog cache overflow metric with the catalog cache high watermark to determine if the catalog cache size is too small. If the catalog cache overflows frequently, you consider increasing the catalog cache size.

>

Correction	If catalog cache overflows are large, the catalog cache may be too small for the workload. Enlarging the catalog cache may improve its performance. If the workload includes transactions that compile a large number of SQL statements referencing many tables, views, aliases, user-defined functions, or stored procedures in a single unit of work, then compiling fewer SQL statements in a single transaction may improve the performance of the catalog cache. Or if the workload includes binding packages that contain many SQL statements referencing many tables, views, aliases, user-defined functions, or stored procedures, you can try splitting packages. If they include fewer SQL statements, performance can improve. Monitor the catalog cache metrics for your normal workload. If you see frequent overflows, you can increase the value of *catalogcache_sz* database configuration parameter to the *catalog cache high watermark/4096* (rounded up to whole integer).

Package Cache Usage Details

The Package Cache Usage Details section shows the detailed package cache usage metrics for the selected item. The metrics presented here include:

- **Package Cache Lookups:** The number of times that an application looked for a section or package in the package cache. At a database level, it indicates the overall number of references since the database was started, or monitor data was reset.
- **Package Cache Inserts:** The total number of times that a requested section was not available for use and had to be loaded into the package cache. This count includes any implicit prepares performed by the system.
- **Package Cache Overflows:** The number of times the package cache overflowed the bounds of its allocated memory
- **Package Cache High Watermark:** The largest size reached by the package cache.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Caches Tab
Metrics	These metrics are used to calculate the package cache hit ratio. They also give you an estimate of package cache access activity occurring on the database. You can use the package cache overflow metric with the package cache high watermark to determine if the package cache size is too small. If the package cache overflows frequently, you may need to increase the package cache size.
Correction	Monitor the package cache metrics for your normal workload. If you see frequent overflows, you can increase the value of the *pckcachesz* database configuration parameter to the *package cache high watermark/4096* (rounded up to whole integer).

Private Workspace Usage Details

This section shows the detailed private workspace usage metrics for the selected item. The metrics presented here include:

- **Private Workspace Lookups:** Lookups of SQL sections by an application in its agents' private workspace.
- **Private Workspace Inserts:** Inserts of SQL sections by an application into the private workspace.
- **Private Workspace Overflows:** The number of times that the private workspaces overflowed the bounds of its allocated memory.
- **Private Workspace High Watermark:** The largest size reached by the Private Workspace.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Caches Tab

Metrics	These metrics are used to calculate the private workspace hit ratio. They also give you an estimate of private workspace access activity occurring on the database. When the private workspace overflows, memory is temporarily borrowed from other entities in agent private memory. This can result in memory shortage errors from these entities or possibly performance degradation.
Correction	You can reduce the chance of overflow by increasing APPLHEAPSZ.

Shared Workspace Usage Details

This section shows the detailed shared workspace usage metrics for the selected item. The metrics presented here include:

- **Shared Workspace Lookups:** Lookups of SQL sections by applications in shared workspaces.
- **Shared Workspace Inserts:** Number of inserts of SQL sections by applications into shared workspaces.
- **Shared Workspace Overflows:** The number of times that shared workspaces overflowed the bounds of their allocated memory.
- **Shared Workspace High Watermark:** The largest size reached by shared workspaces.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Caches Tab
Metrics	These metrics are used to calculate the shared workspace hit ratio. They also give you an estimate of shared workspace access activity occurring on the database. When the shared workspace overflows, memory is temporarily borrowed from other entities in application shared memory. This can result in memory shortage errors from these entities or possibly performance degradation.
Correction	You can reduce the chance of overflow by increasing APPL_CTL_HEAP_SZ.

Memory Pools Tab

The Memory Pools tab offers the following detailed views:

Application Details	Memory Pool Details	Memory Pool Usage
Memory Pool Utilization		

Application Details

This section shows the list of applications connected to the database. You can select an application to show the breakdown of the memory pool usage by the application.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Memory Pools Tab

Memory Pool Details

This section shows the memory pool usage details for the selected application (or database level) in the application details section. The metrics shown here include:

- **Pool ID:** This is the kind of memory pool. Each pool type is only shown once
- **Pool Current Size:** The current size of a memory pool.
- **Pool Max Size:** The internally configured size of a memory pool in DB2.

>

Pool High Watermark: The largest size of a memory pool since its creation.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Memory Pools Tab IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Memory Pools tab
Metrics	The nature of memory_pool data elements varies between platforms. On Windows systems, no memory usage is reported at the database level, while on UNIX systems, memory is reported at the database level. Instead of reporting this memory at the database level, the system monitor for Windows systems reports it in instance -level snapshots. This divergence in reporting is due to differences in the underlying memory architecture between Windows systems and UNIX systems.

Memory Pool Usage

This section shows the memory pool usage details of the selected memory pool over time including the current size of the memory pool, its high watermark, and the maximum size of the memory pool.

Location

- IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Memory Pools Tab

Memory Pool Utilization

This section shows the percentage utilization of the selected memory pool.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Memory Pools Tab IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Memory Pools tab
Metrics	You can use this metric to see if a memory pool is nearly full. You can diagnose the problems with specific memory pools by monitoring their utilization over time.
Correction	If the value of the pool utilization is consistently close to or exceeds 100%, you may need to increase the configuration parameter associated with that pool.

Lock Memory Tab

The Lock Memory tab of the Memory Detail view shows the following statistics:

- Lock List Utilization (dupiicates a statistics available on the Key Ratio Analysis pane of the Home page)
- Lock Escalations

Lock Escalations

This section shows the metrics for lock escalations in the database. The metrics shown here are:

- **Lock Escalations:** The number of times that locks have been escalated from several row locks to a table lock.
- **Exclusive Lock Escalations:** The number of times that locks have been escalated from several row locks to one exclusive table lock, or the number of times an exclusive lock on a row caused the table lock to become an exclusive lock.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Memory > Lock Memory Tab

Metrics	A lock is escalated when the total number of locks held by an application reaches the maximum amount of lock list space available to the application, or the lock list space consumed by all applications is approaching the total lock list space. The amount of lock list space available is determined by the maxlocks and locklist configuration parameters. When an application reaches the maximum number of locks allowed and there are no more locks to escalate, it will then use space in the lock list allocated for other applications. When the entire lock list is full, an error occurs. This data item includes a count of all lock escalations, including exclusive lock escalations. There are several possible causes for excessive lock escalations: *The lock list size (locklist) may be too small for the number of concurrent applications* *The percent of the lock list usable by each application (maxlocks) may be too small* *One or more applications may be using an excessive number of locks.*
Correction	To resolve these problems, you may be able to: *Increase the locklist configuration parameter value. See the Administration Guide for a description of this configuration parameter.* *Increase the maxlocks configuration parameter value. See the Administration Guide for a description of this configuration parameter.* *Identify the applications with large numbers of locks (see locks_held_top), or those that are holding too much of the lock list, using the following formula:* *(((locks held * 36) / (locklist * 4096)) * 100)* *and comparing the value to maxlocks. These applications can also cause lock escalations in other applications by using too large a portion of the lock list. These applications may need to resort to using table locks instead of row locks, although table locks may cause an increase in lock_waits and lock_wait_time.*

I/O

In addition to a Home page, the I/O category of DB2 Performance Analyst includes the following tabbed pages are available on the I/O Detail view:

Applications Tab	Bufferpools Tab	Logging Tab
Tablespaces Tab		

Home

The Home Page of the I/O performance category page displays the following vital DB2 I/O statistic categories:

Key System Analysis pane	Bottleneck Analysis pane	SQL Analysis pane
Non Buffered I/O Activity	Workload Analysis pane	

Key System Analysis pane

The root cause of many database problems can be attributed to I/O-related operations. I/O monitoring is essentially a review of your physical database design. All physical storage characteristics and placements, table and index designs, and the speed with which all processes works are on display when I/O is monitored. Because a database's main index of performance is measured by how fast I/O needs are satisfied, it is your responsibility to quickly determine if a reported database slowdown is I/O related.

The following statistics are used on the Performance Analyst for DB2 I/O Home Page to succinctly communicate the general overall performance levels of I/O:

>

Efficiency

Async Read Ratio	Async Write Ratio	Victim Page Cleaners %

Resource Utilization

Log Space Cleaners Triggered	Victim Page Cleaners Triggered	Threshold Cleaners Triggered

Async Read Ratio

Async read ratio is the percentage of all index and data reads that performed asynchronously by the database manager prefetchers.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	Prefetching refers to the technique whereby the Database Manager can read several pages on the physical reads device simultaneously into a bufferpool in anticipation of access to these pages. Since each physical read operation is costly, a 50% async read ratio is desirable. It shows that prefetchers are working well and read waits are minimal.
Correction	The *num_ioserver* database configuration parameter defines the number of prefetchers that are available for the database. To get an optimal async read ratio, you should set the *num_ioserver* database configuration parameter to 1-2 more than the number of physical storage devices used by the database to ensure that asynchronous I/O is occurring. This parameter should be changed judiciously. Having too many prefetchers invariably results in high system I/O because prefetchers can read many more pages than required into the bufferpools.

Async Write Ratio

Async write ratio is the percentage of all index and data writes that were performed asynchronously by the database manager bufferpool page cleaners.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	Page cleaners write changed pages from bufferpool to disk before the space in the bufferpool is needed by a database agent. If the page cleaners are not working well, the problem may manifest itself in two ways: *The database agents need to synchronously free space in bufferpools resulting in poorer response time.* *If the system crashes, recovery time is greater because there will be a large number of pages that would not have been committed to disk at the time of the crash and they require processing for a large number of log records.* An overall async write ratio of 50% is desirable for most transactional database systems. If your database is 'query only' (i.e., updates are rarely or never performed), it is fine to have an async write ratio of 0 to less than 50%. Even in 'query only' databases if the queries create temp tables on the database for intermediate results, a 30-40% async write ratio is desirable.
Correction	The async write ratio can be tuned using the *num_iocleaners* database configuration parameter, which specifies the number of asynchronous page cleaners for a database. Increasing the number of page cleaners generally results in a higher async write ratio. The following rules of thumb should be followed when tuning this parameter: *For most transactional systems set this parameter to between one and the number of physical storage devices used by the database.* *For workloads that consists mainly of simple reads from the database, set this parameter to zero.* *When you define very large bufferpools on your database you may need to increase the num_iocleaners* Monitor the victim page cleaners% metric when tuning the *num_iocleaners*.

Victim Page Cleaners %

The victim page cleaners % metric is the percentage of times the page cleaners were invoked because synchronous write was needed during the victim buffer replacement in the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	Usually this metric should be maintained around the 50-60% level. If this ratio much lower than 50%, it can indicate that you have defined too many page cleaners. If the *chngpgs_thresh* database configuration parameter is set too low, the page cleaners may be writing out pages that will be dirtied later. Aggressive cleaning defeats one purpose of the bufferpool, that is to defer writing to the last possible moment. If this ratio is high, you may have too few page cleaners defined. This can increase recovery time after failures.
Correction	You may increase or decrease the *chngpgs_thresh* database configuration parameter to bring the victim page cleaner % to optimal range. Usually, for databases with heavy transaction workloads, *chngpgs_thresh* should be set to around 60%. For a database with only a few very large tables you can try and increase this parameter to a higher value to see if you get performance improvements.

NOTE: DB2 version 8.1.4 introduces the DB2_USE_ALTERNATE_PAGE_CLEANING registry variable. If this is set to ON, this metric may not be relevant.

Log Space Cleaners Triggered

The log space cleaners triggered metric is the number of times a page cleaner was triggered because log space usage reached a predefined threshold for the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	The log space cleaners are triggered every time the space used by the DB2 log reaches the limit set in the *softmax* database configuration parameter. This parameter specifies the percentage of primary log size at which the cleaners are triggered. By default this parameter is set to 100. A low rate of log space cleaners getting triggered may indicate that logs are not being written to disk frequently enough and that you may need to process a large number of log records and redundant log records in the event of a crash recovery. High rates of log space cleaners getting triggered can indicate that your primary log is too small or that you have set the *softmax* too high. In either case, a very high rate of log cleaners being triggered may adversely impact database performance. It is important to keep in mind, however, that more page cleaner triggers and more frequent soft checkpoints increase the overhead associated with database logging.This can have an impact on the performance of the Database Manager. Also, more frequent soft checkpoints may not reduce the time required to restart a database, if you have: *Very long transactions with few commit points.* *A very large bufferpool and the pages containing the committed transactions are not written back to disk very frequently. The use of asynchronous page cleaners can help avoid this situation.* In both of these cases, the log control information kept in memory does not change frequently and there is no advantage in writing the log control information to disk, unless it has changed
Correction	Examine this metric together with other page cleaner metrics and the async write ratio to determine if excessive numbers of log page cleaners are being triggered. If this is true, you need to either increase your primary log size or adjust the *softmax* parameter to a higher value.

NOTE: Decreasing the log page cleaner triggering rate can impact the overall crash recovery time so it needs to be done judiciously.

Victim Page Cleaners Triggered

The victim page cleaners triggered statistic is the number of times the page cleaners were invoked because synchronous write was necessary during the victim buffer replacement in the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	If the *chngpgs_thresh* database configuration parameter is set too low, the page cleaners may be writing out pages that will be dirtied later. Aggressive cleaning defeats one purpose of the bufferpool, which is to defer writing to the last possible moment. If this ratio is high, it may indicate that you have too few page cleaners defined. Too few page cleaners increases recovery time after failures.
Correction	You may increase or decrease the *chngpgs_thresh* database configuration parameter to bring the victim page cleaner triggering rate to optimal range. For databases with heavy transaction workloads, *chngpgs_thresh* should be set to around 60%. For a database with only a few very large tables you can try and increase this parameter to a higher value to see performance improvements.

NOTE: DB2 v8.1.4 introduces the DB2_USE_ALTERNATE_PAGE_CLEANING registry variable. If this is set to ON, this metric may not be relevant

Threshold Cleaners Triggered

The threshold cleaners triggered metric is the number of times a page cleaner was invoked because a bufferpool had reached the dirty page threshold criterion for the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > Key System Analysis pane
Metrics	The dirty page threshold is set by the *chngpgs_thresh* database configuration parameter. It is a percentage applied to the bufferpool size. When the number of dirty pages in the pool exceeds this value, the cleaners are triggered. If this value is set too low, pages might be written out too early, requiring them to be read back in. If set too high, then too many pages may accumulate, requiring users to write out pages synchronously.
Correction	For databases with frequent update transactions, you can generally ensure that there are enough clean pages in the bufferpool by setting the *chngpgs_thresh* parameter value to be equal-to or less-than the default value. A percentage larger than the default can help performance if your database has a small number of very large tables.

Bottleneck Analysis pane

High prefetch waits and file closures are good indicators of I/O waits at the database. Often, heavy I/O activity may not be due to non-optimal settings of the database configuration parameters. Rather, the underlying cause can be poor physical layout of a few tables/tablespaces or a few poorly written/suboptimal database applications.

It is usually best to try and tune individual objects/applications before attempting to tune database configuration parameters. Database configuration changes can have a profound impact on every application and object that resides in the database and if there are only a couple of apps/objects dragging down the performance of the entire database, it may be best to first isolate and tune these.

I/O Bottleneck analysis provides information on applications and tablespaces that are performing the highest number of read and write operations on the database at a glance. It also shows the bufferpool I/O metrics that point to direct I/O waits.

The following statistics are used on the Performance Analyst for DB2 I/O home page to succinctly communicate the general overall performance levels of I/O:

Hottest Tablespaces	Hottest Apps	Total Database Files Closed
Total Time Spent Waiting for Prefetchers		

>

Hottest Tablespaces

This section displays the list of top three tablespaces with the highest number of physical reads and writes.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Metrics	Understanding the storage-level hot spots of a database is beneficial for a two reasons. First, you can get a feel for overworked physical disks by viewing I/O statistics at the tablespace level. If a particular tablespace is under too much strain, you can attempt to relocate tablespace containers to other less-used devices. You can also try to create new tablespaces on different disks and move hot objects to them (assuming, of course, you have extra disks available). Second, if you have used standard DBA practices and placed indexes in their own tablespace, you can review the I/O statistics for the hottest tablespaces and see if the indexes are actually being used.
Correction	Some areas to consider when reviewing the hottest tablespaces display are: *A temp tablespace showing high volumes of physical I/O could indicate a problem with excessive sort overflows.* *Quickly review all physical I/O for each drive/file system at the OS level and get a feel for the overworked disks on your server. If you have underutilized disk drives (with their own controllers), then you should consider relocating some tablespaces/containers that exhibit high I/O characteristics to those drives* *To minimize contention at the physical disk level, and to improve performance overall, it is generally recommended that a database have its tablespaces and log files physically separated onto different drives/file systems.*

Hottest Apps

This metric shows the top three applications connected to the database with the highest number of physical reads and writes.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Metrics	A few applications with suboptimal queries/database accesses can drag down the performance of the entire database. The hottest apps section allows you to get a good feel for the applications that are causing the highest number of reads and writes on the database. It is also a good indicator of where you will reap the greatest benefit when tuning the application's performance.
Correction	Drill down to the application I/O pages and see how the high I/O applications are behaving compared to other applications. Check the SQL Snapshot Metrics for the high read/write apps to see what queries these applications are running on the database and try to optimize these queries.

Total Database Files Closed

Both SMS tablespaces and DMS tablespace file containers are treated as files in the Database Manager's interaction with the operating system, and file handles are required. There is a limit to the number of files that can be open concurrently by each database agent and a file is closed when the limit of open files for an agent is reached and the agent needs to open another file. This metric gives the number of times a database file was closed because the limit for concurrently open files was reached.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Metrics	Both SMS table spaces and DMS tablespace file containers are treated as files in the Database Manager's interaction with the operating system. More files are generally used by SMS tablespaces compared to the number of containers used for a DMS file table space. Thus, if you use SMS tablespaces, the database needs to open a larger number of files compared to what would be required for DMS file tablespaces. Closing and opening files is a very expensive operation and a large number of closed files can adversely affect the performance of the database. It is recommended that the number of files closed be as small as possible (preferably 0).

Correction	If you notice that files are being closed during the monitoring, first check the rate at which the files are being closed at the database level, from the I/O drill-down. If you see files being closed frequently, increase the *maxfilop* database configuration parameter. Note that the operating system on which the database is running also has a limit on the number of files that can be opened concurrently. Make sure that you do not increase the *maxfilop* parameter beyond the operating system limit.

Total Time Spent Waiting for Prefetchers

This is the total time database applications spent waiting for prefetchers to finish loading.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > Bottleneck Analysis pane
Correction	If you see a high number for this metric, it indicates there is a bottleneck in the prefetcher operations. Experiment with changing the number of I/O servers *(num_ioserver)* and the I/O server sizes.

SQL Analysis pane

Most of database's I/O performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag down an otherwise well-configured server in terms of user response times. This section shows a snapshot of row level SQL operations that have been performed on the database since you started monitoring. This gives you an insight into the row level of read/write activity that is currently occurring within the database and comparing this to the prefetcher, and cleaner activity allows you to relate the two.

The following statistics are used on the Performance Analyst for DB2 I/O home page to succinctly communicate the general overall performance levels of I/O:

Rows Deleted	Rows Inserted	Rows Selected
Rows Updated	Rows Read	

Rows Deleted

The rows deleted metric is the total number of row deletions attempted on the database. It includes the internal row deletions that were caused by constraints or triggers.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > SQL Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > SQL Analysis pane

Rows Inserted

The rows inserted metric is the total number of row insertions attempted on the database. This metric includes the internal row insertions that were caused by triggers.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > SQL Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > SQL Analysis pane

>

Rows Selected

The rows selected statistic is the total number of rows selected at the database. This metric does not include a count of rows accesses for operations such as COUNT(*) or joins.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > SQL Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > SQL Analysis pane

Rows Updated

The rows updated metric is the total number of row updates attempted on the database. This metric includes the internal row updates that were caused by triggers or constraints

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > SQL Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > SQL Analysis pane

Rows Read

The rows read statistic is the number of rows read from the table. This count is not the number of rows that were returned to the calling application. Rather, it is the number of rows that had to be read in order to return the result set.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > SQL Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > SQL Analysis pane

Non Buffered I/O Activity

Some database operations such as backups, loads, restores, LOB (large object I/O) bypass the bufferpools and read and write directly to storage. Sometimes, even if your database I/O configurations and all of the other activities are within the normal thresholds, the I/O performance may be poor because LOB reads/writes or load/backup/restore operations are causing or waiting on I/O contentions on your storage devices.

The non buffered I/O activity section allows you to quickly identify the level of direct I/O activity on the database. The level of direct I/O activity is succinctly presented using the following metrics:

Direct Reads from the Database	Direct Writes to the Database	Direct I/O Averages

Direct Reads from the Database

The direct reads from the database metric identifies the number of read operations that did not use a bufferpool. Direct reads are performed for the following database operations:

- Reading LONG VARCHAR columns
- Reading LOB (BLOB/CLOB/DBCLOB) columns
- Performing a backup

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > SQL Analysis pane

Direct Writes to the Database

The direct writes to the database metric identifies the number of write operations that did not use a bufferpool. Direct writes are performed for the following database operations:

>

- Writing LONG VARCHAR columns
- Writing LOB (BLOB/CLOB/DBCLOB) columns
- Performing a load
- Performing a restore

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > SQL Analysis pane

Direct I/O Averages

The following values show how well the direct I/O operations are performing:

- Avg. Sectors / Read: Average number of sectors read for each direct read operation
- Avg Sectors / Write: Average number of sectors written for each direct write operation
- Avg Read Time: Average read time per sector for direct reads

Avg Write Time: Average write time per sector for direct writes

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > SQL Analysis pane
Metrics	Understanding the level of direct I/O averages is important. A high level of activity with a slowdown in database performance can indicate that you need to have dedicated and/or high performance storage devices for your LOB objects or backup and load images. These metrics, when viewed with the device metrics at the OS level can also indicate I/O resource conflicts with other processes.
Correction	You can monitor these metrics for the LOB storage devices at the OS level to see if they are inordinately high for the database compared to the OS level. If so, there may be an I/O conflict with another process.

Workload Analysis pane

When experiencing performance slowdowns, some of the first questions you should answer are:

- Who is currently connected to the database?
- What resources are they using?
- What are they currently executing?

The Workload Analysis section of the Performance Analyst I/O page provides insight into the leading resource hogs of a server, as well as a count of total active and inactive processes. Drill-down's are available so you can easily get detailed information into what each leading application is currently involved with.

Workload Analysis statistics includes:

- Top I/O Apps

Top I/O Apps

The Top I/O Apps section gives a list of the apps with performing that are performing the highest number of most aggressive I/O activities The list includes:

- Application with most bufferpool misses
- Application with highest direct reads and writes
- Application with highest sort overflows

>

- Application with highest direct read/write time

Application with highest bufferpool read write time

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Home > Workload Analysis pane
Correction	Drill down to the application I/O level to see the details and compare the metrics with the application's peers, and the database. In a lot of cases simple SQL rewrites, logic changes, and defining indexes based on access patterns can make a world of difference.

Applications Tab

This section shows the list of applications connected to the database and the various I/O metrics for each application. It also shows the rolled up metrics at the database level in the same list. You can select the type of I/O monitoring you intend to perform for the applications from the following:

- Direct I/O: This shows the details of all direct I/O activity for the applications.
- Buffered I/O: This shows the details of all bufferpool I/O activity for the applications connected to the database.
- Extended Storage: This shows the level of extended storage usage by the applications.
- **Sorts & Joins**: This shows the details for the number of files closed for the various applications.
- **Block / Vectored I/O:** This shows the detailed of block and vectored I/O performed by the applications.

Correction

Depending on the type of details you select, you can view the different types of I/O activity associated with the applications. This section is very useful in isolating the applications, and the I/O areas you need to tune by comparing the I/O performance for all the applications and also the overall I/O performance at the database level.

Direct I/O

The Direct I/O Application or Tablespace Details view lets you the following statistics:

Direct Reads and Writes	Direct Read and Write Requests	Direct Read and Writes Time

Direct Reads and Writes

This section shows the rate at which direct read/write operations are occurring for the selected element. Direct I/O operations are those that do not use the bufferpools. Direct read and write operations are performed when:

- Reading/Writing LONG VARCHAR columns
- Reading/Writing LOB data
- Performing a backup operation
- Performing a restore operation
- Performing a load operation

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Applications Tab
Metrics	When using the performance analyst to track I/O, these elements help you distinguish the database I/O from the non-database I/O on storage devices.

>

Direct Read and Write Requests

This section shows the rate at which direct read/write operations requests are occurring for the selected element. Direct I/O operations are those that do not use the bufferpools. Direct read and write operations are performed when:

- Reading/Writing LONG VARCHAR columns
- Reading/Writing LOB data
- Performing a backup operation
- Performing a restore operation
- Performing a load operation

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Applications Tab
Metrics	Use these metrics together with the Direct R/W metrics and Direct R/W request metrics to determine if the rate at which the read/write occurring is similar to the rate at which they are being satisfied. If there is a large discrepancy between the Direct I/O request rates and Direct I/O rates, it may mean that there are I/O contentions on the device(s) where the direct I/O operations are being performed.
Correction	If you suspect there are I/O contentions, you can investigate whether the contentions are being caused by the DB2 processes or other processes by examining the I/O metrics at the operating system level. Consider using dedicated devices for the direct I/O operations to reduce the contention levels.

Direct Read and Writes Time

This section shows the elapsed time to perform the direct I/O operations. Direct I/O operations are those that do not use the bufferpools. Direct read and write operations are performed when:

- Reading/Writing LONG VARCHAR columns
- Reading/Writing LOB data
- Performing a backup operation
- Performing a restore operation

Performing a load operation

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Applications Tab
Metrics	Use these metrics together with the Direct R/W metrics to determine average time it is taking to perform the read write operations. If the level of Direct I/O rates is low and the Direct RW times, it may mean that there are I/O contentions on the device(s) where the direct I/O operations are being performed.
Correction	If you suspect there are I/O contentions, you can investigate whether the contentions are being caused by the DB2 processes or other processes by examining the I/O metrics at the operating system level. Consider using dedicated devices for the direct I/O operations to reduce the contention levels. If the contention levels are low, you may also benefit from using faster storage devices for direct I/O.

Buffered I/O

The Buffered I/O Application Details view lets you see the following statistics:

Hit Ratios	Write/Read Ratios	

Hit Ratios

The DB2 database server reads and updates all data from a bufferpool because memory access is much faster than disk access. Data is copied from disk to a bufferpool as needed by the applications using the database. When the server needs to read/write data and the data is already in the bufferpool, no disk access is required. However, if the data is not in the bufferpool, it needs to be read from the disk, which is significantly slower process.

The buffer pool hit ratio indicates the percentage of time that the database server did not need to load a page to service the request for a page. The calculation takes into account all the index and data pages that were requested

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Applications Tab
Metrics	Avoiding disk I/O is the main issue when trying to improve performance tuning. A high buffer pool hit ratio is desirable because it indicates a lower frequency of synchronous disk I/O. A database where data is accessed uniformly from very large tables will have a poor hit ratio. There is little you can do to improve the performance in such cases.
Correction	The buffer pool hit ratio on the database Home page is the overall hit ratio of the database. First, drill down to the bufferpool level and check the individual buffer pool hit ratios to identify the elements that may require tuning. Increasing the bufferpool size generally improves the buffer pool hit ratio. You can use one of the following methods to increase the bufferpool size depending on the layout of the data: *If the tablespaces using the bufferpool have tables and indexes on them, increase the bufferpool size incrementally until the index page hit ratio stops increasing for that bufferpool. You will not usually get any benefit increasing the bufferpool size after you reach that point.* *If the tablespaces associated with the bufferpool have only indexes on them, increase the bufferpool size until the index page hit ratio stops increasing.* *If the tablespaces associated with the bufferpool have only data on them, increase the bufferpool size until the data page hit ratio stops increasing.* *One general design guideline is to try and use different tablespaces for indexes and tables and associate different bufferpools with them. This generally improves performance and makes bufferpool tuning much easier.*

Write/Read Ratios

This section shows the overall level of read / write activity being performed on the selected element.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Applications Tab

Extended Storage

The Extended Storage Application or Transaction Details view lets you see the following statistics:

Index I/O on Extended Storage	Data I/O on Extended Storage	

Index I/O on Extended Storage

This section shows the Index Pages copied to/from extended storage for the selected element. On 32-bit platforms, virtual addressable memory is usually limited to between 2 and 4 GB. If your 32-bit machine has more real addressable memory than the maximum amount, you can configure any additional real addressable memory beyond virtual addressable memory as an extended storage cache. Any of the defined bufferpools can use an extended storage cache to improve performance. You define the extended storage cache as a number of memory segments.

The bufferpools perform first-level caching, and any extended storage cache is used by the bufferpools as secondary-level caching. Ideally, the bufferpools hold the data that is most frequently accessed, while the extended storage cache hold data that is accessed less frequently.

Pages are copied from the bufferpool to extended storage, when they are selected as victim pages. This copying is required to make space for new pages in the bufferpool.

Required index pages are copied from extended storage to the bufferpool, if they are not in the bufferpool, but are in extended storage. This copying may incur the cost of connecting to the shared memory segment, but saves the cost of a disk read.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Applications Tab
Metrics	Because an extended storage cache is an extension to a bufferpool, it must always be associated with one or more specific bufferpools. Therefore, only the bufferpools that are declared to take advantage of advantage of the extended storage use it. The CREATE and ALTER BUFFERPOOL statements have the attributes NOT EXTENDED STORAGE and EXTENDED STORAGE that control cache usage. By default neither IBMDEFAULTBP nor any newly created bufferpool will use extended storage. Use these metrics together with the Data I/O on Extended Storage metrics to determine the level of I/O being performed on extended storage. If there is significant activity on this section, you may improve performance by increasing the size of extended storage.
Correction	The amount of extended storage for a database is determined by the following two database configuration parameters: ***num_estore_segs*** *defines the number of extended storage memory segments. The default for this configuration parameter is zero, which specifies that no extended storage cache exists.* ***estore_seg_sz*** *defines the size of each extended memory segment. This size is determined by the platform on which the extended storage cache is used.*

Data I/O on Extended Storage

This section shows the DataPages copied to/from extended storage for the selected element.

On 32-bit platforms, virtual addressable memory is usually limited to between 2 and 4 GB. If your 32-bit machine has more real addressable memory than the maximum amount, you can configure any additional real addressable memory beyond virtual addressable memory as an extended storage cache. Any of the defined bufferpools can use an extended storage cache to improve performance. You define the extended storage cache as a number of memory segments.

The bufferpools perform the first-level caching, and any extended storage cache is used by the bufferpools as secondary-level caching. Ideally, the bufferpools hold the data that is most frequently accessed, while the extended storage cache hold data that is accessed less frequently.

Pages are copied from the bufferpool to extended storage, when they are selected as victim pages. This copying is required to make space for new pages in the bufferpool. Required pages are copied from extended storage to the buffer pool, if they are not in the bufferpool, but are in extended storage. This copying may incur the cost of connecting to the shared memory segment, but saves the cost of a disk read.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Applications Tab
Metrics	Because an extended storage cache is an extension to a bufferpool, it must always be associated with one or more specific bufferpools. Therefore, only the bufferpools that are declared to take advantage of advantage of the extended storage use it. The CREATE and ALTER BUFFERPOOL statements have the attributes NOT EXTENDED STORAGE and EXTENDED STORAGE that control cache usage. By default neither IBMDEFAULTBP nor any newly created bufferpool will use extended storage. Use these metrics together with the Index I/O on Extended Storage metrics to determine the level of I/O being performed on extended storage. If there is significant activity on this section, you can improve performance by increasing the size of extended storage.

>

Bufferpools Tab

This section shows the list of bufferpools on the database and the various I/O metrics for each bufferpool. It also shows the rolled up metrics at the database level in the same list. You can select the type of monitoring you intend to perform for the bufferpools from the following:

- Direct I/O: This shows the details of all direct I/O activity for the various bufferpools
- Buffered I/O: This shows the details of all bufferpool I/O activity for the various bufferpools in the database
- Extended Storage: This shows the level of extended storage usage by the various bufferpools that use extended storage
- Total Database Files Closed: This shows the details for the number of files closed for the various bufferpools in the database
- Block/Vectored I/O: This shows the detailed of block and vectored I/O performed by the various bufferpools.

Correction

Depending on the type of details you select, you can view the different types of I/O activity associated with the bufferpools. This section is very useful in isolating the bufferpools, and the I/O areas you need to tune by comparing the I/O performance for all the bufferpools and also the overall I/O performance at the database level.

Block/Vectored I/O

The Block/Vectored I/O section of the Bufferpool Details view displays the following statistics:

Vectored I/O	Block I/O	

Vectored I/O

The Vectored I/O section gives the metrics for vectored I/O requests and reads.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Bufferpools Tab
Metrics	Use this element to determine how often vectored I/O is being done. The number of vectored I/O requests is monitored only during sequential prefetching.

Block I/O

The Block I/O section gives the metrics for block I/O requests and reads. If block-based bufferpool is enabled, this section reports how often block I/O is being done. The number of block I/O requests is monitored only during sequential prefetching when using block-based bufferpools.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Bufferpools Tab

Metrics	If block-based bufferpool is enabled and the block I/O activity is very low, or close to the values in the vectored I/O section, consider changing the block size. This state can be an indication of the following: *The extent size of one or more tablespaces bound to the bufferpool is smaller than the block size specified for the bufferpool.* *Some pages requested in the prefetch request are already present in the page area of the bufferpool.* The prefetcher allows some wasted pages in each bufferpool block, but if too many pages are wasted, then the prefetcher will decide to perform vectored I/O into the page area of the bufferpool. To take full advantage of the sequential prefetch performance improvements that block-based bufferpools provide, it is essential to choose an appropriate value for the block size. This can, however, be difficult because multiple tablespaces with different extent sizes can be bound to the same block-based bufferpool. For optimal performance, it is recommended that you bind tablespaces with the same extent size to a block-based bufferpool with a block size equal to the extent size. Good performance can be achieved when the extent size of the tablespaces are greater than the block size, but not when the extent size is smaller than the block size. For example, if extent size is 2 and block size is 8, vectored I/O would be used instead of block I/O (block I/O would have wasted 6 pages). A reduction of the block size to 2 would solve this problem.
Correction	You can alter the block size for a buffer pool using the ALTER BUFFERPOOL ddl statement.

Logging Tab

The Logging tab of the I/O Detail view displays Logging I/O statistics. This section shows the level of read/write activity being performed by the logger. The information presented here is:

Log Reads: The number of log pages read from the disk by the logger per second.

Log Writes: The number of log pages written to the disk by the logger.

Metrics

You can use this page with the operating system level I/O metrics, to quantify the amount of I/O on a device that is attributable to database activity.

Tablespaces Tab

The Tablespace Tab of the I/O detail view shows the list of tablespaces in the database and the I/O metrics for each tablespace. It also shows the rolled-up metrics at the database level in the same list. You can select the type of monitoring you intend to perform for the tablespaces from the following:

- Direct I/O: Shows the details of all direct I/O activity for the various tablespaces
- Buffered I/O: Shows the details of all bufferpool I/O activity for the various tablespaces in the database
- Extended Storage: Shows the level of extended storage usage by the tablespaces that use extended storage
- Total Database Files Closed: Shows details for the number of files closed for various tablespaces in the database

Correction

Depending on the type of details you select, you can view the different types of I/O activity associated with the tablespaces. This section is very useful in isolating the tablespaces, and the I/O areas you need to tune by comparing the I/O performance for all the tablespaces and also the overall I/O performance at the database level.

Buffered I/O

The following statistics are available on the Buffered I/O section of the Tablespace Details of the I/O view:

Hit Ratios	Async I/O	Write/Read Ratios

Async I/O

Async read ratio is the percentage of all index and data reads that were performed asynchronously by the database manager prefetchers. Async write ratio is the percentage of all index and data writes that were performed asynchronously by the database manager bufferpool page cleaners.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > I/O > Tablespaces Tab
Metrics	Prefetching refers to the technique whereby the Database Manager can read several pages on the physical reads device simultaneously into a bufferpool in anticipation of an access to these pages. Since each physical read operation is costly, a 50% async read ratio is desirable since it shows that the prefetchers are working well and read waits are being minimized. Page cleaners write changed pages from bufferpool to disk before the space in the bufferpool is needed by a database agent. If the page cleaners are not working well, the problem may manifest itself in two ways: *The database agents will need to synchronously free up space in bufferpools resulting in poorer response time.* *If the system crashes, the recovery time of the system will be greater because there will be a large number of pages that would not have been committed to disk at time of crash and will require processing of a large number of log records.* An overall async write ratio of 50% is desirable for most transactional database systems. If your database is 'query only' (i.e., updates are rarely or never performed), it is okay to have an async write ratio of 0 to less than 50%. Even in 'query only' databases if the queries create temp tables on the database for intermediate results, a 30-40% async write ratio is desirable.
Correction	*num_ioserver* database configuration parameter defines the number of prefetchers that are available for the database. To get an optimal async read ratio, you should set the *num_ioserver* database configuration parameter to 1-2 more than the number of physical storage devices used by the database to ensure that asynchronous I/O is occurring. This parameter should be changed judiciously since having too many prefecthers invariably results in high system I/O because prefecthers can read many more pages than required into the bufferpools. Async write ratio can be tuned using the *num_iocleaners* database configuration parameter. This parameter specifies the number of asynchronous page cleaners for a database. Increasing the number of page cleaners generally results in a higher async write ratio. The following rules of thumb can be followed when tuning this parameter: *For most transactional systems set this parameter to between 1 and the number of physical storage devices used by the database.* *For workloads that consists of mainly simple reads from the database, set this parameter to 0.* *When you define very large bufferpools on your database you may need to increase the num_iocleaners.* *Monitor the victim page cleaners% metric when tuning the num_iocleaners.*

Space

While DBAs focus on memory settings and tuning SQL, they often forget just how dangerous and insidious storage problems can be. This is not a good mindset because storage headaches can play a major role in wrecking an otherwise well-running database.

Storage problems generally take one of two forms:

- The 'hit-the-wall' variety that can bring things to a complete standstill
- The 'performance vampire' kind that slowly drains the performance of a database over time

Storage problems can quickly lead to performance problems, such as when a transaction log runs out of free space. But storage problems can also silently work behind the scenes to slowly but surely rob a database of its performance.

In addition to a Home page, the Space category of DB2 Performance Analyst includes the following tabbed pages:

Usage tab	Logging tab	Bufferpools tab
Objects tab	Quiesce tab	Rebalance tab
Reorganization tab	Rollforward tab	

Home

The Home page of the Space performance category view displays statistics i the following categories:

Tablespace Analysis pane	Bottleneck Analysis pane	SQL Analysis pane
Storage Analysis pane	Log Analysis pane	

Tablespace Analysis pane

The Tablespace Analysis section includes a list of all of the tablespaces in the database. The following information is presented for each tablespace:

- Tablespace Name: The name of the tablespace
- Type: system managed (SMS) or database managed (DMS) tablespace
- State: This element describes the current state of the tablespace
- Used Space: The space currently in use on tablespace
- Free Space: The space currently free on the tablespace. (This is not applicable to SMS tablespaces)
- Content Type: The type of data that can be stored on the tablespace. It can be one of the following:
  - Any (any type of data)
  - Long (LOB/ Structured type / Index data)
  - System Temporary (work area used by database manager to perform operations like sorts and joins.
  - User Temporary (Stores declared global temporary tables)
- Page Size: Size of pages used by the tablespace
- Extent Size: Number of pages that will be written to a container before writing to the next container on the tablespace
- Prefetch Size: Number of pages that will be read from the tablespace by the prefetchers in a single prefetch
- Total Space: Total space (Used + Free) used by the tablespace on the storage device(s). This will be the same as used space for SMS tables
- High Watermark: Highest amount of used space for the tablespace (This is not applicable to SMS tablespaces)

>

Other statistics appearing in this section include:

Total Tablespaces	The total tablespaces statistic is the total number of tablespaces in the database.
SMS Tablespaces	The SMS tablespaces statistic is the number of system managed (SMS) tablespaces in the database.
DMS Tablespaces	The DMS tablespaces statistic is the number of database managed (DMS) tablespaces in the database.

Bottleneck Analysis pane

Total Tablespaces	SMS Tablespaces	DMS Tablespaces

Space shortfalls or inaccessible tablespaces can quickly bring a database to a halt. The space bottleneck analysis section provides an at a glance look at the items that may require your attention. You can drilldown to the details pages to see further details of the various metrics.

The metrics presented here are:

DMS Tablespaces at Utilization >80%	Inaccessible Tablespaces	

The following statistics, available on this pane, are duplicates of statistics found on the Bottleneck Analysis pane of the Home page:

Inaccessible Containers	Tablespaces in Abnormal State	

DMS Tablespaces at Utilization >80%

The DMS tablespaces at utilization greater than 80% metric gives the number of DMS tablespaces in the database where the used space on the tablespace is more than 80% of the total available space.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Home > Bottleneck Analysis pane
Metrics	This metric is an indicator that some of the tablespaces in your database may be running out of space. You should drill down to the Space metrics to identify the tablespaces that may be at the critical utilization level. In tablespaces where there are little on no inserts after the initial data is loaded, little or no free space may have been left on purpose. In such cases, it is normal to see a high utilization
Correction	If this metric goes beyond the normal operational value for your database, consider creating additional space on the DMS tablespaces that are running out of space. You can do this by performing one or more one of the following tasks: *Clean up unused space by deleting the tables/indexes that are no longer required.* *Resize or extend the existing containers in the tablespaces.* *Add new containers to the tablespaces.*

Inaccessible Tablespaces

The inaccessible tablespaces statistic is the number of tablespaces that are offline.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Home > Bottleneck Analysis pane

Metrics	Sometimes a problem (such as inaccessible containers) will cause a tablespace to go offline. Even when the problem is corrected, such tablespaces are not accessible by any application and need to be brought online explicitly.
Correction	From Tablespace Analysis section, identify the tablespaces whose state is set to Offline Drill down to space details to make sure all of the containers for these tablespaces are accessible. If they are not, first fix the problem with the containers Disconnect all applications from the database (you can use DBArtisan DB Monitor). Reconnect to the database. Use the ALTER TABLESPACE .. SWITCH ONLINE command to bring the tablespace back online. *If the tablespace can be brought up successfully after issuing the command, or if the tablespace was not in the OFFLINE state to begin with, DB2 returns an SQLCODE of 0.* *If the tablespace cannot be brought up successfully because there are still problems with one or more of the containers, DB2 returns an SQLCODE of --293. You can force the database to restart by using the RESTART ... DROP PENDING TABLESPACE, but have to drop any faulty table spaces afterward.*

SQL Analysis pane

The metrics presented here are:

DDL Statements Executed	Insert/Updates/Deletes Statements Executed	Select Statements Executed
Total SQL Executed		

DDL Statements Executed

This metric indicates the number of SQL Data Definition Language (DDL) statements that have been executed.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Home > SQL Analysis pane
Metrics	You can use this element to determine the level of database activity at the application or database level. DDL statements are expensive to run because of their impact on the system catalog tables. As a result, if the value of this element is high, you should determine the cause, and possibly restrict this activity from being performed. DDL statements can also impact: *The catalog cache, by invalidating table descriptor information and authorization information that are stored there and causing additional system overhead to retrieve the information from the system catalogs* *The package cache, by invalidating sections that are stored there and causing additional system overhead due to section recompilation.*

Insert/Updates/Deletes Statements Executed

This is the number of SQL UPDATE, INSERT, and DELETE statements that were executed.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Home > SQL Analysis pane
Metrics	You can use this element to determine the level of database activity at the application or database level. This information can be useful for analyzing application activity and throughput.

Select Statements Executed

This statistic is the number of SQL SELECT statements that were executed.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Home > SQL Analysis pane
Metrics	You can use this element to determine the level of database activity at the application or database level. This information can be useful for analyzing application activity and throughput.

Total SQL Executed

You can use this element to calculate the total number of SQL statements at the database level.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Home > SQL Analysis pane
Metrics	This element can also help you in determine reasons for poor performance. Failed statements mean time wasted by the database manager, and as a result, lower throughput for the database. You can use this element in conjunction with the other SQL metrics to understand the portion of SQL activity that can be attributed to various statement types.

Storage Analysis pane

While DBAs focus on memory settings and tuning SQL, they frequently forget how dangerous and insidious storage problems can be. This is not a good mindset because storage headaches can play a major role in wrecking an otherwise well-running database. Storage problems generally take one of two forms:

- The ‘hit-the-wall’ variety that can bring things to a complete standstill.
- The ‘performance vampire’ kind that slowly drains the performance of a database over time.

Storage problems have the capability to bring the curtain down on a database very quickly, as in the case of a transaction log running out of free space. But storage problems can also silently work behind the scenes to slowly, but surely rob a database of its performance. For example, a hub table in a busy database might be accessed very quickly when an application is first given life, but over time, if it develops a heavy forwarded row problem, it can cause things to run very differently.

The Storage Analysis section of the Performance Analyst Home page displays the total used and free space of all DMS tablespaces in the database. It also provides the total and available log space for the database.

The following ratios succinctly communicate the general overall performance levels of the datasource:

Log Space	Last Backup	

The following statistics, available on this pane, duplicate statistics available on the Storage Analysis pane of the Home page:

DMS Space	Total SMS Space Used	

Log Analysis pane

The Log Analysis section gives you at-a-glance information about all your log configuration parameters and the high watermarks for the log space usage. You can use this information to quickly determine the logging mechanism in place for your database and if the high watermarks are getting too close to the maximum allocations for the log space.

The metrics presented here are:

Log Configuration	High Watermarks	

Log Configuration

The log configuration metric gives a list of space related log configuration parameters for your database. The values of these parameters determine how the logger uses storage space on the database:

- Block on Log Disk Full(*blk_log_dsk_ful*): this configuration parameter can be set to prevent disk full errors from being generated when DB2 cannot create a new log file in the active log path. Instead, DB2 will attempt to create the log file every five minutes until it succeeds. After each attempt, DB2 will write a message to the administration notification log. The only way to confirm that your application is hanging because of a log disk full condition is to monitor the administration notification log. Until the log file is successfully created, any user application that attempts to update table data is not able to commit transactions. Read-only queries may not be directly affected; however, if a query needs to access data that is locked by an update request or a data page that is fixed in the bufferpool by the updating application, read-only queries also appear to hang.

  Setting *blk_log_dsk_ful* to YES causes applications to hang when DB2 encounters a log disk full error. You are then able to resolve the error and the transaction can continue. A disk full situation can be resolved by moving old log files to another file system, or by increasing the size of the file system so that hanging applications can complete.

  If *blk_log_dsk_ful* is set to NO, a transaction that receives a log disk full error fails and is rolled back. In some cases, the database comes down if a transaction causes a log disk full error.

- Log File Size(*logfilsiz*): specifies the size of each configured log, in number of 4-KB pages.

  There is a 256-GB logical limit on the total active log space that you can configure. This limit is the result of the upper limit on logfilsiz, which is 262144, and the upper limit on (*logprimary* + *logsecond*), which is 256.

  The size of the log file has a direct bearing on performance. There is a performance cost for switching from one log to another. So, from a pure performance perspective, the larger the log file, size the better. This parameter also indicates the log file size for archiving. In this case, a larger log file is size it not necessarily better, since a larger log file size may increase the chance of failure or cause a delay in log shipping scenarios. When considering active log space, it may be better to have a larger number of smaller log files. For example, if there are two very large log files and a transaction starts close to the end of one log file, only half of the log space remains available.

  Every time a database is deactivated (all connections to the database are terminated), the log file that is currently being written is truncated. So, if a database is frequently being deactivated, it is better not to choose a large log file size because DB2 will create a large file only to have it truncated. You can use the ACTIVATE DATABASE command to avoid this cost, and having the bufferpool primed will also help with performance.

  Assuming that you have an application that keeps the database open to minimize processing time when opening the database, the log file size should be determined by the amount of time it takes to make offline archived log copies.

  Minimizing log file loss is also an important consideration when setting the log size. Archiving takes an entire log. If you use a single large log, you increase the time between archiving. If the medium containing the log fails, some transaction information will probably be lost. Decreasing the log size increases the frequency of archiving but can reduce the amount of information loss in case of a media failure since the smaller logs before the one lost can be used.

- Log Retain(*logretain*): if *logretain* is set to RECOVERY, archived logs are kept in the database log path directory and the database is considered to be recoverable, meaning that rollforward recovery is enabled.

  NOTE: The default value for the *logretain* database configuration parameter does not support rollforward recovery and must be changed if you are going to use it.

- User Exit(*userexit*): causes the database manager to call a user exit program for archiving and retrieving logs. The log files are archived in a location that is different from the active log path. If userexit is set to ON, rollforward recovery is enabled.

  NOTE: The *userexit* database configuration parameter must be enabled to set logsecond parameter to-1

- Primary Logs(*logprimary*): specifies the number of primary logs of size *logfilsz* that will be created.

  A primary log, whether empty or full, requires the same amount of disk space. Thus, if you configure more logs than you need, you use disk space unnecessarily. If you configure too few logs, you can encounter a log-full condition. As you select the number of logs to configure, you must consider the size to make each log and whether your application can handle a log-full condition. The total log file size limit on active log space is 256 GB.

  If you are enabling an existing database for rollforward recovery, change the number of primary logs to the sum of the number of primary and secondary logs, plus 1. Additional information is logged for LONG VARCHAR and LOB fields in a database enabled for rollforward recovery.

- Secondary Logs(*logsecond*): specifies the number of secondary log files that are created and used for recovery, if needed.

  If the primary log files become full, secondary log files (of size *logfilsiz*) are allocated, one at a time as needed, up to the maximum specified by this parameter. If this parameter is set to -1, the database is configured with infinite active log space. There is no limit on the size or number of in-flight transactions running on the database.

  NOTE: If this parameter is set to -1, crash recovery time may be increased since DB2 may need to retrieve archived log files.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Home > Log Analysis pane

High Watermarks

The high watermarks statistic shows the maximum total and secondary log space used by the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Home > Log Analysis pane
Metrics	The *total log space used high watermark* in conjunction with the total log space allocated, helps you determine if applications have come close to using the total log space available to your database. The *total secondary log space used high watermark* in conjunction with total log space used high watermark helps you determine your current dependency on secondary logs.
Correction	High values of either high watermarks can indicate that you either need to increase the log space on the database or that the applications using the database need to perform more frequent COMMITS. You can increase the overall log space available by adjusting one or more of the following database configuration parameters: **logfilsiz**, **logprimary**, **logsecond**, and **logretain**.

Usage tab

The Space Detail Usage tab section shows you the list of tablespaces for the database. The following information is presented for each tablespace:

Tablespace Details	Container Details	

Tablespace Details

ID: Internal DB2 Tablespace identifier.

Name: Tablespace Name

Type: Tablespace Type – SMS (System Managed Space) or DMS (Database Managed Space)

Content Type: Shows the type of data that can be stored on the tablespace. The valid types are:

- *Any* meaning can store any type of data

- *Sys Temp* meaning, used internally by DB2 to store intermediate result data
- *User Temp* meaning, can be used to create temporary tables
- *Long* meaning, can be used to store LOB/LONG/Index data

State: Can be one or more of the following:

- Normal (see the definition SQLB_NORMAL in sqlutil.h)
- Quiesced: SHARE
- Quiesced: UPDATE
- Quiesced: EXCLUSIVE
- Load pending
- Delete pending
- Backup pending
- Roll forward in progress
- Roll forward pending
- Restore pending
- Recovery pending (not used)
- Disable pending
- Reorg in progress
- Backup in progress
- Storage must be defined
- Restore in progress
- Offline and not accessible
- Drop pending
- Storage may be defined
- Storage Definition is in 'final' state
- Storage Definition was changed prior to rollforward
- DMS rebalancer is active
- TBS deletion in progress
- TBS creation in progress

Page Size: The size of tablespace pages

Extent Size: Number of pages that will be written to a container before writing to the next container on the tablespace

Prefetch Size: The maximum number of pages the prefetcher gets from the disk at a time

Usable Space: The total space in a table space minus overhead space. (Note: For SMS tablespaces, this element will have the same value as Used Space).

Used Space: The total space in the tablespace that is currently used in the tablespace (including overhead)

>

Free Space: The total space that is free to be used for storage of data. This element is not applicable for SMS tablespace. free space for an SMS tablespace is the OS file system free space for the SMS containers.

Total Space: The total space used by the tablespace

Pending Free Pages: The number of pages in the tablespace that would become free if all pending transactions are committed or rolled back and new space is requested for an object. (Applies only to DMS tablespaces)

 # Containers: The number of containers in the tablespace

Usable Pages: The total pages in a table space minus overhead pages. (Note: For SMS tablespaces, this element will have the same value as Used Pages).

Used Pages: The total pages in the tablespace that is currently used in the tablespace (including overhead)

Free Pages: The total pages that are free to be used for storage of data. This element is not applicable for SMS tablespace. Free pages for an SMS tablespace are dependent on the OS file system free space for the SMS containers.

Total Pages: The total pages used by the tablespace

Page HWM: For DMS, this element represents the page number of the first free extent following the last allocated extent of a table space. Note that this is not really a "high water mark", but rather a "current water mark", since the value can decrease. For SMS, this is not applicable.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Usage tab IBM DB2 (W/U/L) Performance Analyst statistics > Space > Rebalance tab

Container Details

This section shows the list of containers for the selected tablespace. The information presented here includes:

ID: A value that uniquely identifies the container within the tablespace

Name: The name of the container. Typically this is the full path of the file/directory/device where the container exists.

Type: The type of container. An SMS Container will be a directory. A DMS Containers will be a raw device/file/stripped disk/ or stripped file. Together with Container name, and partition, this metric identifies the physical location of the container.

DB Partition: The database partition where the container exists

Usage Pages: Usable pages in the container (Applicable to DMS Tablespaces only).

Total Pages: Total pages in the container

Stripe Set: The stripe set that a container belongs to.

Accessible: This element describes if a container is accessible or not

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Usage tab
Metrics	This section can be used to get a better understanding of the physical layout of a tablespace. Also, it lets you quickly identify if any containers are inaccessible.

Logging tab

The Logging tab of the Space Detail view displays Logging Space statistics:

Log Read/Writes	Indoubt Transactions	Log Usage
Secondary Logs		

Log Read/Writes

This section shows the level of read/write activity being performed by the logger. The information presented here is:

Log Reads: The number of log pages read from the disk by the logger per second.

Log Writes: The number of log pages written to the disk by the logger.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Logging tab
Metrics	You can use this page with the operating system level I/O metrics, to quantify the amount of I/O on a device that is attributable to database activity.

Indoubt Transactions

This section shows the number of outstanding indoubt transactions in the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Logging tab
`Metrics`	Indoubt transactions hold log space for uncommitted transactions, which can cause the logs to become full. When the logs are full, further transactions cannot be completed. The resolution of this problem involves a manual process of heuristically resolving the indoubt transactions. This monitor element provides a count of the number of currently outstanding indoubt transactions that must be heuristically resolved.

Log Usage

This section shows the following log usage details:

Total Used Space: Total space used by log files

Total Free Space: Total space still available to the logger

Total Used Space Highwatermark: The maximum amount of total log space used.

Secondary Space Used High Watermark: The maximum amount of secondary log space used.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Logging tab
Metrics	Use these metrics to check your current log space usage and the highest amount of space used by the logger. If the usage is high and the high watermarks are also near the total available space, you may need to increase the amount of log space or end/commit/rollback the oldest transaction.
Correction	You can use the following database configuration parameters to change the amount of log space available: **logfilsiz**, **logprimary**, and **logsecond.** You can end (force) a connection with a long running transaction using the DBArtisan database monitor.

>

Secondary Logs

This section shows the secondary logs allocations by the database over the monitoring period.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Logging tab
Metrics	When the primary log files become full, the secondary log files are allocated one at a time as needed, up to a maximum number as controlled by the logsecond database configuration parameter. Secondary log files are useful in preventing occasional log fill ups but they may not be as efficient as primary log and they also increase the recovery time. If you see a constant reliance on secondary log files, it may indicate infrequent commits by database applications or insufficient primary log space.
Correction	First examine the database applications to see if secondary log files are due to long running transactions and whether these transactions can be shortened. If the transactions cannot be shortened or you still see a very frequent use of secondary logs after transaction tuning, increase the primary log size (logprimary,logfilsiz Database configuration parameters) to reduce the dependence on secondary logs. If there is an occasional long running transaction, and you see your see transaction log full (SQL0964C) error messages, you can either increase the number of secondary log files (logsecond database configuration parameter) or set the number to -1 (no limit on number of secondary log files).

Bufferpools tab

The Bufferpools tab of the Space Detail View shows the list of tablespaces, with the current bufferpool associated with each tablespace and also the bufferpool used by the database at next startup

Objects tab

The Objects tab of the Space Detail View shows the following statistics:

Tablespace Object Summary	Tablespace Object Details	

Tablespace Object Summary

This section shows information about how tablespaces are being used on the database. The information presented here includes:

Tablespace: Name of the tablespace

Table Count: Number of tables on the tablespace

Page Size: Tablespace page size

Index Count: Number of indexes on the tablespace

Table Space: Space used by tables on the tablespace

Index Space: Space used by indexes on the tablespace

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Bufferpools tab

Tablespace Object Details

This section shows the details of the objects on the tablespaces shown in the Tablespace Object summary section. The user can toggle between the list of tables and indexes for the tablespace.

The information presented here for tables includes:

>

Table Schema: Schema name of the table

Table Name: Name of the tables

Avg Row Length: Approximate length of the table row.

 # Rows: Table Cardinality

#Pages Allocated: Number of pages allocated to the table

Table Size: Total space used by the rows on the table

Available: Total available space on the table

Total: Total space used by the table pages

NOTE: The information presented here is based on the table/index statistics and if the statistics haven't been collected /updated, this information may be out of date.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Bufferpools tab

Quiesce tab

The Quiesce tab of the Space Detail view shows displays the following information:

Quiesced Tablespaces	Quiescer Details	

Quiesced Tablespaces

This section shows a list of quiesced tablespaces. A quiesce is a persistent lock. Its benefit is that it persists across transaction failures, connection failures, and even across system failures (such as power failure, or reboot).

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Quiesce tab
Metrics	A quiesce is owned by a connection. If the connection is lost, the quiesce remains, but it has no owner, and is called a phantom quiesce. For example, if a power outage caused a load operation to be interrupted during the delete phase, the tablespaces for the loaded table would be left in delete pending, quiesce exclusive state. With a database restart, this quiesce would be an unowned (or phantom) quiesce. The removal of a phantom quiesce requires a connection with the same user ID used when the quiesce mode was set. This section is useful in identifying the tablespaces that are in a quiesced state.
Correction	To remove a phantom quiesce: *Connect to the database from DBArtisan with the same user ID used when the quiesce mode was set.* *Re-quiesce the table space using the current quiesce state from DBArtisan.* *Reset the quicesce from DBartisan.*

Quiescer Details

This section shows the list of quiescers for the selected quiesced table.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Quiesce tab
Metrics	Any quiescers whose agent ids are not listed under Users Detail >Attributes are phantom.

Correction	To remove a phantom quiesce: *Connect to the database from DBArtisan with the same user ID used when the quiesce mode was set.* *Re-quiesce the table space using the current quiesce state from DBArtisan.* *Reset the quicesce from DBartisan.*

Rebalance tab

The Rebalance tab of the Space Detail page displays the following information:

- Rebalance Status

The following statistic, available on this pane, are duplicates of statistics on the Space Usage tab.

Tablespace Details

Rebalance Status

This section shows a graph of the percentage of extents that have been processed and the percentage that remains incomplete for the selected tablespace.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Rebalance tab

Reorganization tab

The Reorganization tab of the Space Details View displays the following statistics:

Table Details	Reorganization Status	Reorganization Attributes

Table Details

This section shows the details of all the tables on which reorganization has been attempted since the start of monitoring. It also shows the status of tables that are presently undergoing reorganization. The information presented in this includes:

Table Schema: Table Schema of the table on which reorganization is being/was performed

Table Name: Table Name of the table on which reorganization is being/was performed

Tablespace: Tablespace of the table on which reorganization is being/was performed

Start: Start time of the reorganization

End: End Time of the reorganization

Phase Start: Start time of a phase of reorganization

Current Counter: Percentage of reorganization operation completed (0 if the reorganization status is complete)

Max Phase: The maximum number of reorganization phases that will occur during reorganization processing. This applies to classic (offline) reorganizations only.

Status: The status of an in-place (online) table reorganization. This is not applicable to classic (offline) table reorganizations. An in-place table reorganization can be in one of the following states:

- Started/Resumed

>

- Paused
- Stopped
- Completed
- Truncate

Completion: Table reorganize success indicator.

Phase: Table reorganize phase. This applies to classic (offline) table reorganizations only).

For classic table reorganization, the following phases are possible:

- Sort
- Build
- Replace
- Index Recreate

% Complete: Shows the completion status of the reorganization jobs

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Reorganization tab

Reorganization Status

Shows a graphical status view of the currently selected reorganization.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Reorganization tab

Reorganization Attributes

This section shows the reorganization attributes of the currently selected reorganization. The attributes include:

- Allow Write Access
- Allow Read Access
- Allow No Access
- Recluster Via Index Scan
- Reorg Long Field LOB Data
- No Table Truncation
- Reclustering
- Reclaiming
- Inplace Table Reorg
- Table Reorg
- Recluster Via Table Scan
- Reorg Data Only.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Space > Reorganization tab

Rollforward tab

Recovering database changes is time-consuming. You can use the database system monitor to monitor the progression of a recovery. This section provides the following elements to help you monitor the status of rollforwards:

Timestamp: Timestamp of the log being processed. This is an indicator of the data changes that will be recovered

Tablespace: Tablespace being rolled forward. This element identifies the tablespaces involved in the rollforward

Type: An indicator of whether the recovery is happening at a database or tablespace level

Log#: The log being currently processed. This identifies the log involved in the rollforward.

Status: The status of rollforward recovery. This metric indicates if the recovery is in an undo (rollback) or redo (rollforward) phase.

#Tablespaces: Number of tablespaces involved in the rollforward

Node: Indicates the database partition number for the rollforward operation

Users

The database activity can be examined using key performance metrics and wait/bottleneck-based analysis. In addition to a Home page, the Users category of DB2 Performance Analyst includes the following tabbed pages:

Top Sessions tab	Cursors tab	SQL Activity tab
Attributes tab	Lock Waits tab	Locks tab

Home

The Home page of the Users performance category view displays the vital DB2 user statistics in the following categories:

Key User Analysis pane	Bottleneck Analysis pane	SQL Analysis pane
Connection Analysis pane		

Key User Analysis pane

Database performance analysts typically use one of two methods for examining the performance levels of a database - ratio-based or wait/bottleneck-based. Ratio-based analysis involves examining a number of key database ratios that can be used to indicate how well a database is running. Performance ratios serve as very good roll-up mechanisms for busy DBAs to use for at-a-glance performance analysis. Many DBAs have large database farms to contend with and cannot spend time checking detailed wait-based analysis outputs for each and every database they oversee. Succinctly presented performance ratios can assist in such situations by giving DBAs a few solid indicators that can be quickly scanned to see if any database needs immediate attention.

While there are certainly many opinions as to what rules to follow, there are some standards that should always be adhered to. To start with, many of the formulas that make up ratio-based analysis must be derived from delta measurements instead of cumulative statistics. Many of the global ratios that a DBA will examine come from the DB2 Snapshot Monitor. The snapshot monitor returns metrics for the Instance and particular databases since the start of monitoring. Some metrics are cumulative (e.g., counters) and others are instantaneous (e.g., gauges).

>

A final thing to remember about using ratio-based analysis is that, while there are a number of rules of thumb that can be used as starting points to begin the evaluation of database performance, DBAs must determine each SQL Server's individual 'personality' with respect to the key performance ratios. Some hard and fast rules simply do not apply to every database. The danger in using blanket ratio standards is that they can lead the DBA to haphazardly take action, which can at times contribute nothing to the situation, and sometimes even degrade performance.

The following ratios are used on the Performance Analyst User page to succinctly communicate the general overall performance levels of the monitored database:

Efficiency

Transactions/sec	Static SQL/sec	Dynamic SQL/sec

Activity

Location	Applications Executing	Connections High Watermark

Transactions/sec

Transactions/sec is the number of transactions (units of work) completed per second on the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Key User Analysis pane
Metrics	A small rate of transactional activity on the database can indicate that applications are not doing frequent commits, which may lead to logging and concurrency problems.
Correction	Drill down to the Users Detail>SQL Activity page to check which applications are running their transactions for long periods of time

Static SQL/sec

The static SQL/sec metric is the number of static SQL statement executions attempt on the database each second.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Key User Analysis pane

Dynamic SQL/sec

This metric is the number of dynamic SQL statement executions being attempted on the database each second.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Key User Analysis pane

Applications Idle

The idle applications metric indicates the number of applications that are currently connected to the database for which the database manager is not executing any requests

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Key User Analysis pane
Metrics	This can help you understand the level of activity within a database and the amount of system resource being used.

Applications Executing

The executing applications statistic indicates the number of applications for which the database manager is currently executing requests.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Key User Analysis pane
Metrics	You can use this element to understand how many of the database manager agent tokens are being used by applications connected to this database.

Connections High Watermark

The connections high watermark is the highest number of simultaneous connections to the database since the database was activated.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Key User Analysis pane
Metrics	You can use this element to evaluate the setting of the *maxappls* configuration parameter. If the value of this element is the same as the *maxappls* parameter, it is likely that some database connection requests were rejected, since *maxappls* limits the number of database connections allowed.
Correction	Increase the value of the *maxappls* parameter to allow more connections.

Bottleneck Analysis pane

When a DB2 database is active, the applications accessing the database can either successfully access the database resources, or they can wait for resources to become available. An application that is waiting can indicate normal database activity, or it can be an indicator that a database bottleneck exists. This is where wait-based or bottleneck analysis comes into play. A database administrator can use this form of performance analysis to determine if perceived bottlenecks in a database are contributing to a performance problem.

Performance Analyst Bottleneck analysis section helps a DBA isolate potential problem areas where there are waits, resource overloads or potential failures. If concurrency issues or heavy table scan activity has been dragging a database's performance down, a DBA can use bottleneck analysis to isolate the root cause. Once one or more bottlenecks have been isolated as possible areas of interest, the DBA can drill down and examine details on why the applications or resources are causing the problem. The Performance Analyst Users page identifies the top system and application bottlenecks that may be currently a cause of lessened performance.

The following bottleneck indicators are used on the Performance Analyst Home page to succinctly communicate the general overall performance level of the monitored database:

Top Bottlenecked User Processes	Automatic Rebinds	Internal Rollbacks
Waits	Deadlocks	Timeouts

The following statistics, appearing on this page, duplicate statistics available on the Bottleneck Analysis pane of the Home page:

Lock Escalations	Sort Overflows	Failed SQL

The following statistic, available on this pane, duplicates a statistic on the SQL Analysis pane of the Home page:

Failed SQL

>

Top Bottlenecked User Processes

The top bottlenecked user processes grid shows the top bottlenecked applications in the system. These applications are either performing the top I/O operations or have idled the longest to obtain locks.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Bottleneck Analysis pane

Automatic Rebinds

The automatic rebinds statistic is the number of automatic rebinds (or recompiles) that have been attempted.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	Automatic rebinds are the internal binds the system performs when an package has been invalidated. The rebind is performed the first time the database manager needs to execute an SQL statement from the package. For example, packages are invalidated when you: *Drop an object, such as a table, view, or index, on which the plan depends* *Add or drop a foreign key* *Revoke object privileges on which the plan depends.* Since automatic rebinds can have a significant impact on performance, they should be minimized where possible.

Internal Rollbacks

The internal rollbacks statistic is the total number of rollbacks initiated internally by the database manager.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	An internal rollback occurs when any of the following cannot complete successfully: A reorga*nization* *An import* *A bind or pre-compile* *An application ends as a result of a deadlock situation or lock timeout situation* *An applicat*ion ends without executing an explicit commit or rollback statement (on Windows).

Waits

The lock waits metric is the total number of times that applications or connections waited for locks within the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Bottleneck Analysis pane

Deadlocks

The deadlocks statistic is the total number of deadlocks that have occurred since this instance of Performance Analyst started monitoring the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Bottleneck Analysis pane

Metrics	If a large number of deadlocks are detected, it can indicate that applications are experiencing lock contention problems. Deadlocks are usually caused by one of the following situations: *Lock escalations on the database* *Catalog tables locked for Repeatable Read* *Applications are using inappropriate isolation levels at bind time* *Applications are obtaining the same locks in a different order* *Applications are locking tables explicitly where row level locks are sufficient.*
Correction	You may be able to modify the applications causing lock contentions for better concurrency. To identify the applications that may be causing contentions, go to the Lock Waits tab of the Users detail page and review the Application Details section.

Timeouts

The lock timeouts metric identifies the number of times that a request to lock an object timed out without being granted.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Bottleneck Analysis pane
Metrics	If the number of lock timeouts becomes excessive when compared to the acceptable range for your database, it can indicate that an application is holding locks for long durations. It can also indicate that the amount of time an application waits for a lock before timing out is too short. If you have too few lock timeouts and the average lock wait time is too high, it can indicate that the lock timeout configuration parameter is set to an excessively high value
Correction	First you should examine the lock activity at the application level (Users>Lock, Lock Waits drilldown) to identify any particular application that is causing excessive lock contentions. If so, you can tune the application to provide better concurrency. If lock timeouts are excessive, and average lock wait times are very short, you can increase the *locktimeout* database configuration parameter to make the applications wait longer before timing out.

SQL Analysis pane

Most of database's I/O performance can be attributed to SQL statement execution. Poorly optimized SQL statements can drag down an otherwise well-configured server in terms of user response times.

This section shows a snapshot of row level SQL operations that have been performed on the database since you started monitoring. This gives you an insight into the row level of read/write activity that is currently occurring within the database and comparing this to the prefetcher, and cleaner activity allows you to relate the two.

The following statistics are used on the Performance Analyst for DB2 User home page to succinctly communicate the general overall performance levels of I/O. They duplicate statistics appearing on the SQL Analysis pane of the I/O Home page.:

Rows Deleted	Rows Inserted	Rows Selected
Rows Updated	Rows Read	

Connection Analysis pane

The Connection Analysis section displays the level of application connect/disconnect activity occurring on the database.

>

The following statistics are used on the Performance Analyst for DB2 User home page to succinctly communicate general overall connectivity:

Current Connections	Disconnects	Connections Since Database Activation
High Watermark for Concurrent Connections		

Current Connections

The current connections metric is the number of applications currently connected to the database.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Connection Analysis pane
Metrics	You can use this metric to help you get an overview of the level of database activity and the amount of system resources in use.
Correction	This metric can help you adjust the setting of the *maxappls* and *max_coordagents* configuration parameters. For example, if this value is always the same as *maxappls*, you consider increasing the value of *maxappls*.

Disconnects

This metric shows the number of disconnects from the database since the first connect, activate, or last reset.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Connection Analysis pane

Connections Since Database Activation

The connections since database activation statistic indicates the number of connections to the database since the first connect, activate, or last reset.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Connection Analysis pane
Metrics	Use this metric in conjunction with the database connection time and the database start time to see if the application connection frequency is low. When there are no connections to db2, by default the database is automatically deactivated until the next connection. The connection that reactivates the database encounters slow connection time because it needs to wait for the database to reinitialize. If very few applications are connecting to the database, some connections can be slow because of this
Correction	If there are short periods when your database has no applications connected to it, activate the database explicitly using the ACTIVATE DATABASE command. This significantly improves the connection performance for the database.

High Watermark for Concurrent Connections

The connections high watermark is the highest number of simultaneous connections to the database since the database was activated.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Home > Connection Analysis pane
Metrics	You can use this element to evaluate the setting of the *maxappls* configuration parameter. If the value of this element is the same as the *maxappls* parameter, it is likely that some database connection requests were rejected, since *maxappls* limits the number of database connections allowed.
Correction	Increase the value of the *maxappls* parameter to allow more connections.

>

Workload Analysis pane

When your phone starts ringing with complaints of performance slowdowns, one of the first things you should get a handle on is:

- Who is currently connected to the database?
- What resources are they using?
- What are they currently executing?

The Workload Analysis section of the Performance Analyst Users page provides insight into the leading resource hogs of a server. Drill-downs are available so you can easily get detailed information into what each leading application is currently involved with. The Workload Analysis section presents the following metrics:

Top Lock Escalations	The top lock escalation displays the application with largest number of lock escalations.
Top CPU Process	The top CPU process displays the application with the highest CPU usage.
Top I/O Process	The top I/O process displays the application with the highest I/O read/write time.
Oldest Transaction	The oldest transaction displays the application with the longest running current transaction.

Top Sessions tab

Very often, one renegade client application brings a healthy database down by hogging most of the available resources. This page shows the top resource hogs in each of the resource usage areas to help you quickly identify the worst offenders in each area. You see:

Top I/O Process	Top Memory Process	Top CPU Process
Top Sessions		

Top I/O Process

This section shows the application with the top IO time and compares it with the IO time of all other applications.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Top Sessions tab
Metrics	The metrics are calculated by measuring the total time spent by each application doing reads and writes (buffered and non-buffered).

Top Memory Process

This section shows the application with the largest memory pool usage aggregate. It also shows the total memory pool usage for all other applications.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Top Sessions tab
Metrics	This metric is calculated by aggregating all of the memory pools for each application and showing the top application. This section is useful in determining if the top application is using a lot more resources

>

Top CPU Process

This section shows the application with the top CPU time usage and compares it with the CPU usage of all other applications.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Top Sessions tab
Metrics	Metrics are calculated using the total user and system CPU time for each application connected to the database.

Top Sessions

Shows a list of all applications connected to the database and the resource usage for these applications. The information presented for each application includes:

Auth ID: The authorization ID of the user who invoked the application that is being monitored

Name: The name of the application running at the client as known to the database manager

Handle: A system-wide **unique** ID for the application. On a single-partitioned database, it consists of a 16 bit counter. On a multi-partitioned database, it consists of the coordinating partition number concatenated with a 16 bit counter. In addition, this identifier will be the same on every partition where the application may make a secondary connection.

ID: This identifier is generated when the application connects to the database at the database manager

Client PID: The process ID of the client application that made the connection to the database

I/O Time: The total time spent by application in performing buffered and non-buffered reads and writes

Memory Usage: Total memory pool usage for the application

CPU Time: The total user + system cpu time used by the application agents

 # Agents: Number of agents for the application

Status: Current Status of the application. Values for this field are:

- **Database Connect Pending:** The application has initiated a database connection but the request has not yet completed.
- **Database Connect Completed:** The application has initiated a database connection and the request has completed.
- **Unit of Work Executing**: The database manager is executing requests on behalf of the unit of work.
- **Unit of Work Waiting:** The database manager is waiting on behalf of the unit of work in the application. This status typically means that the system is executing in the application's code.
- **Lock Wait:** The unit of work is waiting for a lock. After the lock is granted, the status is restored to its previous value.
- **Commit Active**: The unit of work is committing its database changes.
- **Rollback Active**: The unit of work is rolling back its database changes.
- **Recompiling:** The database manager is recompiling (that is, rebinding) a plan on behalf of the application.
- **Compiling**: The database manager is compiling an SQL statement or precompiling a plan on behalf of the application.
- **Request Interrupted**: An interrupt of a request is in progress.

>

- **Database Disconnect Pending**: The application has initiated a database disconnect but the command has not yet completed executing. The application may not have explicitly executed the database disconnect command. The database manager will disconnect from a database if the application ends without disconnecting.
- **Decoupled from Agent:** The application has been decoupled from an agent.
- **Transaction Prepared:** The unit of work is part of a global transaction that has entered the prepared phase of the two-phase commit protocol.
- **Transaction Heuristically Committed**: The unit of work is part of a global transaction that has been heuristically committed.
- **Transaction Heuristically Rolled Back**: The unit of work is part of a global transaction that has been heuristically rolled-back.
- **Transaction Ended**: The unit of work is part of a global transaction that has ended but has not yet entered the prepared phase of the two-phase commit protocol.
- **Creating Database**: The agent has initiated a request to create a database and that request has not yet completed.
- **Restarting Database**: The application is restarting a database in order to perform crash recovery.
- **Restoring Database:** The application is restoring a backup image to the database.
- **Backing Up Database:** The application is performing a backup of the database.
- **Data Fast Load:** The application is performing a "fast load" of data into the database.
- **Data Fast Unload:** The application is performing a "fast unload" of data from the database.
- **Wait to Disable Tablespace**: The application has detected an I/O error and is attempting to disable a particular table space. The application has to wait for all other active transactions on the table space to complete before it can disable the table space.
- **Quiescing a Tablespace:** The application is performing a quiesce table space request.
- **Wait for Remote Partition:** The application is waiting for a response from a remote partition in a partitioned database instance.
- **Remote Request Pending**: The application is waiting for results from a federated datasource.

Last UOW Start Time: The date and time that the unit of work first required database resource

Last UOW End Time: The date and time that the most recent unit of work completed, which occurs when database changes are committed or rolled back

Current UOW Elapsed Time: The elapsed execution time of the most recently completed unit of work.

Login Timestamp: The date and time that a connection request was granted.

Inbound Comm Address: This is the communication address of the client.

Client Protocol: The communication protocol that the client application is using to communicate with the server.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Top Sessions tab

Cursors tab

The Cursors tab of the Users Detail view displays the Application SQL Cursor Activity. It details the SQL cursors open for each application. The information presented here includes:

Open Cursors: The number of cursors (local and remote) currently open for an application.

>

Blocking Cursors: The number of blocking cursors (local and remote) currently open for the application

Block Requests Accepted: The percentage of all blocking I/O requests that were accepted by the database

Block Requests Rejected: The percentage of all blocking I/O requests that were rejected by the database

Metrics

The number of blocking cursors for an application should be close to the number of open cursors. If this is not the case the application performance may be affected. If there are many cursors blocking data, the communication heap may become full. When this happens, instead of returning an error, the database stops allocating I/O blocks for blocking cursors.

Correction

If the number of blocking cursors in an application is not close to the number of open cursors, you may be able to improve application performance by tweaking the precompile options for record blocking for the treatment of ambiguous cursors, and by redefining the cursors with FOR FETCH ONLY clause where possible

If a large number of cursors were unable to perform blocking, you may be able to improve performance by increasing the *query_heap* database manager configuration parameter.

SQL Activity tab

The SQL Activity tab of the Users Detail view displays the following statistics:

Application List	SQL Executed	SQL Activity
UOW Details	Binds and Precompiles Executed	

Application List

This section shows a list of applications connected to the database. You can select an application from this list and see details for the SQL activity for this application.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > SQL Activity tab

SQL Executed

This section shows the details for the various types of SQL statements executed by the particular application. It also gives you the total application throughput. The information presented here includes:

Selects: Percentage of statements that are select statements

Updates/Inserts/Deletes: Percentage of statements that modify/write data on the database

DDLs: Percentage of CREATE/DROP/ALTER statements

Others: All other types of statements (Fetch, Open cursor etc)

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > SQL Activity tab
Metrics	Various types of statements have varying impact on the performance and concurrency of the database. For example - DDL statements can cause catalog and package cache invalidations, Insert/update/delete statements usually obtain locks and perform log IO. This section can help you determine the contributing statement types to the activity for an application.

>

SQL Activity

This section gives details on the sql activity for the selected application. It allows you to quick compare the number of static dynamic and failed sql statements with the overall application throughput. The metrics presented here include:

Static SQL: The number of static SQL statements that the application tried to execute

Dynamic SQL: The number of dynamic SQL statements that the application tried to execute

Failed SQL: The number of SQL statements (static and dynamic) that the application tried to execute but failed

SQL Throughput: Number of SQL statements (dynamic and static) that the application executed successfully

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > SQL Activity tab
Metrics	These elements may help you determine the nature of the application. Failed SQL statements metric may also help identify the cause of poor performance on the database since failed statements mean time wasted by the database manager.

UOW Details

Shows the units of work being completed by the application over the refresh time interval.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > SQL Activity tab
Metrics	A low rate of unit of work completion may be an indicator of one of the following: *The application is not committing transactions frequently and may be causing concurrency issues for other applications* *The application is idling and not performing any work on the server*
Correction	Examine the locks obtained by the application (Users Detail >Locks page) to see if a low number here should be a cause for concern.

Binds and Precompiles Executed

This section shows the total number of implicit and explicit binds and precompiles executed by the application.

Implicit rebinds are performed by the system when the application tries to execute a sql statement from a package and the package has been invalidated. A package is usually invalidated when you perform DDL/DCL operations on objects on which the package plan is dependent. Explicit rebinds occur when the user executes the REBIND PACKAGE command.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > SQL Activity tab
Metrics	Bind and precompile operations have significant overhead for the database operations and they should be minimized during periods of heavy usage when possible.

Attributes tab

The Attributes tab of the Users Detail page displays the following statistics:

Application List	Application Attributes	Direct Authorities
Indirect Authorities		

>

Application List

This section shows the list of applications connected to the database. The user can obtain further details for an application from other sections on this page by selecting the application from this list.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Attributes tab

Application Attributes

Shows the main information for the selected application. This information is useful in determining the applications connection attributes and its state.The information shown here includes:

Application Status: The status of the application. Values for this field are:

- **Database Connect Pending:** The application has initiated a database connection but the request has not yet completed.
- **Database Connect Completed:** The application has initiated a database connection and the request has completed.
- **Unit of Work Executing**: The database manager is executing requests on behalf of the unit of work.
- **Unit of Work Waiting:** The database manager is waiting on behalf of the unit of work in the application. This status typically means that the system is executing in the application's code.
- **Lock Wait:** The unit of work is waiting for a lock. After the lock is granted, the status is restored to its previous value.
- **Commit Active**: The unit of work is committing its database changes.
- **Rollback Active**: The unit of work is rolling back its database changes.
- **Recompiling:** The database manager is recompiling (that is, rebinding) a plan on behalf of the application.
- **Compiling**: The database manager is compiling an SQL statement or precompiling a plan on behalf of the application.
- **Request Interrupted**: An interrupt of a request is in progress.
- **Database Disconnect Pending**: The application has initiated a database disconnect but the command has not yet completed executing. The application may not have explicitly executed the database disconnect command. The database manager will disconnect from a database if the application ends without disconnecting.
- **Decoupled from Agent:** The application has been decoupled from an agent.
- **Transaction Prepared:** The unit of work is part of a global transaction that has entered the prepared phase of the two-phase commit protocol.
- **Transaction Heuristically Committed**: The unit of work is part of a global transaction that has been heuristically committed.
- **Transaction Heuristically Rolled Back**: The unit of work is part of a global transaction that has been heuristically rolled-back.
- **Transaction Ended**: The unit of work is part of a global transaction that has ended but has not yet entered the prepared phase of the two-phase commit protocol.
- **Creating Database**: The agent has initiated a request to create a database and that request has not yet completed.
- **Restarting Database**: The application is restarting a database in order to perform crash recovery.
- **Restoring Database:** The application is restoring a backup image to the database.

- **Backing Up Database:** The application is performing a backup of the database.
- **Data Fast Load:** The application is performing a "fast load" of data into the database.
- **Data Fast Unload:** The application is performing a "fast unload" of data from the database.
- **Wait to Disable Tablespace**: The application has detected an I/O error and is attempting to disable a particular table space. The application has to wait for all other active transactions on the table space to complete before it can disable the table space.
- **Quiescing a Tablespace:** The application is performing a quiesce table space request.
- **Wait for Remote Partition:** The application is waiting for a response from a remote partition in a partitioned database instance.
- **Remote Request Pending**: The application is waiting for results from a federated datasource.

ID of code page used by the application: The code page identifier.

Application status change time: The date and time the application entered its current status.

Sequence Number: This is incremented whenever a unit of work ends (that is, when a COMMIT or ROLLBACK terminates a unit of work). Together, the application id and sequence number uniquely identify a transaction.

Client Product/Version ID: The product and version of DB2 that is running on the client.

Database Alias used by the Application: The alias of the database provided by the application to connect to the database.

User login id: The ID that the user specified when logging in to the operating system. This ID is distinct from auth_id, which the user specifies when connecting to the database

DRDA correlation token: The DRDA correlation token is used for correlating the processing between the application server and the application requester. It is an identifier dumped into logs when errors arise, that you can use to identify the conversation that is in error. In some cases, it will be the LUWID of the conversation

Client operating platform: The operating system on which the client application is running.

Database territory code: The territory code of the database for which the monitor data is collected

Application agent priority: The priority of the agents working for this application

Application priority type: Operating system priority type for the agent working on behalf of the application

Coordinating node: In a multinode system, the node number of the node where the application connected or attached to the instance

Connection request start timestamp: The date and time that an application started a connection request.

Connection request completion timestamp: The date and time that a connection request was granted.

Previous unit of work completion timestamp: This is the time the unit of work completed.

Unit of work start timestamp: The date and time that the unit of work first required database resources.

Unit of work stop timestamp: The date and time that the most recent unit of work completed, which occurs when database changes are committed or rolled back.

Most recent unit of work elapsed time: The elapsed execution time of the most recently completed unit of work.

Unit of work completion status: The status of the unit of work and how it stopped.

>

Application idle time: Number of seconds since an application has issued any requests to the server. This includes applications that have not terminated a transaction, for example not issued a commit or rollback.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Attributes tab

Direct Authorities

Shows the highest level of direct authority granted to the application

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Attributes tab

Indirect Authorities

Shows the highest level of indirect authority granted to the application

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Attributes tab

Lock Waits tab

The Lock Waits tab of the Users Detail view shows the details for all applications in the database that are waiting on locks. This information is useful in tracking concurrency issues in the database. The statistics available here are:

Application Details	Applications Waiting on Locks	

The following statistic, available on this tab, duplicates a statistic on the Bottleneck Analysis pane of the Home page:

Avg. Lock Wait Time

Application Details

This section shows the list of applications waiting for locks to be released on objects they need to access. The information presented here includes:

Auth ID: The authorization ID of the user who invoked the application that is being monitored.

Name: Name of the application executable

Handle: A systemwide **unique** ID for the application

ID: This identifier is generated when the application connects to the database at the database manager

Client PID: The process ID of the client application that made the connection to the database

Total Lock Waits: The total number of times that applications or connections waited for locks

Total Lock Waiting Time: This is the total amount of elapsed time that this connection or transaction has waited for a lock to be granted to it.

Locks Waiting: Indicates the number of agents waiting on a lock

UOW Lock Waiting Time: The total amount of elapsed time current unit of work has spent waiting for locks

Wait Start Time: The date and time that this application started waiting to obtain a lock on the object that is currently locked by another application

Lock App Handle: System wide Unique ID of the application holding the lock

>

Lock App ID: ID of the application holding the lock

Lock Sequence ID: The sequence number of the application that is holding a lock on the object that this application is waiting to obtain.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Lock Waits tab
Metrics	AuthID, Name, Handle, ID, and Client PID help you identify the application waiting on the lock. Lock App Handle, Lock App ID, and Lock Sequence number help you identify the application that is holding the waited on lock. You can use this information on Users Detail Locks tab to get further details on the application holding the locks and the locks it's holding.
Correction	If an application is not performing well, and it has a high number of total lock waits, and a high total lock waiting time you may need to improve concurrency of the applications that are holding the locks.

Applications Waiting on Locks

This metric indicates the percentage of all connected applications waiting on locks.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Lock Waits tab
Metrics	If this number is high, the application may have concurrency problems.
Correction	You can identify the applications that are holding locks or exclusive locks for a long time from the Users Detail> Locks page and tune such applications for better concurrency.

Locks tab

The Locks tab of the Users Detail view gives you information for the following statistics:

Application Details	Locks Held Details	Locks Waiting Details

Application Details

This section shows a list of applications connected to the database. You can look at the lock details for an application in the list by selecting the application. The attributes presented for each application include:

Handle: A system-wide **unique** ID for the application

ID: This identifier is generated when the application connects to the database at the database manager

Seq No

Name: Name of the application executable

Auth ID: The authorization ID of the user who invoked the application that is being monitored.

Status

Code Page

Locks Held

Total Wait Time:

UOW Lock Waiting Time

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Locks tab

>

Locks Held Details

This section gives the details of all the locks that the selected application is currently holding. The details presented here include:

Lock Name: Internal binary lock name. This element serves as a unique identifier for locks

Lock Attributes: Gives the lock attribute settings

Release Flags: The lock release flags can either be *untracked* or *SQL Compiler*

Lock Count: The number of locks on the lock being held. This value ranges from 1 to 255. It is incremented as new locks are acquired, and decremented as locks are released. When lock_count has a value of 255, this indicates that a transaction duration lock is being held. At this point, lock_count is no longer incremented or decremented when locks are acquired or released. The lock_count element is set to a value of 255 in one of two possible ways:

1 lock_count is incremented 255 times due to new locks being acquired.

2 A transaction duration lock is explicitly acquired. For example, with a LOCK TABLE statement, or an INSERT.

Hold Count: The number of holds placed on the lock. Holds are placed on locks by cursors registered with the WITH HOLD clause and some DB2 utilities. Locks with holds are not released when transactions are committed.

Object Name: It is the name of the object for table-level locks is the file ID (FID) for SMS and DMS table spaces. For row-level locks, the object name is the row ID (RID). For table space locks, the object name is blank. For buffer pool locks, the object name is the name of the buffer pool.

Object Type: The type of object against which the application holds a lock (for object-lock-level information), or the type of object for which the application is waiting to obtain a lock (for application-level and deadlock-level information).

Tablespace: This is the name of the table space against which the lock is held.

Table Schema: Schema of the table that the lock is on

Table Name: Name of the table that the lock is on. This element is only set if *Object Type* indicates Table

Lock Mode: The type of lock being held.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Locks tab

Locks Waiting Details

This section gives the details of all the locks that the selected application is currently waiting on. The details include:

Subsection: Identifies the subsection associated with the waited on lock

Agent ID: The application handle of the agent holding a lock for which this application is waiting

App ID: The application ID of the application holding a lock on the object that this application is waiting to obtain.

Lock Name: Internal binary lock name. This element serves as a unique identifier for locks.

Lock Attributes: Gives the lock attribute settings

Release Flags: Lock release flags

Object Type: The type of object against which the application holds a lock

Lock Mode: The type of lock being held

Lock Mode Requested: The lock mode being requested by the application

Tablespace: This is the name of the table space against which the lock is held.

Table Schema: Schema of the table that the lock is on

Table Name: Name of the table that the lock is on. This element is only set if *Object Type* indicates Table

Wait Start Timestamp: The date and time that this application started waiting to obtain a lock on the object that is currently locked by another application

Escalation: Indicates whether a lock request was made as part of a lock escalation.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Users > Locks tab

Cfg (Database Configuration)

The Database Configuration view displays the following tabbed pages:

Configuration tab	History tab	Indoubt Transactions tab
Objects tab		

Configuration tab

This tab lets you view the current values for the DB2 database configuration parameters.

History tab

The History tab of the Database Configuration Details viewlets you browse the history file entries for the database. The history file contains a record of recovery and administrative events. Recovery events include full database and table space level backup, incremental backup, restore, and rollforward operations. Additional logged events include create, alter, drop, or rename table space, reorganize table, drop table, and load.

On this tab, you'll find the following information:

History Details	SQLCA Structure	Tablespaces List
DDL View		

History Details

This section allows you to select the operation for which you want to see the history details. You can also narrow the history list to operations since a certain date or operations involving a certain object name. You can show the history for the following types of operations:

Backups	Dropped Tables	Load
Quiesce	Reorganizations	Rollforward
Tablespace Alters	Tablespace Renames	Tablespace Creates
Tablespace Drops		

You can see further detail on some operations in the lower sections, by selecting that operation from the history detail operation list.

SQLCA Structure	Shows the SQLCA structure for the selected operation from the history detail. This information is not available for Quiesce, Rollforward, and Reorganization operations.

Tablespaces List	Shows the list of tablespaces involved in the selected operation in history detail.
DDL View	Shows the DDL for the selected operation in the history detai

Indoubt Transactions tab

The Indoubt Transactions tab provides a list of transactions that are in doubt. The Indoubt Transaction Detail metrics displayed for each indoubt transaction include:

Application ID: Specifies the application identifier assigned by the database manager for this transaction

Sequence Number: Specifies the sequence number assigned by the database manager as an extension to the application ID

Status: Indicates the status of this indoubt transaction. Valid values are:

SQLXA_TS_PREP	The transaction is prepared. The connected parameter can be used to determine whether the transaction is waiting for the second phase of normal commit processing or whether an error occurred and resynchronization with the transaction manager is required.
SQLXA_TS_HCOM	The transaction has been heuristically committed.
SQLXA_TS_HROL	The transaction has been heuristically rolled back.
SQLXA_TS_MACK	The transaction is missing commit acknowledgement from a node in a partitioned database.
SQLXA_TS_END	The transaction has ended at this database. This transaction may be re-activated, committed, or rolled back at a later time. It is also possible that the transaction manager encountered an error and the transaction will not be completed. If this is the case, this transaction requires heuristic actions, because it may be holding locks and preventing other applications from accessing data.

Timestamp: Specifies the time when the transaction entered the indoubt state.

Auth ID: Specifies the authorization ID of the user who ran the transaction.

Log Full: Indicates whether or not this transaction caused a log full condition

Originator

XID: Specifies the XA identifier assigned by the transaction manager to uniquely identify a global transaction

Objects tab

Tables and indexes consume the storage in all databases. The Database Object Detail grid displays object space details for the database selected in the Database Object Summary grid.

The following statistics are presented for tables:

Table Schema: The owner of the object

Table Name: The name of the object

File Group: File group where the object resides

Reserved (KB): The amount of space (in KB) reserved by the object

Used (KB): The amount of space (in KB) used by the object

Free (KB): The amount of free space (in KB) used by the object

Percent Used: The percentage of space used by the object

>

Instance

The Embarcadero Performance Analyst Instance view lets you review availability and overall performance of all monitored databases from a single window.

In addition to a Home page, the Instance category of DB2 Performance Analyst includes the following tabbed pages:

Configuration tab	Agents & Connections tab	Utilities tab
FCM Resources tab	FCM Throughput tab	Memory Pools tab
Sorts & Joins tab		

Home

Statistics on the Instance home view are organized into the following categories:

Agent Analysis pane	Bottleneck Analysis pane	Location
Instance Information pane	Key Ratio Analysis pane	

Agent Analysis pane

The Agent Analysis section presents the following metrics:

Connections	Connections Executing	Agents Registered
Agents Waiting	Idle Agents	

Connections

The connections statistic gives the total Number of connections to databases in the current instance.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > Agent Analysis pane
Metrics	Shows the number of connections from remote/local clients to databases in this instance. This changes frequently, so you may need to sample it at specific intervals over an extended period of time to get a realistic view of system usage.
Correction	This element can help you adjust the setting of the *max_coordagents* configuration parameter.

Connections Executing

The connections executed statistic is the number if applications that are currently connected to a database and are currently processing a unit of work within the Database Manager instance being monitored.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > Agent Analysis pane
Metrics	The connections executing metric can help you determine the level of concurrent processing occurring on the Database Manager. This value can change frequently, so you may need to sample it at specific intervals over an extended period of time to get a realistic view of system usage
Correction	This element can help you adjust the setting of the *maxcagents* configuration parameter.

Agents Registered

The agents registered metric is the number of agents registered in the Database Manager instance that is being monitored (coordinator agents and subagents).

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > Agent Analysis pane
Correction	You can use this element to help evaluate your setting for the *maxagents* configuration parameter.

Agents Waiting

The number of agents waiting for a token so they can execute a transaction in the database manager.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > Agent Analysis pane
Metrics	Each application has a dedicated coordinator agent to process database requests within the Database Manager. Furthermore, each agent has to get a token before it can execute a transaction. The maximum number of agents that can execute Database Manager transactions is limited by the configuration parameter *maxcagents*.
Correction	You can use this element to help evaluate your setting for the *maxcagents* configuration parameter.

Idle Agents

The idle agents metric is the number of agents in the agent pool that are currently unassigned to an application and are, therefore, "idle."

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > Agent Analysis pane
Metrics	Having idle agents available to service requests for agents can improve performance.
Correction	You can use this element to help set the *num_poolagents* configuration parameter.

Bottleneck Analysis pane

The Bottleneck Analysis section presents the following metrics:

Sorts and Joins

Post Threshold Sorts	Piped Sorts Rejects	Hash Join Thresholds

Agent Bottlenecks

Agents Waiting on a Token	Agents Created Due to Empty Pool	

Post Threshold Sorts

This metric shows the number of sorts that have requested heaps after the sort heap threshold has been exceeded.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > Bottleneck Analysis pane

Metrics	Under normal conditions, the Database Manager allocates sort heap using the value specified by the *sortheap* configuration parameter. If the amount of memory allocated to sort heaps exceeds the sort heap threshold (*sheapthres* configuration parameter), the database manager allocates sort heap using a value less than that specified by the *sortheap* configuration parameter. Each active sort on the system allocates memory, which may result in sorting taking up too much of the system memory available. Sorts that start after the sort heap threshold has been reached may not receive an optimum amount of memory to execute. As a result, however, the entire system may benefit.
Correction	*By modifying the sort heap threshold and sort heap size configuration parameters, the performance of sort operations and/or the overall system can be improved. If this element's value is high, you can:* *Increase the sort heap threshold (sheapthres) or,* *Adjust applications to use fewer or smaller sorts via SQL query changes.*

Piped Sorts Rejects

The piped sorts rejects statistic is the number of piped sorts that were rejected by the database manager.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > Bottleneck Analysis pane
Metrics	Each active sort allocates memory, and can result in sorting taking up too much available system memory. A piped sort is not accepted if the sort heap threshold is exceeded when the sort heap is allocated for the sort. The sort list heap (*sortheap*) and sort heap threshold (*sheapthres*) configuration parameters help control the amount of memory used for sort operations. These parameters are also used to determine whether a sort will be piped. Since piped sorts may reduce disk I/O, allowing more piped sorts can improve the performance of sort operations and possibly the performance of the overall system.
Correction	If piped sorts are being rejected, you might consider decreasing your sort heap or increasing your sort heap threshold. You should be aware of the possible implications of either of these options. If you increase the sort heap threshold there is the possibility that more memory will remain allocated for sorting. This could cause paging memory to disk. If you decrease the sort heap, you might require an extra merge phase that could slow down the sort

Hash Join Thresholds

The hash join thresholds metric is the total number of times that a hash join heap request was limited because of concurrent use of shared or private sort heap.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > Bottleneck Analysis pane
Metrics	Each active sort allocates memory, and can result in sorting taking up too much of the available system memory. A piped sort is not accepted if the sort heap threshold is exceeded when the sort heap is allocated for the sort. The sort list heap (*sortheap*) and sort heap threshold (*sheapthres*) configuration parameters help control the amount of memory used for sort operations. These parameters are also used to determine whether a sort will be piped. Since piped sorts may reduce disk I/O, allowing more piped sorts can improve the performance of sort operations and possibly the performance of the overall system.
Correction	If this value is large (greater than 5% of Hash Join Overflows), the sort heap threshold should be increased.

Agents Waiting on a Token

The agents waiting statistic is the percentage of agents registered that are waiting for a token so they can execute a transaction in the Database Manager.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > Bottleneck Analysis pane

Metrics	Each application has a dedicated coordinator agent to process database requests within the Database Manager. Each agent has to get a token before it can execute a transaction. The maximum number of agents that can execute database manager transactions is limited by the configuration parameter *maxcagents*. This parameter is used to control the load on the system during periods of high simultaneous application activity. For example, you may have a system requiring a large number of connections but with a limited amount of memory to serve those connections. Adjusting this parameter can be useful in such an environment, where a period of high simultaneous activity could cause excessive operating system paging. This parameter does not limit the number of applications that can have connections to a database. It only limits the number of Database Manager agents that can be processed concurrently by the Database Manager at any one time, thereby limiting the usage of system resources during times of peak processing.
Correction	In cases where the high concurrency of applications is causing problems, you can use benchmark testing to tune this parameter to optimize the performance of the database.

Agents Created Due to Empty Pool

The agents created due to empty pool metric shows the number of agents created because the agent pool was empty. It includes the number of agents started when the current instance started (*num_initagents*).

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > Bottleneck Analysis pane

FCM Analysis pane

The fast communication manager (FCM) provides communication support for Enterprise Server Editions. Each database partition server has one FCM thread to provide communications between database partition servers to handle agent requests and to deliver message buffers. The FCM thread starts when you start the instance.

The following ratios are used on the FCM Analysis section to succinctly communicate the general overall FCM performance levels of the monitored instance:

Buffer Throughput	Buffer Utilization	Connection Entry Utilization
Message Anchor Utilization	Request Block Utilization	

Buffer Throughput

The buffer throughput metric is the total number of buffers that have been sent and received by this database partition since the start of monitoring.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > FCM Analysis pane
Metrics	This gives the level of traffic to and from the selected partition.
Correction	If this traffic level is high, you should consider redistributing one or more databases on the instance or move tables to reduce internode traffic.

Buffer Utilization

This element indicates the percentage of all FCM buffers currently being used by the fast communication manager.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > FCM Analysis pane
Correction	You can use this information to tune *fcm_num_anchors*. If the utilization percentage is high, you should increase the *fcm_num_anchors* to ensure that operations do not run out of FCM message anchors. If the utilization is low, you can decrease *fcm_num_anchors* to conserve system resources.

Connection Entry Utilization

The connection entry utilization element indicates the percentage of all connection entries that are currently being used by the fast communication manager.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > FCM Analysis pane
Correction	You can use this information to tune *fcm_num_connect.* If the utilization percentage is high, you should increase the *fcm_num_connect* to ensure that operations do not run out of FCM connection entries. If the utilization is low, you can decrease *fcm_num_connect* to conserve system resources.

NOTE: This metric is only applicable to DB2 version 7. In DB2 version 8, the maximum connection entries are adjusted dynamically and automatically

Message Anchor Utilization

This metric shows the percentage of all message anchors currently being used by the fast communication manager.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > FCM Analysis pane
Correction	You can use this information to tune *fcm_num_buffers.* If the utilization percentage is high, you should increase the *fcm_num_buffers* to ensure that operations do not run out of FCM buffers. If the utilization is low, you can decrease *fcm_num_buffers* to conserve system resources.

NOTE: This metric is only applicable to DB2 version 7. In DB2 version 8, the maximum message anchors are adjusted dynamically and automatically

Request Block Utilization

This element indicates the percentage of all request blocks that are currently being used by the fast communication manager.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > FCM Analysis pane
Correction	You can use this information to tune *fcm_num_rqb.* If the utilization percentage is high, you should increase the *fcm_num_rqb* to ensure that operations do not run out of FCM request blocks. If the utilization is low, you can decrease *fcm_num_rqb* to conserve system resources.

NOTE: This metric is only applicable to DB2 version 7. In DB2 version 8, the maximum request blocks are adjusted dynamically and automatically

Instance Information pane

Here you see the following information:

Product Name	FixPack Level	Instance Type
Operating System	Instance Name	Number of Partitions
Active Local Databases		

Key Ratio Analysis pane

The Key Ratio Analysis section presents the following metrics:

>

Efficiency

Agents Registered	Agents Waiting	Monitor Heap
FCMBP	Private Sort Memory	

Agents Registered

The agents registered metric shows the percentage of total allowable agents (coordinator agents and subagents) that are currently working in the instance.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > Key Ratio Analysis pane
Metrics	When this value is close to 100% it means that the number of agents currently working is close to the hard limit set in the database manager configuration. If this value stays high for long periods of time it means that there is high level of activity occurring in the databases on the instance.
Correction	Use this element to help you evaluate your setting for the *maxagents* database manager configuration parameter. When you increase the value of this parameter, make sure that your system has enough memory available to accommodate the additional agents that may be created.

Agents Waiting

The agents waiting statistic is the percentage of agents registered that are waiting for a token so they can execute a transaction in the Database Manager.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > Key Ratio Analysis pane
Metrics	Each application has a dedicated coordinator agent to process database requests within the Database Manager. Each agent has to get a token before it can execute a transaction. The maximum number of agents that can execute database manager transactions is limited by the configuration parameter *maxcagents*. This parameter is used to control the load on the system during periods of high simultaneous application activity. For example, you may have a system requiring a large number of connections but with a limited amount of memory to serve those connections. Adjusting this parameter can be useful in such an environment, where a period of high simultaneous activity could cause excessive operating system paging. This parameter does not limit the number of applications that can have connections to a database. It only limits the number of Database Manager agents that can be processed concurrently by the Database Manager at any one time, thereby limiting the usage of system resources during times of peak processing.
Correction	In cases where the high concurrency of applications is causing problems, you can use benchmark testing to tune this parameter to optimize the performance of the database.

Monitor Heap

This metric shows the percentage of total monitor heap that is currently being used by the Database Manager.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > Key Ratio Analysis pane
Metrics	The amount of memory required for monitoring activity depends on the number of monitoring applications (applications taking snapshots or event monitors), which switches are set, and the level of database activity. Memory is allocated from the monitor heap when you perform database monitoring activities such as taking a snapshot, turning on a monitor switch, resetting a monitor, or activating an event monitor. If this metric is running at over 70% utilization and you have many applications performing snapshot /event monitoring, you may need to increase the size of the monitor heap.

Correction	Increase the value of *mon_heap_sz* database manager configuration parameter to increase the available monitor heap size.

FCMBP

The fast communication manager bufferpool (FCMBP) is the percentage of total amount of the FCMBP that is currently being used. This metric is only applicable to multipartitioned database environments.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > Key Ratio Analysis pane
Metrics	You can use this information to tune the *fcm_num_buffers* database manager configuration parameter.

NOTE: FCM buffers are used for internal communication both among and within database servers. The FCMBU metric gives the percentage utilization of the available fcm buffers in the Database Manager.

Private Sort Memory

The private sort memory metric is the percentage of the total private sort memory heap that is currently in use.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Home > Key Ratio Analysis pane
Metrics	You can use this information to tune the *sheapthres* database manager configuration parameter. If the value of this element is close to 100%, it means that the sorts for some databases are not getting the full sort heap as defined by the *sortheap* database configuration parameter.

Configuration tab

The Configuration tab of the Instance Detail view lets you view and change the instance configuration parameter. It also allows you to look at some information about the instance such as Product Name, Type, Fixpak Level, Operating System, instance name, partitions, and number of active databases.

Agents & Connections tab

The Agents & Connections tab of the Instance Detail view gives you the following analyses:

Connection Analysis	Agent Pool Analysis	Agent Analysis

Connection Analysis

This section presents the metrics for the current local and remote connections to databases in the instance. The information presented here includes:

Local Connections: The number of local applications that are currently connected to a database within the database manager instance being monitored.

Remote Connections: The current number of connections initiated from remote clients to the instance of the database manager that is being monitored.

Local Connections Executing: The number of local applications that are currently connected to a database within the database manager instance being monitored and are currently processing a unit of work.

>

Remote Connections Executing: The number of remote applications that are currently connected to a database and are currently processing a unit of work within the database manager instance being monitored.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Agents & Connections tab
Metrics	These metrics can help you determine the level of concurrent processing occurring in the database manager. These metrics will change frequently so it is important to obtain a sampling over an extended period. These metrics are useful when tuning the *max_coordagents,* and *maxcagents* database manager configuration parameters.

Agent Pool Analysis

The metrics in this section help determine how often the database agents are created due to empty pool and how often they are assigned from the agent pool.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Agents & Connections tab
Metrics	If the number of agents created due to empty pool is high compared to the agents assigned from pool, it may mean that the number of pool agents should be increased. It may also mean that the overall workload is too high.
Correction	You can change the *maxcagents* and *num_poolagents* database manager configuration parameters to tune these metrics.

Agent Analysis

The metrics presented here give a detailed view of how the database manager agents are working. The metrics presented here include:

Registered Agents: The number of agents registered in the database manager instance that is being monitored (coordinator agents and subagents).

Agents Waiting for Token: The number of agents waiting for a token so they can execute a transaction in the database manager.

Idle Agents: The number of agents in the agent pool that are currently unassigned to an application and are, therefore, "idle".

Stolen Agents: The number of times that agents are stolen from an application. Agents are stolen when an idle agent associated with an application is reassigned to work on a different application.

Maximum Agents Registered: The maximum number of agents that the database manager has ever registered, at the same time, since it was started (coordinator agents and subagents).

Maximum Agents Waiting: The maximum number of agents that have ever been waiting for a token, at the same time, since the database manager was started.

Maximum Coordinating Agents: The maximum number of coordinating agents working at one time.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Agents & Connections tab

Utilities tab

The Utilities tab of the Instance Detail View gives information on the utilities currently executing in the database manager (available in DB2 version 8 only). The Utility Execution Details are as follows:

Utility ID: The unique identifier corresponding to the utility invocation

>

Type: Class of the utility (Rebalance, Backup, Restore, Reorg etc..)

Utility Description: A brief description of the work the utility is performing

Utility Start Time: The time the utility was invoked

DBName: Database on which the utility is operating.

FCM Resources tab

The FCM Resources tab of the Instances Detail View presents statistics on the following:

Node Details	FCM Buffer Utilization	FCM Message Anchor Utilitization
FCM Connection Entries Utilization	FCM Connection Entries Utilization	

Node Details

This section gives a detailed view of inter- node communication in a multipartition environment. The information shown here includes the connection status and buffers sent and received between various combinations of partitions in a multi-partitioned environment. The columns presented here include:

Source Node: Partition that sends the information

Target Node: Partition that received the information

Buffers Sent: Number of Buffers sent from source node to the target node

Buffers Received: Number of Buffers sent from source node to the target node

Connection Status

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > FCM Resources tab IBM DB2 (W/U/L) Performance Analyst statistics > Instance > FCM Throughput tab

FCM Buffer Utilization

This section provides details on the current utilization and low watermark of FCM buffers.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > FCM Resources tab
Metrics	You can monitor the buffer utilization for the selected partition together with the low watermark to determine the usage trend of the fcm buffers.
Correction	If you notice a very small value for the low watermark and similar values for free FCM buffers over a long period, you may need to increase the *fcm_num_buffers* database manager configuration parameter to increase the number of available FCM buffers. If on the other hand you notice a very high value for the low watermark – you may have allocated too many buffers and you can reduce the value of *fcm_num_buffers* database manager configuration parameter.

FCM Message Anchor Utilitization

This section provides details on the current utilization and low watermark of FCM message anchors.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > FCM Resources tab
Metrics	You can monitor the message anchor utilization for the selected partition together with the low watermark to determine the usage trend of the message anchors.
Correction	If you notice a very small value for the low watermark and similar values for free message anchors over a long period, you may need to increase the *fcm_num_anchors* database manager configuration parameter to increase the number of available message anchors. If on the other hand you notice a very high value for the low watermark – you may have allocated too many message anchors and you can reduce the value of *fcm_num_anchors* database manager configuration parameter.

NOTE: DB2 version 8 automatically adjusts *fcm_num_anchors* parameter, so you do not need to adjust this manually in DB2 version 8.

FCM Connection Entries Utilization

This section provides details on the current utilization and low watermark of FCM connection entries.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > FCM Resources tab
Metrics	You can monitor the connection entry utilization for the selected partition together with the low watermark to determine the usage trend of connection entries.
Correction	If you notice a very small value for the low watermark and similar values for free connection entries over a long period, you may need to increase the *fcm_num_connect* database manager configuration parameter to increase the number of available connection entries. If on the other hand you notice a very high value for the low watermark – you may have allocated too many connection entries and you can reduce the value of *fcm_num_connect* database manager configuration parameter.

NOTE: DB2 version 8 automatically adjusts *fcm_num_rqb* parameter, so you do not need to adjust this manually in DB2 version 8.

FCM Request Block Utilization

This section provides details on the current utilization and low watermark of FCM request blocks.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > FCM Resources tab
Metrics	You can monitor the request block utilization for the selected partition together with the low watermark to determine the usage trend of the request blocks.
Correction	If you notice a very small value for the low watermark and similar values for free request blocks over a long period, you may need to increase the *fcm_num_rqb* database manager configuration parameter to increase the number of available request blocks. If on the other hand you notice a very high value for the low watermark – you may have allocated too many request blocks and you can reduce the value of *fcm_num_rqb* database manager configuration parameter.

NOTE: DB2 version 8 automatically adjusts *fcm_num_rqb* parameter, so you do not need to adjust this manually in DB2 version 8.

>

FCM Throughput tab

The FCM Throughput tab of the Instance Detail VIew includes the following statistics:

- FCM Throughput

The following statistic, available on this tab, duplicates a statistic available on the Instance FCM Resources tab:

- Node Details

FCM Throughput

This section shows the delta statistics for the FCM buffers sent and received between the selected nodes in the Node Details section

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > FCM Throughput tab
Metrics	If the total number of FCM buffers sent or received between selected nodes is high, you may want to redistribute the database, or move tables to reduce the inter-node traffic.

Memory Pools tab

The Memory Pools tab is divided into three panes. Unique to this tab is the Memory Pool List pane. The Memory Pool Details and the Memory Pool Utilization are duplicates of panes available in the Memory category.

Memory Pool List pane

This section shows the memory pool usage details of the selected memory pool over time.

Sorts & Joins tab

The Sorts & Joins tab of the Instance Detail View presents the following information:

Sort Heap Details	Sort Threshold Details	Hash Join Threshold Details
Piped Sort Details		

Sort Heap Details

This section shows the metrics for the total number of allocated pages of sort heap space for all sorts. It represents the sum of sort heap space allocated for all sorts in all active databases in the database manager.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Sorts & Joins tab
Metrics	Normal memory estimates do not include sort heap space. If excessive sorting is occurring, the extra memory used for the sort heap should be added to the base memory requirements for running the database manager. Generally, the larger the sort heap, the more efficient the sort. Appropriate use of indexes can reduce the amount of sorting required.
Correction	You may use the information in this section to help you tune the *sheapthres* configuration parameter. If the element value is greater than or equal to *sheapthres*, it means that the sorts are not getting the full sort heap as defined by the *sortheap* parameter.

Sort Threshold Details

This section details the number of sorts that have requested heaps after the sort heap threshold has been exceeded.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Sorts & Joins tab
Metrics	Under normal conditions, the database manager will allocate sort heap using the value specified by the *sortheap* configuration parameter. If the amount of memory allocated to sort heaps exceeds the sort heap threshold (*sheapthres* configuration parameter), the database manager will allocate sort heap using a value less than that specified by the sortheap configuration parameter. Each active sort allocates memory, which may result in sorting taking up too much of the system memory available. Sorts that start after the sort heap threshold has been reached may not receive an optimum amount of memory to execute, but, as a result, the entire system may benefit. By modifying the sort heap threshold and sort heap size configuration parameters, the performance of sort operations and/or the overall system can be improved.
Correction	If this element's value is high, you can: *Increase the sort heap threshold (sheapthres) or,* *Adjust applications to use fewer or smaller sorts via SQL query changes.*

Hash Join Threshold Details

This section shows the total number of times that a hash join heap request was limited due to concurrent use of shared or private sort heap space.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Sorts & Joins tab
Correction	If this value is large (greater than 5% of hash join overflows), the sort heap threshold should be increased

Piped Sort Details

This section gives the details on the piped sorts that have been requested and accepted.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > Instance > Sorts & Joins tab
Metrics	Each active sort on the system allocates memory, which may result in sorting taking up too much of the available system memory. When the number of accepted piped sorts is low compared to the number requested, you can improve sort performance by adjusting one or both of the following configuration parameters: `sortheap` and **sheapthres**.
Correction	If piped sorts are being rejected, you might consider decreasing your sort heap or increasing your sort heap threshold. You should be aware of the possible implications of either of these options. If you increase the sort heap threshold, then there is the possibility that more memory will remain allocated for sorting. This could cause the paging of memory to disk. If you decrease the sort heap, you might require an extra merge phase that could slow down the sort.

OS

In many scenarios, an optimally tuned database may not perform well because there are constraints imposed by the system where the database is running. These constraints may include processes competing with the database sever for resources (CPU, I/O, or Memory), a slow CPU, insufficient or slow I/O devices, and insufficient memory. The OS Statistics page of Performance Analyst lets you examine operating system metrics for the following platforms:

- AIX

- HP-UX

  NOTE: To view processor info and swap disk info on an HP-UX box, you need to login as ROOT in the OS login.

- Linux
- Solaris
- Unix
- Windows XP and 2000

  NOTE: The statistics available on the OS page depend on the operating system platform.

  TIP: If you magnify a graph to change back to the original size or minimize the size, close the OS Page and then reopen it again. Performance Analyst displays the graph in its original form and size.

In addition to a Home page, the OS category of DB2 Performance Analyst includes the following tabbed pages:

CPU Tab	I/O Tab	Memory Tab
Network Tab	Processes Tab	Space Tab

Home

The OS home page includes the following sections:

Key Resource Usage pane	Bottleneck Analysis pane	Memory Analysis pane
Disk Analysis panes	Workload Analysis pane	

Key Resource Usage pane

The following ratios are used on the Performance Analyst OS home page to communicate the general overall performance levels of the operating system:

Disk Time	Load Average	Processor Time
Paged Memory Used (Windows)	Swap Memory Used (AIX, HP-UX, Linux, Solaris, Unix)	

TIP: To open the CPU tab, right-click Processor Time. To open the I/O tab, right-click Disk Time. To open the Memory tab, right-click Paged Memory Used.

Disk Time

The Disk Time statistic is the percentage of elapsed time that the selected disk drive/device was busy servicing read or write requests.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	You should avoid consistently seeing values for this statistic greater then 90%.
Correction	Add more disk drives and partition the files among all of the drives.

Load Average

The Load Average statistic represents the system load averages over the last 1, 5, and 15 minutes.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	High load averages usually mean the system is being used heavily and the response time is correspondingly slow.

Processor Time

The Processor Time statistic indicates the percentage of time the processor is working. This counter is a primary indicator of processor activity.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	If your computer seems to be running sluggishly, this statistic could be displaying a high percentage.
Correction	Upgrade to a processor with a larger L2 cache, a faster processor, or install an additional processor.

Paged Memory Used

The Paged Memory Used statistic is the ratio of Commit Memory Bytes to the Commit Limit. Committed memory is where memory space has been reserved in the paging file if it needs to be written to disk. The commit limit is determined by the size of the paging file. As the paging file increases, so does the commit limit.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	This value displays the current percentage value only and not an average. If the percentage of paged memory used is above 90%, you may be running out of memory.
Correction	Increase the size of page file.

NOTE: This statistic is available for the Windows platform.

Swap Memory Used

The Swap Memory Used statistic is the percentage of swap space currently in use.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Key Resource Usage pane
Metrics	If the percentage of swap memory used is above 90%, you may be running out of memory.
Correction	Increase the size of your swap files.

Bottleneck Analysis pane

The following ratios are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Average Disk Queue Length	Interrupts/Sec	Network Output Queue Length/Network Queue Length
Page Faults/Sec	Processor Queue Length	Processor Speed
Processor		

TIP: To open the I/O tab, right-click any Details menu item. To open the CPU tab, right-click the Item Processor Queues Length Details menu item. To open the Network tab, right-click Network Output Queues Length.

NOTE: The statistics available in this section depend on the platform of operating system.

Average Disk Queue Length

The Average Disk Queue Length statistic is the average number of both read and write requests that were queued for the selected disk during the sample interval.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Bottleneck Analysis pane
Metrics	This metric is useful in identifying I/O related bottlenecks. If the disk queue lengths for certain disks are consistently much higher than others, you may need to redistribute the load among available disks. If the disk queues lengths for all disks are consistently large, and you see a high amount of I/O activity, your disks may be inefficient.
Correction	Some things you can do if you have problems with this statistic include: *Redistribute the data on the disk with the large average disk queue to other disks.* *Upgrade to faster disk(s).*

Interrupts/Sec

Interrupts/Sec is the average rate, in incidents per second, at which the processor received and serviced hardware interrupts. This value is an indirect indicator of the activity of devices that generate interrupts, such as the system clock, the mouse, disk drivers, data communication lines, network interface cards, and other peripheral devices. These devices normally interrupt the processor when they have completed a task or require attention. Normal thread execution is suspended. The system clock typically interrupts the processor every 10 milliseconds, creating a background of interrupt activity. This statistic shows the difference between the values observed in the last two samples, divided by the duration of the sample interval.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Bottleneck Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > OS > CPU Tab
Metrics	A high value indicates possible excessive hardware interrupts; justification is dependent on device activity.

Network Output Queue Length/Network Queue Length

The Network Output Queue Length statistic is the number of threads in the processor queue.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Bottleneck Analysis pane
Metrics	Unlike the disk counters, this counter shows ready threads only, not threads that are running. There is a single queue for processor time even on computers with multiple processors. Therefore, if a computer has multiple processors, you need to divide this value by the number of processors servicing the workload. A sustained processor queue of less than 10 threads per processor is normally acceptable, dependent of the workload.
Correction	A sustained high value in the Processor Queue Length could indicate that a processor bottleneck has developed due to threads of a process requiring more process cycles than are available. If this is the case, you should look at installing a faster (or an additional) processor.

NOTE: The name of this statistic depends on the platform of the operating system.

Page Faults/Sec

The Page Faults/Sec statistic is the overall rate faulted pages are handled by the processor. It is measured in numbers of pages faulted per second. A page fault occurs when a process requires code or data that is not in its working set. This counter includes both hard faults and soft faults.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Bottleneck Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab
Metrics	This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
Correction	If the number of page faults remains consistently high, you can check with your Windows System Administrator for further investigation. Often, large numbers of page faults are not a problem so long as they are soft faults. However, hard faults, that require disk access, can cause delays.

Processor Queue Length

The Processor Queue Length statistic is the number of threads in the processor queue.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Bottleneck Analysis pane IBM DB2 (W/U/L) Performance Analyst statistics > OS > CPU Tab
Metrics	Unlike the disk counters, this counter shows ready threads only, not threads that are running. There is a single queue for processor time even on computers with multiple processors. Therefore, if a computer has multiple processors, you need to divide this value by the number of processors servicing the workload. A sustained processor queue of less than 10 threads per processor is normally acceptable, dependent of the workload.
Correction	A sustained high value in the Processor Queue Length could indicate that a processor bottleneck has developed due to threads of a process requiring more process cycles than are available. If this is the case, you should look at installing a faster (or an additional) processor.

Processor Speed

The Processor Speed statistic displays the speed of the active processor in MHz. The speed is approximate.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Bottleneck Analysis pane

Processor

The Processor Statistic displays the type of processor currently in use, for example, GenuineIntel.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Bottleneck Analysis pane

Memory Analysis pane

The following metrics are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Available Paged Memory (Windows)	Available Physical Memory	Available Swap Memory (AIX, HP-UX, Linux, Solaris, Unix)
Total Paged Memory/Total Swap Memory	Total Physical Memory	

TIP: To open the Memory tab, right-click any Details menu item.

NOTE: The statistics available in this section depend on the platform of operating system.

Available Paged Memory

The Available Paged Memory statistic shows the amount of virtual memory available for the processes.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	If the available virtual memory is less than 10% of the total virtual memory, your system may run out of memory.
Correction	Increase the size of page file.

NOTE: This statistic is available for the Windows platform.

Available Physical Memory

The Available Physical Memory statistic represents the amount of RAM available to all processes.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	This counter displays the last observed value only and not an average. Use this value with the Total physical memory and paging metrics (Memory details page). If the available physical memory is very small compared to this value, and the paging activity is high, your system may be running low on memory.
Correction	Some things you can do if you have problems with this statistic include: *Check the running processes to see if there are any memory leaks.* *Stop any services that are not required.* *Install additional RAM.*

Available Swap Memory

The Available Swap Memory statistic represents the amount of virtual memory available for the processes.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Memory Analysis pane

Metrics	If the Available Swap Memory is less than 10% of the total Swap Memory, your system may run out of memory.
Correction	Increase the size of swap files.

Total Paged Memory/Total Swap Memory

The Total Paged Memory statistic shows the maximum amount of virtual memory available to all processes.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	It is recommended that this be 1½ - 3 times the amount of RAM on the system.

NOTE: The name of this statistic depends on the platform of the operating system.

Total Physical Memory

The Total Physical Memory statistic shows the amount of physical memory installed on your computer.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Memory Analysis pane
Metrics	This is an informational metric and displays the total amount installed on the machine. Use this value with the available physical memory and paging metrics (Memory details page). If the available physical memory is very small compared to this value, and the paging activity is high, your system may be running low on memory.

Disk Analysis pane

The following ratios are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Free Disk Space	Total Disk Space	Used Disk Space

TIP: To open the Space tab, right-click any Details menu item.

NOTE: The statistics available in this section depend on the platform of operating system.

Free Disk Space

The Free Disk Space statistic shows the unallocated space, in megabytes on all logical disk drives.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Disk Analysis pane
Correction	There are many things a DBA can do to ensure that a database does not encounter a space problem due to physical space limitations: *If a database currently resides on a disk that has little free space, you can add more files to the database. Of course, you should add the new files to other physical hard disks that can accommodate a growing database.* *You should examine hard disks with shrinking disk space to see if you can relocate or delete files to allow more free space.*

>

Total Disk Space

Total Disk Space displays the total allocated and unallocated space, in megabytes on all logical disk drives.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Disk Analysis pane
Correction	There are many things a DBA can do to ensure that a database does not encounter a space problem due to physical space limitations, here are two: *If a database currently resides on a disk that has little free space, you can add more files to the database. Of course, you should add the new files to other physical hard disks that can accommodate a growing database.* *You should examine hard disks with shrinking disk space to see if you can relocate or delete files to allow more free space.*

Used Disk Space

The Used Disk Space statistic shows the amount of allocated space, in megabytes on all logical disk drives.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Disk Analysis pane
Correction	There are many things a DBA can do to ensure that a database does not encounter a space problem due to physical space limitations: *If a database currently resides on a disk that has little free space, you can add more files to the database. Of course, you should add the new files to other physical hard disks that can accommodate a growing database.* *You should examine hard disks with shrinking disk space to see if you can relocate or delete files to allow more free space.*

Workload Analysis pane

The following ratios are used on the Performance Analyst OS home page to succinctly communicate the general overall performance levels of the operating system:

Number of Logins	Number of Processes	Top CPU Process
Top I/O Process	Top Memory Process	

TIP: To open the Processes tab, right-click any Details menu item.

NOTE: The statistics available in this section depend on the platform of operating system.

Number of Logins

This statistic displays the total number of logins on the server.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Workload Analysis pane

Number of Processes

This statistic displays the total number of processes on the server.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Disk Analysis pane

>

Top CPU Process

Top CPU Process shows the current process that is consuming the most amount of CPU. The information displayed is dependent on the platform of the operating system. Information displayed includes the name of the process, process ID, amount of memory consumed expressed in KB, amount of CPU expressed as a percentage, the amount of Major Page Faults, and the amount of I/O expressed in KB/sec.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Disk Analysis pane
Metrics	If the amount of CPU time used by this process is close to 100% and the CPU usage is very high, this process may be the bottleneck on the server.
Correction	Investigate the process further to see if it is in an inconsistent state. Also, look at minimum requirements for CPU speed for the process. You may need to upgrade your CPU.

Top I/O Process

The Top I/O Process statistic shows the current process that is consuming the most amount of CPU. The information displayed is dependent on the platform of the operating system. Information displayed includes the name of the process, process ID, amount of memory consumed expressed in KB, amount of CPU expressed as a percentage, the amount of Major Page Faults, and the amount of I/O expressed in KB/sec.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Disk Analysis pane

Top Memory Process

Top Memory Process shows the current process that is consuming the most amount of memory. The information displayed is dependent on the platform of the operating system. Information displayed includes the name of the process, process ID, amount of memory consumed expressed in KB, amount of CPU expressed as a percentage, the amount of Major Page Faults, and the amount of I/O expressed in KB/sec.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Home > Disk Analysis pane
Metrics	If you are running out of memory on the system, this is a quick way to identify the top memory user. If the displayed process is using a significant portion of the total memory, it could be causing the memory issues.

Processes Overview

The Processes Overview of the OS Summary includes the following sections:

Top CPU Process	Top I/O Process	Top Memory Process

CPU Tab

The CPU tab of the OS Detail includes the following sections:

Context Switches/Sec	CPU Events	CPU Utilization
Load Averages		

>

The following statistics, available on this tab, are duplicates of statistics on the OS Home page:

Processor Queue Length	Interrupts/Sec	

NOTE: The sections available on this tab depend on the platform of operating system.

CPU Utilization

The CPU Utilization section includes the following information:

- % Privileged Time
- % User Time

% Privileged Time

The metric shows the percentage of elapsed time that the process threads spent executing code in privileged mode.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > CPU Tab
Metrics	The ideal range should be 0-40% (less than 40% indicates excessive system activity).
Correction	If your CPU consistently runs at less than 40% you may need to upgrade your system to include a faster processor(s).

NOTE: For Windows systems, when a Windows system service is called, the service will often run in privileged mode to gain access to system-private data. Such data is protected from access by threads executing in user mode. Calls to the system can be explicit or implicit, such as page faults or interrupts. These kernel commands, are considered privileged to keep the low-level commands executing and prevent a system freeze. Unlike some early operating systems, Windows uses process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. Some work done by Windows on behalf of the application might appear in other subsystem processes in addition to the privileged time in the process.

% User Time

The % User Time statistic is the percentage of elapsed time the processor spends in the user mode. User mode is a restricted processing mode designed for applications, environment subsystems, and integral subsystems. The alternative, privileged mode, is designed for operating system components and allows direct access to hardware and all memory. The operating system switches application threads to privileged mode to access operating system services. This counter displays the average busy time as a percentage of the sample time.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > CPU Tab
Metrics	If the Privileged Time is high in conjunction with Physical Disk Reads, consider upgrading the disk I/O subsystem.

CPU Events

The CPU Events section includes the following information:

Context Switches/Sec	System Calls/Sec	

NOTE: The statistics available in this section depend on the platform of operating system.

Context Switches/Sec

The Context Switches/Sec section shows the combined rate at which all processors on the computer are switched from one thread to another. Context switches occur when a running thread voluntarily relinquishes the processor, is preempted by a higher priority ready thread, or switches between user-mode and privileged (kernel) mode to use an Executive or subsystem service.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > CPU Tab
Metrics	The ideal range should be between 0-10,000. GA number greater then 10,000 may indicate too many threads contending for resources.

System Calls/Sec

System calls are functions which require the services of the operating system. Usually they involve access to some data that users must not be permitted to corrupt or even change.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > CPU Tab
Correction	Try to find a good ratio between the amount of time the operating system runs (fulfilling system calls and doing process scheduling), and the amount of time the application runs. Your goal should be running 60% to 75% of the time in application mode, and 25% to 40% of the time in operating system mode. If you find that the system is spending 50% of its time in each mode, then you should investigate to determine what is wrong.

Load Averages

The Load Average statistic represents the system load averages over the last 1, 5, and 15 minutes.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > CPU Tab
Metrics	High load averages usually mean the system is being used heavily and the response time is correspondingly slow.

Processes Tab

The Processes tab of the OS Detail page succinctly communicates the general overall performance levels of processes. The columns available in this table depend on the platform of operating system. The table below describes the information available in the table on this tab:

Column	Description
Process	The name of the process.
User	The user of the process.
ID	The ID Process is the unique identifier of this process. ID Process numbers are reused, so they only identify a process for the lifetime of that process.
CPU	The CPU is the percentage of elapsed time that all of process threads used the processor to execution instructions.
User Mode	The User Mode is the percentage of elapsed time that the process threads spent executing code in user mode.
Memory **WINDOWS ONLY**	Memory is the current size, in bytes, of the virtual address space the process is using. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is finite, and the process can limit its ability to load libraries.

Column	Description
Memory (MB)	Memory is the current size, in bytes, of the virtual address space the process is using. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is finite, and the process can limit its ability to load libraries.
Memory	Memory is the percentage of the memory used of the total memory.
Active Memory	Active Memory is the amount of committed virtual memory, in bytes for this process. Active memory is the physical memory which has space reserved on the disk paging file(s). There can be one or more paging files on each physical drive. This counter displays the last observed value only; it is not an average.
I/O Data	The rate at which the process is reading and writing bytes in I/O operations. This counter counts all I/O activity generated by the process to include file, network and device I/Os.
Elapsed Time	The total elapsed time, in seconds, that this process has been running.
Thread Count	The number of threads currently active in this process. An instruction is the basic unit of execution in a processor, and a thread is the object that executes instructions. Every running process has at least one thread.
Handle Count	The total number of handles currently open by this process. This number is equal to the sum of the handles currently open by each thread in this process.
Priority	The current base priority of this process. Threads within a process can raise and lower their own base priority relative to the process' base priority.
Creating Proc ID	The Creating Process ID value is the Process ID of the process that created the process. The creating process may have terminated, so this value may no longer identify a running process.
Page Faults/Sec	Page Faults/Sec is the rate at which page faults by the threads executing in this process are occurring. A page fault occurs when a thread refers to a virtual memory page that is not in its working set in main memory. This may not cause the page to be fetched from disk if it is on the standby list and hence already in main memory, or if it is in use by another process with whom the page is shared.
Page File	Page File is the current number of kilobytes that this process has used in the paging file(s). Paging files are used to store pages of memory used by the process that are not contained in other files. Paging files are shared by all processes, and the lack of space in paging files can prevent other processes from allocating memory.
Private	Private is the current size, in kilobytes, of memory that this process has allocated that cannot be shared with other processes.

I/O Tab

The I/O tab includes the following sections:

Average Disk Queue/Total Disk Queue	Bytes per I/O Operation	Disk I/O Time
Disk Transfers/Sec	I/O Details	Transfer Rate

NOTE: The sections available on this tab depend on the platform of operating system.

>

I/O Details

The table below describes the information available in this section:

Column	Description
Disk	The disk number assignment.
Reading (KB/s)	The amount of bytes read from the device.
Writing (KB/s)	The amount of bytes written to the device.
Disk Read Time	Disk Read Time is the percentage of elapsed time that the selected disk drive was busy servicing read requests.
Disk Write Time	Disk Write Time is the percentage of elapsed time that the selected disk drive was busy servicing write requests.
Disk Time	Disk Time is the percentage of elapsed time that the selected disk was busy servicing requests.
Avg. Read Queue	Avg. Disk Read Queue Length is the average number of read requests that were queued for the selected disk during the sample interval.
Avg. Write Queue	Avg. Disk Write Queue Length is the average number of write requests that were queued for the selected disk during the sample interval.
Disk Reads/Sec	Disk Reads/Sec is the rate of read operations on the disk.
Disk Writes/Sec	Disk Writes/Sec is the rate of write operations on the disk.

NOTE: The columns available in this table depend on the platform of operating system.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > I/O Tab

Transfer Rate

The Transfer Rate section succinctly communicates the general overall performance levels of the server's I/O. The table below describes the statistics in this section:

Statistic	Description	Metrics
Reading	The average number of bytes transferred from the disk during read operations.	It is measured in number of KBs per second.
Writing	The average number of bytes transferred to the disk during write operations.	It is measured in number of KBs per second.
Transferred	The amount of time required to transfer the data between the device's cache and the actual device media.	N/A

NOTE: The statistics available in this section depend on the platform of operating system.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > I/O Tab

>

Bytes per I/O Operation

The Bytes per I/O Operation section of the OS Detail succinctly communicates the general overall performance levels of the server's I/O. The table below describes the statistics in this section:

Statistic	Description	Metrics
Reading	The average number of bytes transferred from the disk during read operations.	It is measured in number of bytes per second.
Writing	The average number of bytes transferred to the disk during write operations.	It is measured in number of bytes per second.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > I/O Tab

Average Disk Queue/Total Disk Queue

The Average Disk Queue section of the OS Detail page succinctly communicates the general overall performance levels of the server's I/O. The table below describes the statistics in this section:

Statistic	Description	Information
Average Read Queue	Average Read Queue Length is the average number of read requests that were queued for a particular disk.	Sustained loads on any of the disk queue lengths could represent problems within your system.
Average Write Queue	Average Write Queue is the average number of write requests that were queued for a particular disk.	Sustained loads on any of the disk queue lengths could represent problems within your system. Contact your Windows System Administrator for assistance in troubleshooting these problems.
Average Queue	Average time in ms. to service I/O requests which for disks includes seek, rotational latency and data transfer time).	Sustained loads on any of the disk queue lengths could represent problems within your system. Contact your Unix System Administrator for assistance in troubleshooting these problems.

NOTE: The statistics available in this section depend on the platform of operating system.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > I/O Tab
Correction	Contact your Windows System Administrator for assistance in troubleshooting these problems.

Disk I/O Time

The Disk I/O Time section of the OS Detail page succinctly communicates the general overall performance levels of the disk's I/O. The statistics available in this section depend on the platform of operating system. The table below describes the statistics in this section:

Statistic	Description	Metrics
% Disk Read Time	% Disk Read Time is the percentage of time during the sample interval that the disk is busy servicing read requests.	Use this in conjunction with the % Disk Write Time to identify the disks that may be running at capacity and may require upgrades. This metric may also be used to balance the load between under/over utilized disks.

>

Statistic	Description	Metrics
% Disk Write Time	% Disk Write Time is the percentage of elapsed time that the selected disk drive was busy servicing write requests.	Use this in conjunction with the % Disk Read Time to identify the disks that may be running at capacity and may require upgrades. This metric may also be used to balance the load between under/over utilized disks.
% Disk Time	% Disk Time is the percentage of time during the sample interval that the disk is busy servicing requests.	Use this statistic to identify the disks that may be running at capacity and may require upgrades. This metric may also be used to balance the load between under/over utilized disks.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > I/O Tab

Disk Transfers/Sec

The Disk Transfers/Sec section of the OS Detail page succinctly communicates the general overall performance levels of the disk transfers. The table below describes the statistics in this section

Statistic	Description
Reads/Sec	The the rate of read operations on the disk.
Writes/Sec	The rate of write operations on the disk.
Transfers/Sec	The rate of transfer operations per second.

NOTE: The statistics available in this section depend on the platform of operating system.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > I/O Tab

Memory Tab

The Memory tab of the OS Detail page includes the following sections:

Buffer Size	Cache Efficiency	Cache Size
Memory Available	Paging Activity	Page Replacements

The following statistic, available on this tab, duplicates a statistic on the OS Memory Tab.

- Page Faults/Sec

NOTE: The sections available on this tab depend on the platform of operating system.

>

Paging Activity

The Paging Activity section includes the following statistics:

Blocks Input/Sec	Blocks Output/Sec	Paged In
Paged Outt	Pages Input/Sec	Pages Output/Sec

NOTE: The statistics available in this section depend on the platform of operating system.

Blocks Input/Sec

The Blocks Input/sec statistic is the number of blocks sent to a block device per second.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab
Correction	Although it never hurts to have as much physical memory as your system can handle, here are two things you can check within your system to alleviate the memory bottleneck: *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your swap file. Normally, the bigger the initial size of your swap file, the better, in performance terms.*

Blocks Output/Sec

The Blocks Output/sec statistic is the number of blocks received from a block device per second).

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab
Correction	Although it never hurts to have as much physical memory as your system can handle, here are two things you can check within your system to alleviate the memory bottleneck. *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your swap file. Normally, the bigger the initial size of your swap file, the better, in performance terms.*

Paged In

The number of disk block paged into core from disk which occurs when the server is experiencing a shortage of RAM. Any nonzero value indicates excessive swap activity. In any virtual memory environment, it is not uncommon to see RAM pages moved out to a swap disk. This is a special disk area in UNIX that's reserved for holding memory pages so that the processor is capable of addressing RAM in excess of its full capability. While page-out of operations are a normal part of any server's operations, page-in operations indicate that the real amount of available RAM has been exceeded and that additional RAM pages are required on the server.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab
Correction	High values indicate the possible need to increase the amount of RAM on the system.

>

Paged Out

The number of disk block paged out of core from disk. In any virtual memory environment, it is not uncommon to see RAM pages moved out to a swap disk. This is a special disk area in UNIX that's reserved for holding memory pages so that the processor is capable of addressing RAM in excess of its full capability. While page-out of operations are a normal part of any server's operations, page-in operations indicate that the real amount of available RAM has been exceeded and that additional RAM pages are required on the server.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab

Pages Input/Sec

The Pages Input/Sec statistic is the number of pages read from disk to resolve hard page faults. Hard page faults occur when a process requires code or data that is not in its working set or elsewhere in physical memory, and must be retrieved from disk.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab
Metrics	This value was designed as a primary indicator of the kinds of faults that cause system-wide delays. It includes pages retrieved to satisfy faults in the file system cache (usually requested by applications) and in non-cached mapped memory files. This counter counts numbers of pages, and can be compared to other counts of pages, such as Memory: Page Faults/sec, without conversion. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
Correction	Although it never hurts to have as much physical memory as your system can handle, there are some things you can check within your system to alleviate the memory bottleneck. *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your page file. Normally, the bigger the initial size of your page file, the better, in performance terms.*

Pages Output/Sec

The Pages Output/Sec statistic is the number of pages written to disk to free up space in physical memory. Pages are written back to disk only if they are changed in physical memory. A high rate of pages output might indicate a memory shortage.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab
Metrics	Windows NT writes more pages back to disk to free up space when low in physical memory. This counter counts numbers of pages, and can be compared to other counts of pages, without conversion. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.
Correction	Although it never hurts to have as much physical memory as your system can handle, there are some things you can check within your system to alleviate the memory bottleneck. *Check to see if you have any drivers or protocols that are running but not being used. They use space in all memory pools even if they are idle.* *Check to see if you have additional space on your disk drive that you could use to expand the size of your page file. Normally, the bigger the initial size of your page file, the better, in performance terms.*

Memory Available

The Memory Available section of the OS Detail page succinctly communicates the general overall performance levels of the server's memory. The statistics available in this section depend on the platform of operating system. The table below describes these statistics.

Statistic	Description	Metrics/Correction
Physical Memory	The Physical Memory statistic is the amount of committed virtual memory, in megabytes. Committed memory is where memory space has been reserved in the paging file if it needs to be written to disk.	This counter displays the last observed value only and not an average.
Physical	The Physical statistic is the total physical memory available.	N/A
Paged Memory	The Paged Memory statistic is the amount of committed virtual memory, in bytes. Used Memory is the physical memory which has space reserved on the disk paging file(s). There can be one or more paging files on each physical drive.	This counter displays the last observed value only; it is not an average.
Swap Memory	The Swap Memory statistic is the free, unreserved swap space.	Increase the size of your swap file or add additional files if you consistently see a shortage of free, unreserved swap space.
Swap	The Swap statistic is the total swap memory being used.	N/A
Total Physical	The Total Physical statistic is the total physical memory available.	N/A
Total Swap	The Total Swap statistic is the total swap memory available.	N/A
Swap-ins	The number of memory pages paged in from swapspace.	If the system is consistently low on free memory (over a 30 second average), the memory scheduler will start to swap out processes. Increase the size of your swap file or add additional files.
Swap-outs	The number of memory pages paged in from swapspace.	If the system is consistently low on free memory (over a 30 second average), the memory scheduler will start to swap out processes. Increase the size of your swap file or add additional files.
Buffers	The Buffers statistic represents the amount of buffers used by the kernel in kilobytes.	N/A
Page Faults/Sec	The Page Faults/Sec statistic is the overall rate faulted pages are handled by the processor. It is measured in numbers of pages faulted per second. A page fault occurs when a process requires code or data that is not in its working set. This counter includes both hard faults and soft faults.	This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval. If the number of page faults remains consistently high, you can check with your Windows System Administrator for further investigation. Often, large numbers of page faults are not a problem so long as they are soft faults. However, hard faults, that require disk access, can cause delays.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab

>

Buffer Size

The Buffer Size statistic represents the amount of cache used by the kernel in kilobytes.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab

Cache Size

The Cache Size statistic represents the amount of cache used by the kernel in kilobytes.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab

Page Replacements

The following statistics are available in this section:

Memory Freed (Pages/sec)	Clock Algorithm Scans (Pages/sec)	

Memory Freed (Pages/sec)

Freed or destroyed (Kb/s).

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab

Clock Algorithm Scans (Pages/sec)

The VMM uses a technique known as the clock algorithm to select pages to be replaced. This technique takes advantage of a referenced bit for each page as an indication of what pages have been recently used (referenced). When the page-stealer routine is called, it cycles through the PFT, examining each page's referenced bit.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab

Cache Efficiency

The Cache Efficiency section of the Memory tab succinctly communicates the general overall performance levels of the server's memory. The following statistics are available in this section:

Copy Read Hits %	Data Map Hits %	MDL Read Hits %
Pin Read Hits %		

Copy Read Hits %

The Copy Read Hits % statistic is the percentage of cache copy read requests that hit the cache and does not require a disk read to provide access to the page in the cache.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab
Metrics	When the page is pinned in the memory, the page's physical address in the file system cache will not be altered. A copy read is a file read operation where a page in the cache is copied to the application's buffer. Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.

>

Correction	Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate

Data Map Hits %

The Data Map Hits % statistic is the percentage of data maps in the file system cache that could be resolved without having to retrieve a page from the disk.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab
Metrics	Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.
Correction	Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate

MDL Read Hits %

The MDL Read Hits % statistic is the percentage of Memory Descriptor List Read requests to the file system cache that hit the cache and does not require disk access to provide memory access to the pages in the cache.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab
Metrics	Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.
Correction	Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate.

Pin Read Hits %

The Pin Read Hits % statistic is the percentage of pin read requests that hit the file system cache and does not require a disk read in order to provide access to the page in the file system cache.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Memory Tab
Metrics	Because this value reflects hits, it ideally should be close to 100%. An amount below 100% indicates misses.
Correction	Adding physical memory to a server results in a larger file system cache, which is generally more efficient. Defragmenting disks also helps, putting related pages in the cache together and thereby improving the cache hit rate.

Space Tab

The Space tab of the OS Detail page includes the following sections:

- Device Summary
- Device Detail

Disk Space Free

The Disk Space Free metric displays the amount of free disk space in megabytes.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Space Tab

Disk Space Detail

The Disk Space Detail section of the Space tab succinctly communicates the general overall performance levels of the server's disks and space allotment. The table below describes the statistics in this section:

Statistic	Description
Partition	The drive letter of the disk.
Local Filesystem	The name of the file system.
Type	The type of file system.
Total Space	Total size of the disk/device's capacity expressed in MBs.
Used Space	Amount of MBs currently allocated on the particular disk/device.
Free Space	Amount of MBs currently unallocated and free on the particular disk/device.
Capacity	The percentage of space used on the device.
Mounted On	The mount point of the device.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Space Tab

Device Summary

The Device Summary section of the Space tab displays a bar chart for all devices. The Device Summary section succinctly communicates the general overall performance levels of the server's disks and space allotment. The table below describes the statistics in this section:

Statistic	Description
Used	Amount of MBs currently allocated on the particular disk/device.
Free	Amount of MBs currently unallocated and free on the particular disk/device.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Space Tab

>

Device Detail

The Device Detail section of the Space tab succinctly communicates the general overall performance levels of the server's disks and space allotment. The table below describes the statistics in this section:

Statistic	Description
Partition	The drive letter of the disk.
Local Filesystem	The name of the file system.
Type	The type of file system.
Total Space	Total size of the disk/device's capacity expressed in MBs.
Used Space	Amount of MBs currently allocated on the particular disk/device.
Free Space	Amount of MBs currently unallocated and free on the particular disk/device.
Capacity	The percentage of space used on the device.
Mounted On	The mount point of the device.

NOTE: The statistics available in this section depend on the platform of operating system.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Space Tab

Network Tab

The Network tab of the OS Detail page succinctly communicates the general overall performance levels of the server's networking. The Network tab of the OS Detail page includes the following sections:

Network Details	Packet Collisions	Packet Discards
Packet Errors	Transmission Rate	Transmission Rate (Bytes)
Transmission Queue Length		

NOTE: The sections available on this tab depend on the platform of operating system.

Network Details

The statistics available in this section depend on the platform of operating system. The table below describes the information available in this section:

Column	Description
Network Interface	The name of network interface.
INET Address/Address	The IP address assigned to the network interface.
Pkts Sent/Sec	The number of packets sent per second.
Pkts Received/Sec	The number of packets received per second.
Sent (KB/Sec)	The number of bytes sent per second.
Received (KB/Sec)	The number of bytes received per second.
Out Pkts Discarded	The number of outbound packets discarded.
In Pkts Discarded	The number of inbound packets discarded.
Out Pkt Errors	The number of outbound packet errors.

>

Column	Description
In Pkt Errors	The number of inbound packet errors.
Queue Length	The queue length.
Collisions	The number of collisions.
Packets Discarded	The number of packets discarded.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Network Tab

Transmission Rate

The Transmission Rate section of the Network tab succinctly communicates the packet transmission rate. The following statistics are available in this section:

Packets Sent/Sec	The Packets Sent/Sec statistic is the number of packets sent over each network adapter per second.
Packets Received/Sec	The Packets Received/Sec statistic is the number of packets received over each network adapter per second.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Space Tab

Transmission Rate (Bytes)

The Transmission Rate (Bytes) section of the Network tab succinctly communicates the packet transmission rate. The following statistics are available in this section:

Sent (KB/Sec)	The Sent (KB/Sec) statistic is the rate at which bytes are sent over each network adapter including framing characters
Received (KB/Sec)	The Received (KB/Sec) statistic is the rate at which bytes are received over each network adapter, including framing characters.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Space Tab

Transmission Queue Length

The number of pending outgoing packets on either the software transmit queue or the hardware transmit queue

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Space Tab

>

Packet Collisions

A collision occurs when two devices on the network sense that the network is idle and end up trying to send data at the same time. Some collisions are normal on a repeated network, but excess collisions can cause serious performance problems.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Space Tab

Correction

Contact your network administrator if there is high rate of persistent Packet Collisions.

Packet Discards

The Packet Discards section of the Network tab succinctly communicates the packet discard rate. The following statistics are available in this section:

Outbound Packets Discarded	The Outbound Packets Discarded statistic is the number of outbound packets that were discarded even though no errors had been detected to prevent from being transmitted.
Inbound Packets Discarded	Received Packets Discarded is the number of received packets that were discarded even though no errors had been detected to prevent from being transmitted.
Packet Discards	The Packet Discards statistic represents the number of network packets discarded.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Space Tab
Correction	A possible reason for discarding such a packet could be to free up buffer space. Contact your network administrator if there is high rate of persistent packet discards.

Packet Errors

The Packet Errors section of the Network tab succinctly communicates the packet error rate. The following statistics are available in this section:

Outbound Packet Errors	The Outbound Packet Errors statistic is the outbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.
Inbound Packet Errors	The Packets Received Errors statistic is the received packets that contained errors preventing them from being deliverable to a higher-layer protocol.

Location	IBM DB2 (W/U/L) Performance Analyst statistics > OS > Space Tab

Top SQL

One or two bad queries can cause a lot of trouble for the remaining sessions in a database, and so can one or two bad queries. It is important to find them before they get into a production system, but sometimes a few sneak through.

By applying custom filters and performance-related thresholds, the Top SQL view locates inefficient SQL. By applying filters to certain I/O and statistical counters, you hope to isolate queries that far exceed their nearest competitors in the same area (like disk reads). When you find them, you should reduce the number of sorts a query performs. Or, for a query that returns only a few records, you should try to minimize the number of rows a query processes.

Metrics

When you begin to look for inefficient SQL in a database, there are two primary questions you need to answer:

- What has been the worst SQL that's historically been run in my database?
- What is the worst SQL that's running right now in my database?

When troubleshooting a slow system, you should be on the lookout for any query that shows an execution count that is significantly larger than any other query on the system. It may be that the query is in an inefficient loop, or other problematic programming construct. Only by bringing the query to the attention of the application developers will you know if the query is being mishandled from a programming standpoint.

The following tabs are available on the Top SQL page:

Summary tab	Criteria tab	SQL Details tab
Dynamic SQL Cache tab		

Summary tab

The Summary tab is divided into two sections:

- Top SQL Statements
- Top SQL Details

Top SQL Statements

A lot of a database's overall performance can be attributed to SQL statement execution. Poorly optimized SQL statements or stored procedures can drag an otherwise well-configured database down in terms of user response times.

Before you can identify problem SQL in your database, you have to ask the question of what 'bad SQL' is. What criteria do you use when you begin the hunt for problem SQL in your critical systems? Understand that even the seasoned experts disagree on what constitutes efficient and inefficient SQL; so there is no way to sufficiently answer this question to every SQL Server professional's satisfaction. The Top SQL results grid displays the top SQL statements across various performance categories including Reads, Writes, CPU statistics, and Elapsed Time.

Top SQL Details

This section shows the details of the top SQL statements. The columns presented here include:

Usage: % of all operations of this type performed by the selected sql statement.

SQL text: Text of the selected SQL statement.

Total Reads: Total number of read operations performed by the selected SQL statement.

Total Writes: Total number of write operations performed by the selected SQL statement.

Executions: Number of times the selected SQL statement was executed.

First Start Time: First time the selected SQL statement was executed.

>

Avg CPU Time: Avg CPU time used by the selected SQL statement.

Avg Elapsed Time: Avg Elapsed time for selected SQL statement.

Criteria tab

The Top SQL view displays all SQL running or executed in the current datasource. You can use the lists and boxes to reorder the display according to your specific needs. The Criteria tab of the Top SQL page includes the following sections that you can mix and match to give you the information you want to find.

- Demographic Criteria
- Statistical Criteria
- Sort Criteria

SQL Details tab

The SQL Details tab presents the following sections:

Top SQL Results	Selected SQL Text	Selected SQL Last Snapshot Details

Top SQL Results

This section shows the details of the SQL statements and (their metrics) captured at each refresh of the performance analyst. The information presented here includes

SQL text: Partial text of the captured SQL statement. The full text is shown in the selected SQL text section.

Agents Top: The maximum number of agents that were used when executing the selected statement at the time of last refresh

Max Agents Top: the maximum number of agents that were used when executing the selected statement over all captured executions

Fetch Count: Number of successful fetches performed for the specific statement at the time of last refresh

Avg Fetch Count: Avg number of successful fetches for the selected statement over all its captured executions

Int. Rows Deleted: Number of internal rows deleted for the selected statement as a result of internal activity

Avg Int Rows Deleted: over all its captured executions

Int Rows Inserted

Avg Int Rows Inserted over all its captured executions

Int. Rows Updated

Avg Internal Rows Updated over all its captured executions

#Agents

#Rows Estimate

Avg # Rows Estimate over all its captured executions

Cost Estimate

Avg Cost Estimate over all its captured executions

Rows Read

Avg Rows Read over all its captured executions

Rows Written

Avg Rows Written over all its captured executions

Sorts

Avg Sorts over all its captured executions

Sys CPU Time

Avg Sys CPU Time over all its captured executions

Usr CPU Time

A**vg Usr CPU Time** over all its captured executions

Total Sort Time

Avg Total Sort Time over all its captured executions

Buffer pool Hit Ratio

Overall Buffer pool Hit Ratio over all its captured executions

Temp Data Hit Ratio

Overall Temp Data Hit Ratio over all its captured executions

Selected SQL Text

The Selected SQL Text window allows you to select any of the Top SQL result rows and get the complete SQL text. From here you can open an Explain Plan or Tune the SQL.

Selected SQL Last Snapshot Details

This displays attributes for the last snapshot of the selected SQL statement

Dynamic SQL Cache tab

The Dynamic SQL Cache tab of the Top SQL view, displays the following sections:

Dynamic SQL Cache Details	Dynamic SQL Statement	Execution and Preparation Details
Bufferpool Analysis	Execution Time Details	

Dynamic SQL Cache Details

The DB2 statement cache stores packages and statistics for frequently used SQL statements. This section shows the statements in this cache and the metrics associated with each statement.

Metrics

By examining the contents of this cache, you can identify the dynamic SQL statements that are most frequently executed and the queries that consume the most resources. Using this information, you can examine the most commonly executed and most expensive SQL operations, to determine if SQL tuning could result in better database performance.

Dynamic SQL Statement

The Dynamic SQL Text window allows you to select any of the Dynamic SQL Cache Detail and get the complete SQL text. From here you can open an Explain Plan or Tune the SQL.

Execution and Preparation Details

Here you see the number of times the selected SQL statement was compiled and executed.

Metrics

You can use the execution information to identify the most commonly executed SQL statements in your system. You will probably see the greatest performance benefit from tuning them.

You can use the compilation information to determine the number of different compilations for the selected SQL statement. Use this information together with the number of executions to determine if a bad compilation environment is skewing the metrics.

Bufferpool Analysis

This section shows the various bufferpool usage metrics for the selected SQL statement.

Metrics

This information helps you determine the scope when you tune bufferpools.

Execution Time Details

The information presented here may be used to identify the longest running SQL.

Metrics

The longest running SQL statements together with the most frequently used SQL statements help you identify the most likely candidates for SQL tuning.

Space Analyst

Space management plays a key role in a database's uptime and overall performance. Because of this, it's important that a database administrator understand the various types of space bottlenecks that can occur in a database.

Embarcadero Space Analyst provides sophisticated diagnostic capabilities to help you effectively locate and troubleshoot bottlenecks and performance inefficiencies that result from the challenges of space management. Running inside the Embarcadero DBArtisan console, Space Analyst automates all diagnostic and correction procedures so any space-related headache can be quickly detected and corrected before it reaches critical mass.

The table below describes the major sections of the help available for the Space Analyst.

Section	Description
Welcome to Space Analyst	Provides information about opening, and licensing Space Analyst.
Space Analyst Product Design	Provides conceptual information about Space Analyst's product design and navigation.
Using Space Analyst	Shows how to use each component of the application. Using Space Analyst breaks each process and procedure into simple step-by-step instructions.
Space Analyst for Oracle Tutorial	Helps you get started with Embarcadero Space Analyst. Provides the foundation you need to explore the many features and benefits of Space Analyst.
Space Analyst for Microsoft SQL Server Tutorial	Helps you get started with Embarcadero Space Analyst. Provides the foundation you need to explore the many features and benefits of Space Analyst.
Space Analyst for Sybase ASE Tutorial	Helps you get started with Embarcadero Space Analyst. Provides the foundation you need to explore the many features and benefits of Space Analyst.
Space Analyst for IBM DB2 for Linux, Unix, and Windows Tutorial	Helps you get started with Embarcadero Space Analyst. Provides the foundation you need to explore the many features and benefits of Space Analyst.
Space Analyst Expert Guide	Provides detailed information on each tab of the Space Analyst and discusses fragmentation concepts.

>

Welcome to Space Analyst

Space Analyst provides sophisticated diagnostic capabilities to troubleshoot bottlenecks and performance problems that result from poor space management. Space Analyst helps ensure that database storage is being used as efficiently as possible.

Space Analyst also includes a Reorganization Wizard as part of the Space Management Wizard that handles full storage structure and object reorganization needs.

For more information, see:

- What You Should Know Before Installation
- Opening an Analyst
- Licensing Space Analyst
- Using Space Analyst
- Using Help
- Embarcadero Technologies Technical Support

What You Should Know Before Installation

When you install the Analyst, several Embarcadero-generated stored objects are pushed onto the database server.

These objects are as follows:

SQL Server

Procedures

sp_embt_sa_database_map	sp_embt_sa_database_overview	sp_embt_sa_dbbottlenecks
sp_embt_sa_filegroups	sp_embt_sa_files	sp_embt_sa_index_space_general
sp_embt_sa_index_space_reorg	sp_embt_sa_index_space_reorg_summary	sp_embt_sa_server_overview
sp_embt_sa_table_space_general	sp_embt_sa_table_space_reorg	sp_embt_sa_table_space_reorg_summary
sp_embt_sa_index_reorg_summary	sp_embt_sa_table_reorg_summary	sp_embt_sa_version

SYBASE

Procedures

sp_embt_sa_dbspace	sp_embt_sa_device_summary	sp_embt_sa_device_detail
sp_embt_sa_database_overview	sp_embt_sa_db_obj_space	sp_embt_sa_database_detail
sp_embt_sa_db_frag_summary	sp_embt_sa_index_diag_summary	sp_embt_sa_index_diag
sp_embt_sa_server_overview	sp_embt_sa_table_diag_summary	sp_embt_sa_table_diag
sp_embt_sa_version		

NOTE: No procedures are installed in Space Analyst for Oracle or DB2.

Opening an Analyst

To open the Space Analyst, do the following:

1 On the **Analyst** menu, click **Space Analyst**.

NOTE: The first time that you open the Space Analyst, DBArtisan alerts you that you need to install the Analyst's procedures. When you install the Analyst, several procedures and tables are pushed onto the database server. For a list of these procedures and tables, see What You Should Know Before Installation.

Uninstalling a Repository

To uninstall an Analyst repository, do the following:

1 On the **Command Menu**, select **Uninstall Repository**.

The Analyst Repository Uninstallation Wizard opens.

Analyst Repository Uninstallation Wizard - Panel 1

The table below describes the options and functionality on this panel of the Analyst Repository Uninstallation Wizard:

Option	Description
Administrator Login Information	Displays the login name, password, and login as information.
Database to hold the procedures	Displays the database that stores the procedures.
Uninstall	Uninstalls the repository and opens Panel 2. For more information, see Analyst Repository Uninstallation Wizard - Panel 2.

Analyst Repository Uninstallation Wizard - Panel 2

The table below describes the options and functionality on this panel of the Analyst Repository Uninstallation Wizard:

Option	Description
Repository Login	Displays repository login.
Execution Log	Displays the repository execution log and links to ETRepoLog.txt.
Execution Error(s)	Displays any errors.
Finish	Closes wizard.

Licensing Space Analyst

The Space Analyst is licensed as a component add-on to DBArtisan. You are issued one license key for the number of databases/servers you want to license. For example, if you purchase ten licenses for Space Analyst, you only need to obtain one license key from Embarcadero.

For more information, see Space Analyst License Manager.

Space Analyst License Manager

The Analyst License Manager lets you manage licenses.

The table below describes the options and functionality on the Analyst License Manager:

Option	Description
License Summary	Displays basic licensing metadata for the product.
Available Datasources	Lets you select a datasource from the list of datasource available to license.
License	Click to add a license to a datasource.
Licensed Datasource(s)	Displays datasources already licensed.
Remove	Click to remove a license from a datasource.
Apply	Click to make you license choices permanent.
Cancel	Click to close the Licensing Dialog box. **NOTE:** If unsaved license changes have been made, the Analyst asks if you want to discard their changes before exiting.

Completing the Space Analyst License Manager

To complete the Analyst License Manager, do the following:

1. Open DBArtisan.
2. On the **Help** menu, click **Licensing**.

   DBArtisan opens the License Configuration Dialog box.
3. Click **License**.

   DBArtisan opens the Analyst License Manager.
4. Select options.
5. To make your license choices permanent, click **Apply**.
6. To close the Analyst License Manager, click **Cancel**.

For more information, see Space Analyst License Manager.

Space Analyst Add an Analyst License Dialog Box

The Add an Analyst License Details dialog box lets you add a license. The table below describes the options and functionality on the Add a Analyst License Details dialog box:

Option	Description
Select a DBMS platform	Lets you select the DBMS platform.

Using Help

DBArtisan provides comprehensive, context-sensitive Help. Help includes detailed information on using the application. DBArtisan Help includes the following standard conventions:

>

Fonts

The table below describes the conventions in DBArtisan Help to distinguish elements of text

Font	Description
7	Sequential list of steps to complete a task.
Result Font	Result of an action in the previous step.
Bold Font	Elements of the application, including dialog box and window titles, menus, commands on menus, dialog box options, buttons, icons.
Italic Font	Variables.
`Monospace Font`	Code samples.
UPPERCASE Font	Keyboard combinations. SQL keywords and elements.
Underlined Blue Font	Hyperlinked words and phrases. Click to jump to a new location with related information.

Opening DBArtisan Help

Although the most common way to access Help is through context-sensitive Help, DBArtisan also offers three additional ways to access Help:

1. On the **Help** menu, click **Help Topics**.
2. On the Windows Start button, point to **Programs**, **Embarcadero DBArtisan**, and then click **DBArtisan Help**.
3. On the Windows desktop, open the **Embarcadero DBArtisan** folder, and then double-click **Help**.

> **TIP:** To toggle between Help and the application, on the Windows Start button, open DBArtisan Help, and then press ALT+TAB.

For more information, see Using DBArtisan Help.

Using DBArtisan Help

Help is a customizable window that lets you access DBArtisan Help. The table below describes the tabs in the DBArtisan Help window:

Tab	Description
Contents	The Table of Contents for the Help system.
Index	A comprehensive index of DBArtisan Help.
Search	A search engine for every word in the DBArtisan Help. DBArtisan Help features an advanced full-text search, letting you search using Boolean, wildcard, and nested expressions.
Favorites	A personalized list where you can access your favorite DBArtisan Help topics. DBArtisan saves this list so it displays every time you open DBArtisan Help.

Embarcadero Technologies Technical Support

Space Analyst is fully supported by the Embarcadero Technologies Technical Support team. The Technical Support team is available to assist you with any problems you have with this application. When contacting Embarcadero Technologies Technical Support please provide the following:

- Product name.
- Product version.
- Operating system (server and client).
- Database platform in which the error occurs (including version).
- Error message (exact text or a screen shot if possible).
- Detailed description of the problem including steps to reproduce.

Every Embarcadero Technologies application includes sixty days free technical support. We strongly recommend that you purchase a maintenance contract to continue receiving technical support after this period.

Contacting Embarcadero Technologies Technical Support

Embarcadero Technologies offers three ways to contact Embarcadero Technical Support:

1 Web

2 Phone

3 E-mail

Web

Embarcadero Technologies Technical Support on the Web is available through:

- Embarcadero Technologies Web Site
- DBArtisan Help

Embarcadero Technologies Web Site

- Complete a technical support form at the Embarcadero Technologies Technical Support Web page.

DBArtisan Help

To open Embarcadero Technologies Technical Support on the Web using DBArtisan:

1 On the **Help** menu, point to **Embarcadero on the Web**, and then click **Online Technical Support**.

DBArtisan connects to the Web and opens the Embarcadero Technologies Technical Support Request Web page.

2 Complete the online form.

3 Click **Submit**.

Phone

The Embarcadero Technologies Technical Support phone number is (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 licensing issues, including Product Unlock Codes, 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., Saturday and Sunday 8 A.M. to 5 P.M. Pacific time.

The Embarcadero Technologies Technical Support fax number is (415) 495-4418.

E-mail

Embarcadero Technologies offers two ways to contact Technical Support via e-mail:

1. Complete a technical support form at the Embarcadero Technologies Technical Support Web page.

   OR

2. Send an e-mail message to support@Embarcadero.com.

   Please include the following in your e-mail message:

   - Product name.
   - Product version.
   - Operating system (client and server).
   - Database platform in which the error occurs (including version).
   - Error message (exact text or a screen shot if possible).
   - Detailed description of the problem including steps to reproduce.

>

Space Analyst Application Basics

The Space Analyst utility assists DBAs with space diagnostics and reorganizations by displaying pertinent reorganization diagnostics and providing a Reorganization Wizard as part of the Space Management Wizard that builds scripts that you can execute or schedule to perform reorganizations.

For more information, see Product Design.

Space Analyst Product Design

The Space Analyst product design includes tabs for Oracle, Microsoft SQL Server and IBM DB2.

Oracle Space Analyst Tabs

The Space Analyst includes the following tabs:

Database Overview Tab	Datafiles Tab	Tablespace Map Tab
Tables Tab	Indexes Tab	

Microsoft SQL Server Space Analyst Tabs

The Space Analyst includes the following tabs:

Storage Overview Tab	Files/File Groups Tab	Map Tab
Databases Tab	Objects Tab	

IBM DB2 Space Analyst Tabs

The Space Analyst includes the following tabs:

Database Overview Tab	Tablespaces Tab	Tablespace Map Tab
Tables Tab	Indexes Tab	

Sybase Space Analyst Tabs

The Space Analyst includes the following tabs:

Storage Overview Tab	Database Devices	Databases Tab
Tables Tab	Indexes Tab	

Database Overview Tab

The Database Overview Tab displays the storage overview for the DB2 datasource.

There are four distinct sections displayed on the first **Database Overview Tab** of Space Analyst:

- Database Overview

- Log Overview (DB2 Only)
- Tablespace Space Overview
- Tablespace Fragmentation Overview
- Tablespace Objects Overview

For more information, see Space Analyst Tabs.

Data Files Tab

The Data Files Tab displays datafile fragmentation information, and displays how much space in each datafile is assigned to tables, indexes, clusters, and rollback segments.

For more information, see Space Analyst Tabs.

Tablespace Map Tab

The Tablespace Map Tab allows for tablespace mapping.

For more information, see Space Analyst Tabs.

Tables Tab

The Tables Tab lets you identify tables with wasted space. The Space Analyst computes a block inefficiency ratio that quickly communicates the severity of the space problem. Space Analyst provides a summary count of all tables in a database that violate a customizable block inefficiency threshold, details on the tables, and the extent of wasted space.

For more information, see Space Analyst Tabs.

Indexes Tab

The Indexes Tab lets you select index reorganization candidates. The tab also offers the option to obtain deleted leaf row information for each index, although such data gathering can take some time, especially on large systems.

Space Analyst runs diagnostics against your database and graphically display a count of indexes that are found to contain various types of space problems.

For more information, see Space Analyst Tabs.

Tablespaces Tab

The Tablespaces Tab displays the Tablespace Container allocations and Database Partition Breakdown for each Tablespace.

There are three distinct sections displayed on the **Tablespaces Tab** of Space Analyst:

- Space Allocation by Container
- Space Allocation by Node
- Tablespaces

Using Space Analyst

Using Space Analyst covers how to use each component of the application. The Using chapter breaks each process and procedure into simple step-by-step instructions.

The Space Analyst includes the following tabs, by platform:

Space Analyst for Oracle

Database Overview Tab	Datafiles Tab	Tablespace Map Tab
Tables Tab	Indexes Tab	

Space Analyst for Microsoft SQL Server

Storage Overview Tab	Files/File Groups Tab	Map Tab
Databases Tab	Objects Tab	

Space Analyst for Sybase ASE

The Space Analyst includes the following tabs:

Storage Overview Tab	Database Devices Tab	Databases Tab
Tables Tab	Indexes Tab	

Space Analyst for IBM DB2 for Linux, Unix, and Windows

The Space Analyst includes the following tabs:

Storage Overview Tab	Datafiles Tab	Tablespace Tab
Tables Tab	Indexes Tab	

Space Analyst for Oracle and IBM DB2 for Linux, Unix, and Windows - Database Overview Tab

The Database Overview Tab lets you view the global storage picture of your database.

The table below describes the options and functionality on the Database Overview Tab of Space Analyst for Oracle:

Option	Description
Database Overview	Shows how much space in KB is taken up in the database by various object types (tables, indexes, etc.) along with a pie chart that represents the percentage of each object type. You can switch between the default graphical view and a grid view. **NOTE:** Objects owned by SYS and SYSTEM are not included in this breakdown.
Tablespace Space Overview	When in graphical mode, this section contains a bar chart graph that displays all tablespaces and how much total space each has allocated. When in grid mode, Space Analyst displays detailed information about all tablespaces (status, total, used, free, percent free, etc.) **ORACLE ONLY:** When in grid mode, Space Analyst displays tablespace status. You can switch between the default graphical view and a grid view.

Option	Description
Tablespace Fragmentation Overview	Contains fragmentation information for all tablespaces and displays how badly fragmented a tablespace or datafile is. Higher scores for the tablespace's fragmentation index indicate little or no fragmentation, while lower scores indicate fragmentation. You can switch between graphical and grid views of tablespace fragmentation. If fragmentation is present, you can view detailed object and fragmentation information for the tablespace in the tablespace map. You can correct fragmentation problems via Space Analyst's Reorganization Wizard. You can switch between the default graphical view and a grid view.
Tablespace Objects Overview	Contains a grid that displays how much space in each tablespace is assigned to tables, indexes, clusters, and rollback segments. **ORACLE ONLY:** When in grid mode, Space Analyst displays tablespace LOB data values.

The table below describes the options and functionality on the Database Overview Tab of Space Analyst for IBM DB2 for Linux, Unix, and Windows:

Option	Description
Database Overview	Shows how much space in KB is taken up in the database by various object types (tables, indexes, etc.) along with a pie chart that represents the percentage of each object type. You can switch between the default graphical view and a grid view.
Log Space Overview	Displays all of the log options for the database: Total Used Total Free Block Log When Disk Full Log File Size User Exit Primary Logs Secondary Logs Maximum Secondary Log Space Used Maximum Total Log Space Used
Tablespace Space Overview	When in graphical mode, this section contains a bar chart graph that displays all tablespaces and how much total space each has allocated. When in grid mode, displays detailed information about all tablespaces: Name Type State Used Space (MB) Free Space (MB) Content Type (Any, Long, Temp, System temp) Number of containers **NOTE:** You can switch between the default graphical view and a grid view.
Tablespace Objects Overview	Contains a grid that displays how much space in each tablespace is assigned to each object. Also displays tablespace LOB data values. Displays detailed information about all tablespace objects: Tablespace Name Table - total space used by tables Index - total space used by indexes Extent Size - tablespace extent size Page size - Tablespace page size

>

Available Functionality

The following functionality is available in the Command menu, shortcut menu, and toolbar:

NOTE: Functionality differ by platform.

Analysis Report	Convert to Dictionary-Managed	Convert to Locally-Managed
Extract	Estimate Size	Map

Open Editor. For information on modifying database objects, see Modifying objects using editors.

Open Wizard	Reorganize	Update Statistics
Tablespace Wizard for IBM DB2 (IBM DB2 LUW)	Threshold Options	

For more information, see:

Correcting Tablespace 'Out of Space' Problems

Space Analyst Tabs

Space Analyst for Oracle and IBM DB2 for Linux, Unix, and Windows - Datafiles Tab

The Data Files Tab lets you view fragmentation and object placement information for data files.

The table below describes the options and functionality on the Datafiles Tab of the Space Analyst:

Option	Description
Datafile Fragmentation	Displays fragmentation about datafiles. Higher scores for a datafile's fragmentation index indicate little or no fragmentation, while lower scores indicate fragmentation. If fragmentation is present, you can view detailed object and fragmentation information for the datafile's tablespace in the tablespace map. You can correct fragmentation problems via Space Analyst's Reorganization Wizard. **NOTE:** Datafiles for temporary tablespaces are not included in the Datafile Fragmentation grid.
Datafile Objects Overview	Contains a grid that displays how much space in each datafile is assigned to tables, indexes, clusters, and rollback segments.

TIP: To open a Tablespace Editor, right-click a datafile in any grid, and then select Open. To extract the tablespace DDL in an ISQL editor, right-click a datafile in any grid, and then select Extract.

Available Functionality

The following functionality is available in the Command menu, shortcut menu, and toolbar:

NOTE: Functionalists differ by platform.

Analysis Report	Convert to Dictionary-Managed	Convert to Locally-Managed
Estimate Size	Extract	Map

Open Editor. For information on modifying database objects, see Modifying objects using editors.

>

Reorganize

Threshold Options

For more information, see Space Analyst Tabs.

Space Analyst for Oracle - Tablespace Map Tab

The Tablespace Map Tab lets you view a graphical layout regarding the physical placement of every object inside a selected tablespace. The tablespace map also helps identify fragmentation problems for a tablespace (honeycomb and bubble). An object grid is also present that lets you view detailed information about every object's storage properties.

The Space Info section contains information regarding the total, used, and free space (in MB) for the tablespace.

The View By section allows you to switch between two different tablespace map views:

Option	Description
Object Type	The normal tablespace map mode, which is color-coded by object type. Pressing the legend button shows the various colors assigned to the object types contained within the tablespace.
Extent Problems	When in this mode, the tablespace map highlights in red any object containing one or both of these object extent problems:· - Max extent problem - the object has reached or is near its maximum extent limit. - Extent deficit - the object may not be able to allocate its next extent because the tablespace does not currently have enough contiguous free space. This problem will be highlighted even if the tablespace has its autoextend property for one or more datafiles set to true.

When in Object Type mode, the accompanying grid of the tablespace map contains the following data:

Data	Description
Name	The name of the object in owner.object name.partition name format.
Extents	The total number of allocated extents for the object.
Blocks	The total number of allocated blocks for the object.
Bytes	The total number of allocated bytes for the object.
Type	The object type (table, index, etc.)
Initial Extent	The initial space extent (in bytes) allocated to the object.
Next Extent	The next extent (in bytes) that the object will attempt to allocate when more space for the object is required.
Percent Increase	The percent increase property for the object.

When in Extent Problems mode, the accompanying grid of the tablespace map contains the following data:

Data	Description
Name	The name of the object in owner.object name.partition name format.
Extents	The total number of allocated extents for the object.
Blocks	The total number of allocated blocks for the object.
Bytes	The total number of allocated bytes for the object.
Type	The object type (table, index, etc.)

Data	Description
Initial Extent	The initial space extent (in bytes) allocated to the object.
Next Extent	The next extent (in bytes) that the object will attempt to allocate when more space for the object is required.
Percent Increase	The percent increase property for the object.
Max Extent Percent	A percentage indicating how close the object is to reaching its maximum extent limit. For example, a 100 indicates an object that has reached its maximum extent limit.
Extendibility	Indicates if the object can allocate its next extent. If 'N', this is because the tablespace does not currently have enough contiguous free space. This problem will be highlighted even if the tablespace has its autoextend property for one or more datafiles set to true.

When viewing a tablespace in the graphical map, there are two broad categories of fragmentation that you will want to be on the lookout for: free space honeycombs and free space bubbles. Honeycomb fragmentation is really not difficult to detect or handle and occurs when two free space extents in a tablespace reside next to one another. For example, rather than having one large free space extent of 6MB, a tablespace might have two extents that are next to each other (a honeycomb) that make up 3MB and 3MB. If an incoming object required a single 4MB extent allocation, Oracle would have to coalesce the honeycomb to make room for the object. The necessary coalesce operation would cause a small degree of performance degradation, however the situation would be resolved and no space-related bottleneck would remain.

Fragmentation bubbles, however, are another story and are the traditional reason DBAs have had to perform full tablespace reorganizations. Bubbles occur when objects are continuously added and dropped in a tablespace, with the end result being free space pockets (or bubbles) being interspersed between objects that remain in the tablespace. This can become problematic when, for example, an incoming object requires a single 5MB extent allocation, but there is not enough contiguous space in the tablespace to meet the demand. Even if the total amount of free space in the tablespace equals the required 5MB, it will do no good if that 5MB is made up of 5 non-contiguous blocks of free space. In this situation, Oracle has no alternative but to either extend the datafile to make room for the object (if the tablespace's datafile has the AUTOEXTEND property enabled), or return an error and deny the object entry into the tablespace.

Space Analyst makes it easy to recognize both of these bottleneck headaches. When the map is viewed in Object Type mode, pockets of green seen throughout the tablespace indicate bubble fragmentation, while two or more blocks of green segments viewed next to one another in the map indicate honeycomb fragmentation.

The Tablespace Map Tab includes the following buttons:

Button	Description
Legend	Opens the Tablespace Map Legend.
Display	Opens the Display window. The Display window lets you specifies the magnification level for the Tablespace Map.
Fit to Window	Fits the Tablespace Map to the window.
Full Screen	Fits the Tablespace Map to the full screen.
Print Map	Opens the Print dialog box to print the Tablespace Map.

Available Functionality

The following functionality is available in the Command menu, shortcut menu, and toolbar:

Analysis Report	Convert to Dictionary-Managed	Convert to Locally-Managed
Estimate Size	Extract	Open Editor
Open Wizard	Reorganize	Update Statistics

Tablespace Wizard for IBM DB2 (IBM DB2 LUW)	Threshold Options	

For more information, see Space Analyst Tabs.

Space Analyst for Oracle, Sybase, and IBM DB2 for Linux, Unix, and Windows - Tables Tab

The Tables Tab displays diagnostic information regarding table space inefficiencies that may exist in your database.

The Table Reorganization Candidate Summary section uses a bar chart to count the various types of table space inefficiencies in your database. Note that the reorganization threshold options are used to compute the counts for each type of space inefficiency.

The table below describes the four distinct types of table space inefficiencies that Space Analyst computes for Oracle and Sybase:

Problem	Description
Block Inefficiencies	These tables have wasted space that exceeds the allowable percentage of wasted space set in the customizable thresholds of Space Analyst (to see all the thresholds, click the toolbar option labeled Reorganization Threshold Options). Tables with wasted space can take longer to scan and therefore contribute to slower response times.
Next Extent Problems	These tables cannot obtain their next extent of free space due to a lack of needed contiguous free space in the tablespace. Tables suffering from this problem may produce error conditions in the database if more data is added to them. This problem will be highlighted even if the tablespace that contains the table has its autoextend property for one or more datafiles set to true.
Chain Problems	These tables suffer from a chained or migrated row problem. The percentage of chained or migrated rows to overall rows contained in the table exceeds the allowable percentage of chained rows set in the customizable thresholds of Space Analyst (to see all the thresholds, click the toolbar option labeled Reorganization Threshold Options). Because chained/migrated rows cause more I/O to occur when they are accessed, overall response times can be degraded when these tables are referenced in SQL queries.
Max Extent Problems	These tables are nearing their maximum extent limit and have exceeded the allowable maximum extent percentage set in the customizable thresholds of Space Analyst (to see all the thresholds, click the toolbar option labeled Reorganization Threshold Options). Tables suffering from this problem may cause space-related errors to result if more data is added to them.

The Table Reorganization Candidate Details section provides filters and a grid that you can use to view detailed space diagnostics for all or selected tables. The table below describes the grid data:

Data	Description
Table Owner	The owner of the table.
Table Name	The name of the table.
Segment Type	Indicates if the object is a table, table partition, etc.
Table (KB)	The total allocated space of the table in KB.
Number of Rows	The number of rows that the table contains. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.

Data	Description
Number of Blocks	The total number of blocks allocated to the table. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Empty Blocks	The number of empty blocks that the table contains. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Highwater Mark	The last block that Oracle will scan up to when searching the table for data. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Average Used Blocks	The average number of blocks that the table is using. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Block Inefficiency	A percentage that indicates how much wasted space (space that the table is not using) that the table contains. High percentages indicate tables that may be candidates for reorganization. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Chained Row Percent	A percentage that indicates how many rows in the table suffer from chaining or migration. High percentages indicate tables that may be candidates for reorganization. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Max Extent Percent	A percentage indicating how close an object is to reaching its maximum extent limit. High percentages indicate tables that may be candidates for reorganization.
Number of Extents	The number of extents that the table consumes.
Max Extents	The maximum extent limit for the table.
Extent Deficit	Indicates if the object can allocate its next extent. If 'N', this is because the tablespace does not currently have enough contiguous free space. This problem will be highlighted even if the tablespace has its autoextend property for one or more datafiles set to true.
Next Extent	The next extent (in bytes) that the table will attempt to allocate when more space for the table is required.
Max Free Space	The largest contiguous block of free space that exists in the tablespace where the table resides.
Tablespace Name	The tablespace that contains the table.

IBM DB2 for Linux, Unix, and Windows

The table below describes the four distinct types of Reorganization Candidate Summary that Space Analyst computes for IBM DB2 for Linux, Unix, and Windows:

Problem	Description
Overflow Problems	Displays the number of tables that fall outside the preset limits for the thresholds of this area. Can be addressed with reorganization of the affected table(s).
Space Utilization Problem	Displays the number of tables that fall outside the preset limits for the thresholds of this area. Can be addressed with reorganization of the affected table(s).
Page Utilization Problem	Displays the number of tables that fall outside the preset limits for the thresholds of this area. Can be addressed with reorganization of the affected table(s).
Index Contention	Displays the number of tables where the index space is shared by the Tablespace. To resolve, use the Space Management Wizard to change the Index Tablespace in table definition.

>

The table below describes the options and functionality on the Table Reorganization Candidate Details section:

Option	Description
Tablespace	Displays all Tablespaces (regular) in the database and the value [All].
Table Schema	Displays all schema in the database and the value [All].
Display	All Tables Tables with Overflow Problems - Displays a list of all tables where overflow threshold is exceeded. Tables with Space Utilization problems - Displays a list of all tables where space utilization threshold is exceeded. Tables with Page Utilization problems - Displays a list of all tables where the page utilization threshold is exceeded. Tables with Index Contention - Displays a list of all tables where the table and indexes on the table share the same Tablespace.
Refresh	Refreshes the table grid with criteria selected in the above lists.

The Table Reorganization Candidate Details section provides filters and a grid that you can use to view detailed space diagnostics for all or selected tables. The table below describes the grid data:

Data	Description
Table Schema	The schema name for the table.
Table Name	The name of the table.
Cardinality	The number of rows in the table.
Size	The size of Table in bytes.
Overflow	The number of overflow rows. Overflow rows can be created when rows are updated and the new rows contain more bytes than the old ones (VARCHAR fields), or when columns are added to existing tables.
Pages with Data	The number of allocated pages that contain data.
Total Pages	The total number of allocated pages.
Active Blocks	The number of Active Blocks in Tables with Multidimensional clustering).
Overflow %	Percentage of overflow rows. This should typical be a small value (5%).
Space Utilization%	A low space utilization is an indicator of excessive fragmentation (internal and external). This can occur when there are many pages with free space and/or many free pages. This should typically be > 68%.
Page Utilization %	A low page utilization is an indicator of excessive external fragmentation. This value should typically be more than 80%.
Reorg	This is an indicator of how severe the fragmentation problem for the table is. --- means Overflow%, Space Utilization% and Page Utilization % are all with in the thresholds. *-* means that Overflow%, and Page Utilization % are outside the threshold but Space Utilization% is with in the threshold.
Tablespace	Tablespace where the table data resides.
Indexspace	Tablespace where the index data resides.

Available Functionality

NOTE: Functionality differ by platform.

The following functionality is available in the Command menu, shortcut menu, and toolbar:

Analysis Report	Analyze	DBCC

Estimate Size	Extract	Map
Table Wizard for IBM DB2 (IBM DB2 LUW)	Open Editor	Open Wizard
Place	Quick Reorganize	Reorganize
Threshold Options	Update Statistics	For more information, see:
Finding Space-Related Table Problems	Space Analyst Tabs.	

Space Analyst for Oracle - Tables Tab

The Tables Tab displays diagnostic information regarding table space inefficiencies that may exist in your database.

The Table Reorganization Candidate Summary section uses a bar chart to count the various types of table space inefficiencies in your database. Note that the reorganization threshold options are used to compute the counts for each type of space inefficiency.

The table below describes the four distinct types of table space inefficiencies that Space Analyst computes for Oracle and Sybase:

Problem	Description
Block Inefficiencies	These tables have wasted space that exceeds the allowable percentage of wasted space set in the customizable thresholds of Space Analyst (to see all the thresholds, click the toolbar option labeled Reorganization Threshold Options). Tables with wasted space can take longer to scan and therefore contribute to slower response times.
Next Extent Problems	These tables cannot obtain their next extent of free space due to a lack of needed contiguous free space in the tablespace. Tables suffering from this problem may produce error conditions in the database if more data is added to them. This problem will be highlighted even if the tablespace that contains the table has its autoextend property for one or more datafiles set to true.
Chain Problems	These tables suffer from a chained or migrated row problem. The percentage of chained or migrated rows to overall rows contained in the table exceeds the allowable percentage of chained rows set in the customizable thresholds of Space Analyst (to see all the thresholds, click the toolbar option labeled Reorganization Threshold Options). Because chained/migrated rows cause more I/O to occur when they are accessed, overall response times can be degraded when these tables are referenced in SQL queries.
Max Extent Problems	These tables are nearing their maximum extent limit and have exceeded the allowable maximum extent percentage set in the customizable thresholds of Space Analyst (to see all the thresholds, click the toolbar option labeled Reorganization Threshold Options). Tables suffering from this problem may cause space-related errors to result if more data is added to them.

The Table Reorganization Candidate Details section provides filters and a grid that you can use to view detailed space diagnostics for all or selected tables. The table below describes the grid data:

Data	Description
Table Owner	The owner of the table.
Table Name	The name of the table.
Segment Type	Indicates if the object is a table, table partition, etc.
Table (KB)	The total allocated space of the table in KB.

Data	Description
Number of Rows	The number of rows that the table contains. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Number of Blocks	The total number of blocks allocated to the table. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Empty Blocks	The number of empty blocks that the table contains. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Highwater Mark	The last block that Oracle will scan up to when searching the table for data. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Average Used Blocks	The average number of blocks that the table is using. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Block Inefficiency	A percentage that indicates how much wasted space (space that the table is not using) that the table contains. High percentages indicate tables that may be candidates for reorganization. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Chained Row Percent	A percentage that indicates how many rows in the table suffer from chaining or migration. High percentages indicate tables that may be candidates for reorganization. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Max Extent Percent	A percentage indicating how close an object is to reaching its maximum extent limit. High percentages indicate tables that may be candidates for reorganization.
Number of Extents	The number of extents that the table consumes.
Max Extents	The maximum extent limit for the table.
Extent Deficit	Indicates if the object can allocate its next extent. If 'N', this is because the tablespace does not currently have enough contiguous free space. This problem will be highlighted even if the tablespace has its autoextend property for one or more datafiles set to true.
Next Extent	The next extent (in bytes) that the table will attempt to allocate when more space for the table is required.
Max Free Space	The largest contiguous block of free space that exists in the tablespace where the table resides.
Tablespace Name	The tablespace that contains the table.

Available Functionality

NOTE: Functionality differ by platform.

The following functionality is available in the Command menu, shortcut menu, and toolbar:

Analysis Report	Analyze	DBCC
Estimate Size	Extract	Map
Table Wizard for IBM DB2 (IBM DB2 LUW)	Open Editor	Open Wizard

Place	Quick Reorganize	Reorganize
Threshold Options	Update Statistics	

For more information, see Space Analyst Tabs.

Space Analyst - Tables Tab

The Tables Tab displays diagnostic information regarding table space inefficiencies that may exist in your database.

The Table Reorganization Candidate Summary section uses a bar chart to count the various types of table space inefficiencies in your database. Note that the reorganization threshold options are used to compute the counts for each type of space inefficiency.

The table below describes the four distinct types of table space inefficiencies that Space Analyst computes for Sybase:

Problem	Description
Block Inefficiencies	These tables have wasted space that exceeds the allowable percentage of wasted space set in the customizable thresholds of Space Analyst (to see all the thresholds, click the toolbar option labeled Reorganization Threshold Options). Tables with wasted space can take longer to scan and therefore contribute to slower response times.
Next Extent Problems	These tables cannot obtain their next extent of free space due to a lack of needed contiguous free space in the tablespace. Tables suffering from this problem may produce error conditions in the database if more data is added to them. This problem will be highlighted even if the tablespace that contains the table has its autoextend property for one or more datafiles set to true.
Chain Problems	These tables suffer from a chained or migrated row problem. The percentage of chained or migrated rows to overall rows contained in the table exceeds the allowable percentage of chained rows set in the customizable thresholds of Space Analyst (to see all the thresholds, click the toolbar option labeled Reorganization Threshold Options). Because chained/migrated rows cause more I/O to occur when they are accessed, overall response times can be degraded when these tables are referenced in SQL queries.
Max Extent Problems	These tables are nearing their maximum extent limit and have exceeded the allowable maximum extent percentage set in the customizable thresholds of Space Analyst (to see all the thresholds, click the toolbar option labeled Reorganization Threshold Options). Tables suffering from this problem may cause space-related errors to result if more data is added to them.

The Table Reorganization Candidate Details section provides filters and a grid that you can use to view detailed space diagnostics for all or selected tables. The table below describes the grid data:

Data	Description
Table Owner	The owner of the table.
Table Name	The name of the table.
Segment Type	Indicates if the object is a table, table partition, etc.
Table (KB)	The total allocated space of the table in KB.
Number of Rows	The number of rows that the table contains. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Number of Blocks	The total number of blocks allocated to the table. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Empty Blocks	The number of empty blocks that the table contains. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.

Data	Description
Highwater Mark	The last block that Oracle will scan up to when searching the table for data. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Average Used Blocks	The average number of blocks that the table is using. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Block Inefficiency	A percentage that indicates how much wasted space (space that the table is not using) that the table contains. High percentages indicate tables that may be candidates for reorganization. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Chained Row Percent	A percentage that indicates how many rows in the table suffer from chaining or migration. High percentages indicate tables that may be candidates for reorganization. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Max Extent Percent	A percentage indicating how close an object is to reaching its maximum extent limit. High percentages indicate tables that may be candidates for reorganization.
Number of Extents	The number of extents that the table consumes.
Max Extents	The maximum extent limit for the table.
Extent Deficit	Indicates if the object can allocate its next extent. If 'N', this is because the tablespace does not currently have enough contiguous free space. This problem will be highlighted even if the tablespace has its autoextend property for one or more datafiles set to true.
Next Extent	The next extent (in bytes) that the table will attempt to allocate when more space for the table is required.
Max Free Space	The largest contiguous block of free space that exists in the tablespace where the table resides.
Tablespace Name	The tablespace that contains the table.

IBM DB2 for Linux, Unix, and Windows

The table below describes the four distinct types of Reorganization Candidate Summary that Space Analyst computes for IBM DB2 for Linux, Unix, and Windows:

Problem	Description
Overflow Problems	Displays the number of tables that fall outside the preset limits for the thresholds of this area. Can be addressed with reorganization of the affected table(s).
Space Utilization Problem	Displays the number of tables that fall outside the preset limits for the thresholds of this area. Can be addressed with reorganization of the affected table(s).
Page Utilization Problem	Displays the number of tables that fall outside the preset limits for the thresholds of this area. Can be addressed with reorganization of the affected table(s).
Index Contention	Displays the number of tables where the index space is shared by the Tablespace. To resolve, use the Space Management Wizard to change the Index Tablespace in table definition.

The table below describes the options and functionality on the Table Reorganization Candidate Details section:

Option	Description
Tablespace	Displays all Tablespaces (regular) in the database and the value [All].
Table Schema	Displays all schema in the database and the value [All].

Option	Description
Display	All Tables Tables with Overflow Problems - Displays a list of all tables where overflow threshold is exceeded. Tables with Space Utilization problems - Displays a list of all tables where space utilization threshold is exceeded. Tables with Page Utilization problems - Displays a list of all tables where the page utilization threshold is exceeded. Tables with Index Contention - Displays a list of all tables where the table and indexes on the table share the same Tablespace.
Refresh	Refreshes the table grid with criteria selected in the above lists.

The Table Reorganization Candidate Details section provides filters and a grid that you can use to view detailed space diagnostics for all or selected tables. The table below describes the grid data:

Data	Description
Table Schema	The schema name for the table.
Table Name	The name of the table.
Cardinality	The number of rows in the table.
Size	The size of Table in bytes.
Overflow	The number of overflow rows. Overflow rows can be created when rows are updated and the new rows contain more bytes than the old ones (VARCHAR fields), or when columns are added to existing tables.
Pages with Data	The number of allocated pages that contain data.
Total Pages	The total number of allocated pages.
Active Blocks	The number of Active Blocks in Tables with Multidimensional clustering).
Overflow %	Percentage of overflow rows. This should typical be a small value (5%).
Space Utilization%	A low space utilization is an indicator of excessive fragmentation (internal and external). This can occur when there are many pages with free space and/or many free pages. This should typically be > 68%.
Page Utilization %	A low page utilization is an indicator of excessive external fragmentation. This value should typically be more than 80%.
Reorg	This is an indicator of how severe the fragmentation problem for the table is. --- means Overflow%, Space Utilization% and Page Utilization % are all with in the thresholds. *-* means that Overflow%, and Page Utilization % are outside the threshold but Space Utilization% is with in the threshold.
Tablespace	Tablespace where the table data resides.
Indexspace	Tablespace where the index data resides.

Available Functionality

NOTE: Functionality differ by platform.

The following functionality is available in the Command menu, shortcut menu, and toolbar:

Analysis Report	Analyze	DBCC
Estimate Size	Extract	Map
Table Wizard for IBM DB2 (IBM DB2 LUW)	Open Editor	Open Wizard

Place	Quick Reorganize	Reorganize
Threshold Options	Update Statistics	

For more information, see Space Analyst Tabs.

Space Analyst for Sybase and IBM DB2 for Linux, Unix, and Windows - Indexes Tab

The Indexes Tab displays diagnostic information regarding index space inefficiencies that may exist in your database.

The Index Reorganization Candidate Summary section uses a bar chart to count the various types of index space inefficiencies in your database. Note that the reorganization threshold options are used to compute the counts for each type of space inefficiency.

The table below describes the three distinct types of table space inefficiencies that Space Analyst computes:

Problem	Description
BLevel	These indexes have extended beyond the allowable B-level size set in the customizable thresholds of Space Analyst (to see all the thresholds, click the toolbar option labeled Reorganization Threshold Options). Indexes with extended levels can take longer to scan and can therefore contribute to slower response times.
Next Extent Problems	These indexes cannot obtain their next extent of free space due to a lack of needed contiguous free space in the tablespace. Indexes suffering from this problem may cause space-related errors if more data is added to their parent table. This problem will be highlighted even if the tablespace that contains the table has its autoextend property for one or more datafiles set to true.
Max Extent Problems	These indexes are nearing their maximum extent limit and have exceeded the allowable maximum extent percentage set in the customizable thresholds of Space Analyst (to see all the thresholds, click the toolbar option labeled Reorganization Threshold Options). Indexes suffering from this problem may cause space-related errors to result if more data is added to their parent table.

The Index Reorganization Candidate Details section provides filters and a grid that you can use to view detailed space diagnostics for all or selected indexes. The grid contains the following data:

Data	Description
Index Owner	The owner of the index.
Index Name	The name of the index.
Index Type	Indicates if the type of index (index, index partition, etc.)
Index (KB)	The total allocated space of the index in KB.
Number of Rows	The number of rows that the parent table contains. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Clustering Factor	Indicates the amount of order of the rows in the table based on the values of the index. If the value is near the number of blocks, then the table is very well ordered. If the value is near the number of rows, then the table is very randomly ordered. In this case, it is unlikely that index entries in the same leaf block point to rows in the same data blocks.
Index Blevel	Indicates the level of the index. Numbers beyond 3-4 can indicate a fragmented index.

Data	Description
Number of Blocks	The total number of blocks allocated to the index. Note: this column may report NULL if the index has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Max Extent Percent	A percentage indicating how close an object is to reaching its maximum extent limit. High percentages indicate indexes that may be candidates for reorganization.
Number of Extents	The number of extents that the index consumes.
Max Extents	The maximum extent limit for the index.
Can Extent Space	Indicates if the object can allocate its next extent. If 'N', this is because the tablespace does not currently have enough contiguous free space. This problem will be highlighted even if the tablespace has its autoextend property for one or more datafiles set to true.
Next Extent	The next extent (in bytes) that the index will attempt to allocate when more space for the index is required.
Max Free Space	The largest contiguous block of free space that exists in the tablespace where the index resides.
Tablespace Name	The tablespace that contains the index.

NOTE: The grid can also contain a column for deleted leaf rows if the Get Deleted Leaf Row Statistics option has been selected. Be aware that the calculation of this statistic can take a lot of time for databases with large numbers of indexes. Large numbers of deleted leaf rows can indicate an index in need of reorganization.

IBM DB2 for Linux, Unix, and Windows

The Index Reorganization Candidate Summary section uses a bar chart to represent the various types of index space inefficiencies in your database:

- Indexes with Cluster Ratio Problems
- Indexes with Used Space% problems
- Indexes with Balance Problems
- Indexes with Deleted RID problems
- Indexes with Empty Leaf Problems
- Indexes with Low Uniqueness

NOTE: Space Analyst uses the options you select in the Reorganization Wizard to compute each type of space inefficiency.

TIP: All of the problems in the Index Reorganization Candidate Summary section except low cardinality can be corrected with index reorganization. The Low cardinality is an indicator of an index where the columns on which it is defined have low level of uniqueness. Such indexes should either be dropped or changed to include columns that will make the index more unique.

The table below describes the options and functionality on the Index Reorganization Candidate Details section:

Option	Description
Tablespace	Displays all Tablespaces (regular/large) in the database and the value [All].
Index Schema	Displays all schema in the database and the value [All].

Option	Description
Display	All Indexes Indexes with Cluster Ratio Problems Indexes with Used Space% problems Indexes with Balance Problems Indexes with Deleted RID problems Indexes with Empty Leaf Problems Indexes with Low Uniqueness
Refresh	Refreshes the index grid with criteria selected in the above lists.

The Index Reorganization Candidate Details section provides filters and a grid that you can use to view detailed space diagnostics for all or selected indexes. The grid contains the following data:

Data	Description
Index Schema	The schema name for the index.
Index Name	The name of the index.
Table Schema	The schema name for the base table
Table Name	The base table name.
Table Cardinality	The number of rows in the base table.
Avg Key Size	The index key size, calculated from the average column length of all columns participating in the index.
Leafs	The total number of index leaf pages.
Pseudo Empty Leafs	The number of pseudo empty index leaf pages. A pseudo empty index leaf page is a page on which all the RIDs are marked as deleted, but have not been physically removed.
Pseudo Deleted RIDs	A RID that is marked deleted. This statistic reports pseudo deletes RIDs on leaf pages that are not pseudo empty. It does not include RIDs marked as deleted on leaf pages where all the RIDs are marked deleted.
Levels	The number of index levels.
Unique Entries	The number of unique index entries that are not marked deleted.
Page Size	The page size of the table space in which the table indexes reside, specified at the time of table creation.
% Free	Specifies the percentage of each index page to leave as free space, a value that is assigned when defining the index.
Cluster Ratio	The clustering ratio of an index should be greater than 80 percent. When multiple indexes are defined on one table, some of these indexes have a low cluster ratio. (The index sequence is not the same as the table sequence.) This cannot be avoided. Be sure to specify the most important index when reorganizing the table. The cluster ratio is usually not optimal for indexes that contain many duplicate keys and many entries.
Used Space%	Less than 50 percent of the space reserved for index entries should be empty.
Balance	To determine if recreating the index would result in a tree having fewer levels. This formula checks the ratio between the amount of space in an index tree that has one less level than the current tree, and the amount of space needed. If a tree with one less level could be created and still leave PCTFREE available, then a reorganization is recommended. The actual number of index entries should be more than 90% (or 100-PCTFREE) of the number of entries an NLEVELS-1 index tree can handle (only checked if NLEVELS>1).
Deleted RIDs	The number of pseudo-deleted RIDs on non-pseudo-empty pages should be less than 20 percent.

>

Data	Description
Empty Leaf	The number of pseudo-empty leaf pages should be less than 20 percent of the total number of leaf pages.
Uniqueness	100 * Index Cardinality / Table Cardinality - A low index uniqueness may indicate the need to Alter or Drop the index.
Reorg	An indicator of how severe the fragmentation problem for the index is. Each hyphen (-) displayed in this column indicates that the calculated results were within the set bounds of the corresponding formula, and each asterisk (*) indicates that the calculated result exceeded the set bounds of its corresponding formula: - or * on the left column corresponds to Cluster Ratio - or * in the second from left column corresponds to Used Space% - or * in the middle column corresponds to Balance - or * in the second column from the right corresponds to Deleted RIDs - or * on the right column corresponds to Empty Leaf
Tablespace	The tablespace where the table data resides.
Indexspace	The tablespace where the index data resides.

Available Functionality

The following functionality is available in the Command menu, shortcut menu, and toolbar:

NOTE: Functionality differ by platform.

Analysis Report	Analyze	DBCC
Drop	Extract	Estimate Size
Map	Index Wizard for IBM DB2 (IBM DB2 LUW)	Open Editor
Open Wizard	Place	Quick Reorganize
Reorganize	Threshold Options	Update Statistics

For more information, see:

Finding Space-Related Table Problems

Space Analyst Tabs

Space Analyst for Oracle - Indexes Tab

The Indexes Tab displays diagnostic information regarding index space inefficiencies that may exist in your database.

The Index Reorganization Candidate Summary section uses a bar chart to count the various types of index space inefficiencies in your database. Note that the reorganization threshold options are used to compute the counts for each type of space inefficiency.

The table below describes the three distinct types of table space inefficiencies that Space Analyst computes:

Problem	Description
BLevel	These indexes have extended beyond the allowable B-level size set in the customizable thresholds of Space Analyst (to see all the thresholds, click the toolbar option labeled Reorganization Threshold Options). Indexes with extended levels can take longer to scan and can therefore contribute to slower response times.
Next Extent Problems	These indexes cannot obtain their next extent of free space due to a lack of needed contiguous free space in the tablespace. Indexes suffering from this problem may cause space-related errors if more data is added to their parent table. This problem will be highlighted even if the tablespace that contains the table has its autoextend property for one or more datafiles set to true.
Max Extent Problems	These indexes are nearing their maximum extent limit and have exceeded the allowable maximum extent percentage set in the customizable thresholds of Space Analyst (to see all the thresholds, click the toolbar option labeled Reorganization Threshold Options). Indexes suffering from this problem may cause space-related errors to result if more data is added to their parent table.

The Index Reorganization Candidate Details section provides filters and a grid that you can use to view detailed space diagnostics for all or selected indexes. The grid contains the following data:

Data	Description
Index Owner	The owner of the index.
Index Name	The name of the index.
Index Type	Indicates if the type of index (index, index partition, etc.)
Index (KB)	The total allocated space of the index in KB.
Number of Rows	The number of rows that the parent table contains. **NOTE:** This column may report NULL if the table has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Clustering Factor	Indicates the amount of order of the rows in the table based on the values of the index. If the value is near the number of blocks, then the table is very well ordered. If the value is near the number of rows, then the table is very randomly ordered. In this case, it is unlikely that index entries in the same leaf block point to rows in the same data blocks.
Index Blevel	Indicates the level of the index. Numbers beyond 3-4 can indicate a fragmented index.
Number of Blocks	The total number of blocks allocated to the index. Note: this column may report NULL if the index has not been analyzed via the ANALYZE command or through the DBMS_STATS package.
Max Extent Percent	A percentage indicating how close an object is to reaching its maximum extent limit. High percentages indicate indexes that may be candidates for reorganization.
Number of Extents	The number of extents that the index consumes.
Max Extents	The maximum extent limit for the index.
Can Extent Space	Indicates if the object can allocate its next extent. If 'N', this is because the tablespace does not currently have enough contiguous free space. This problem will be highlighted even if the tablespace has its autoextend property for one or more datafiles set to true.
Next Extent	The next extent (in bytes) that the index will attempt to allocate when more space for the index is required.
Max Free Space	The largest contiguous block of free space that exists in the tablespace where the index resides.
Tablespace Name	The tablespace that contains the index.

NOTE: The grid can also contain a column for deleted leaf rows if the Get Deleted Leaf Row Statistics option has been selected. Be aware that the calculation of this statistic can take a lot of time for databases with large numbers of indexes. Large numbers of deleted leaf rows can indicate an index in need of reorganization.

Available Functionality

The following functionality is available in the Command menu, shortcut menu, and toolbar:

NOTE: Functionality differ by platform.

Analysis Report	Analyze	DBCC
Drop	Extract	Estimate Size
Map	Index Wizard for IBM DB2 (IBM DB2 LUW)	Open Editor
Open Wizard	Place	Quick Reorganize
Reorganize	Threshold Options	Update Statistics

For more information, see Space Analyst Tabs.

Space Analyst - Storage Overview Tab

The Storage Overview Tab displays an overview of the storage/space situation on the target server:

The table below describes the sections on the Storage Overview Tab of Space Analyst for Microsoft:

Section	Description
Server Overview	Overview of total database and log space allocated on the SQL Server, as well as a count of databases, files, and file groups.
Log Overview	Contains a graphical and grid control that works in the following way: - When the graph is displayed, a bar chart is shown that lists the total percent used of all logs on the server. - When the grid is displayed, detailed storage information for each database log is presented.
Database Overview	Contains a graphical and grid control that works in the following way: - When the graph is displayed, a bar chart is shown that lists the total space usage of all databases on the server. - When the grid is displayed, detailed storage information for each database is presented.
Server Disk Overview	Contains an overview of total free space, by physical server drive, of the server that houses SQL Server.

The table below describes the sections on the Storage Overview Tab of Space Analyst for Sybase:

Section	Description
Server Overview	Overview of total database and log space allocated on the Sybase Server, as well as a count of databases and devices.

Section	Description
Log Overview	Contains a graphical and grid control that works in the following way: - When the graph is displayed, a bar chart is shown that lists the total percent used of all logs on the server. - When the grid is displayed, detailed storage information for each database log is presented. Database name Total log space Free log space Used log space Log percent free
Database Overview	Contains a graphical and grid control that works in the following way: - When the graph is displayed, a bar chart is shown that lists the total space usage of all databases on the server. - When the grid is displayed, detailed storage information for each database is presented: Database name Total database space Database free space Database used space Database percent free
Device Overview	Contains a graphical and grid control that works in the following way: - When the graph is displayed, a bar chart is shown that lists the total space used (in MB) of all devices defined to the server. - When the grid is displayed, detailed storage information for each database device is presented: Device name Physical name Total space Used space Free space Percent free Percent used

Available Functionality

The following functionality is available in the Command menu, shortcut menu, and toolbar for Microsoft:

Analysis Report	Extract	Map
Database Wizard (SQL Server)	Object Details	Open Editor
Open Wizard	DBCC	Shrink
Threshold Options		

Available Functionality

The following functionality is available in the Command menu, shortcut menu, and toolbar for Sybase:

Analysis Report	Extract	Index Details
Open Editor	Open Wizard	Database Wizard (Sybase)

>

Move Log	DBCC	Table Details
Threshold Options	For more information, see:	Pinpointing 'Out of Space' Problems
Space Analyst Tabs		

Space Analyst for Microsoft - File Groups/Files Tab

The Files/File Groups Tab displays a graph and grid for filegroups on the server and a grid for files.

> **TIP:** You can toggle between the graphs and grid views.

The File Groups/Files Tab reports space usage on SQL Server by file group and by file.

When the File Groups radio button is pressed, Space Analyst presents a graphical layout of file group space usage across all databases. For large servers, not all file groups will initially be shown, but the graph is scrollable so you can click the graph and scroll to see all file groups on the chart. When in grid mode, more detailed information is shown for each file group including its auto-growth properties, how many files make up the group, the amounts (in MB) taken up by tables and indexes, and more.

When the Files radio button is pressed, detailed information is displayed for each individual file for every database on the server.

Available Functionality

The following functionality is available in the Command menu, shortcut menu, and toolbar:

Analysis Report	Extract	Map
Database Wizard (SQL Server)	Object Details	Open Editor
Open Wizard	DBCC	Shrink
Open Wizard	Threshold Options	

For more information, see Space Analyst Tabs.

Space Analyst for Microsoft - Map Tab

The Map Tab displays a graphical file map that displays extent placement for objects within database files.

The Map Tab lets you view a graphical layout regarding the physical placement of every object inside a selected database and file. The map helps identify allocated and unused extents throughout a database file. An object grid is also present that lets you view detailed information about every object's storage usage in that file. Keep in mind that an object can span more than one file.

The Space Info section at the top contains information regarding the total number of extents allocated to a database file, a count of extents that contain table data, a count of extents that contain index data, and a count of extents that contain text data.

Databases are comprised of logical pages, each of which is a fixed 8KB in size. Note that the 8KB page is also a unit of I/O, which affects performance and locking. Extents are collections of database pages. An extent is made up of eight, 8K pages and is therefore 64KB in size. Extents are the graphical "blocks" displayed in the map.

Extents come in two flavors - uniform and mixed. Uniform extents ascribe all eight database pages to a single object, whereas mixed extents have the potential to hold multiple objects. With mixed extents, SQL Server tries to preserve space by not letting very small objects take up an entire extent and therefore waste a lot of space. By placing your mouse pointer over an extent "block" in the map, Space Analyst will present a pop-up dialog displaying what object(s) is/are contained in that extent.

>

The table below describes the options and functionality on the Map Tab of Space Analyst:

Option	Description
Space Info	Displays total extent count, extents with tables count, extents with indexes count, and extents with text count. Note that an extent can contain more than one object and more than one object type.
Map	Displays object extent data from beginning to end of the database or file. All blocks in the map are the same size. A Microsoft SQL Server extent is 64KB in size or 8, 8K pages. Purple - Indicate extents that have been allocated to objects (used), and extents that are free. Green - Free extents.
Owner	Displays table owner.
Table name	Displays table name.
Object type	Displays object type.
Index name	Displays index name.
Pages used	Displays pages used.
Pages allocated	Displays pages allocated.
In extents	How many extents each object is in.
Table id	Displays table id.
Index id	Displays index id.

The Tablespace Map Tab includes the following buttons:

Button	Description
Legend	Opens the Tablespace Map Legend.
Display	Opens the Display window. The Display window lets you specifies the magnification level for the Tablespace Map.
Fit to Window	Fits the Tablespace Map to the window.
Full Screen	Fits the Tablespace Map to the full screen.
Print Map	Opens the Print dialog box to print the Tablespace Map.

Available Functionality

The following functionality is available in the Command menu, shortcut menu, and toolbar:

Analysis Report	Extract	Database Wizard (SQL Server)
Object Details	Open Editor	Open Wizard
DBCC	Shrink	Threshold Options

For more information, see Space Analyst Tabs.

Space Analyst for Microsoft - Databases Tab

The Databases Tab lets you analyze databases.

TIP: You can toggle between the graphs and grid views.

The Databases Tab displays information regarding space usage across all databases, potential space-related bottlenecks that are threatening databases, and recent database backup information.

When in graph form, the Database Summary section displays a graphical view of table, index, and free space usage (in MB) across all databases. For large servers, not all databases will initially be shown, but the graph is scrollable so you can click the graph and scroll to see all databases on the chart. When in grid form, the Database Summary section displays much more detail regarding database and log usage, object counts, and more.

The Database Bottlenecks section provides a graphical count of various space-related bottlenecks that are threatening the availability of your databases (some of which are controlled by Space Analyst's customizable thresholds). When in grid form, each database is listed along with information that relays whether or not a potential bottleneck exists for that database.

The Database Backups section details the most recent backup for each database and the size (in MB) of the backup.

Available Functionality

The following functionality is available in the Command menu, shortcut menu, and toolbar:

Analysis Report	Extract	Map
Database Wizard (SQL Server)	Object Details	Open Editor
Open Wizard	DBCC	Shrink
Threshold Options		

For more information, see Space Analyst Tabs.

Space Analyst for Microsoft - Objects Tab

The Objects Tab displays database objects, including tables and indexes.

You can choose whether to view table or index information by clicking either the Tables or Indexes ratio buttons. You can also choose whether to view general space information for tables and indexes or reorganization metrics by selecting the General Space or Space Diagnostics radio buttons.

> **NOTE:** To obtain space diagnostic information for each object, Space Analyst uses the SQL Server DBCC SHOWCONTIG command. This command may take a while to run on large databases. Also, the command is not as efficient in terms of performance on version 7.0 of SQL Server as it is on SQL Server 2000, so understand that running space diagnostics on most versions of SQL Server will take some time.

The **Object summary** presents one of two views:

- When the **General Space** radio button is pressed, a graph is presented that visually displays that amount of allocated space (in KB) for tables, indexes, and free space. For large servers, not all objects for a database will initially be visible on the chart, but the graph is scrollable so you can click the graph and scroll to see all database objects on the chart.
- When the **Space Diagnostics** radio button is pressed, Space Analyst presents a count of all objects (tables or indexes) that exceed your customizable thresholds for various forms of fragmentation of other space-related inefficiency. See below for a thorough discussion of these types of space problems.

The **Object Details** grid presents one of two views:

- When the **General Space** radio button is pressed, general space usage for tables or indexes is shown. Base object metadata (object name, type, partitions, etc.) is shown along with the amounts of total, used, and free space (in KB) for the object. The percent of space used for the object is also shown along with the percent of space in the database that the object consumes.

>

- When the **Space Diagnostics** radio button is pressed, Space Analyst present fragmentation and other space-related inefficiency information for all objects in the selected database. Note that definitions for all the various forms of object fragmentation and space inefficiencies follow later in the help.

The table below describes the options and functionality on the Objects Tab of Space Analyst:

View	Description
Database	Lets you switch between databases.
Tables	Lets you switch between viewing space for tables and indexes. Displays object details when you select.
Indexes	Lets you switch between viewing space for tables and indexes. Displays object details when you select.
General Space View	Displays general space information for tables and indexes in the selected database. Displays object details when you select.
Space Diagnostic View	Displays fragmentation and space information to help the DBA decide if a reorg is necessary. Displays object details when you select.
File Group	Lets you view object data for all filegroups for the database or only a selected filegroup.
Owner	Lets you view object data for all table owners or only a selected owner.
Display	Lets you view space diagnostic information for tables or indexes.

Available Functionality

The following functionality is available in the Command menu, shortcut menu, and toolbar:

Analysis Report	Extract	Database Wizard (SQL Server)
Open Editor	Open Wizard	DBCC
Threshold Options	Update Statistics	

For more information, see Space Analyst Tabs.

Space Analyst for Sybase ASE - Database Devices Tab

The Database Devices Tab displays database device information.

TIP: You can toggle between graph and grid views.

The table below describes the sections on the Database Devices Tab of Space Analyst:

Section	Description
Device Summary	Contains summary information for all Sybase database devices. Displays the amount of data, log, data and log, and free space amounts (in MB).
Device Details	Displays detailed information for selected devices. Lets you click in a bar chart and populate the lower section of the tab with that device's information.
Database Fragments	Displays database fragments.

>

Available Functionality

The following functionality is available in the Command menu, shortcut menu, and toolbar:

Analysis Report	Disk Resize	Extract
Index Details	Move Log	Database Wizard (Sybase)
Database Device Wizard (Sybase)	Segment Wizard (Sybase)	Open Editor
Open Wizard	DBCC	Reinitialize
Refit	Table Details	Threshold Options

For more information, see Space Analyst Tabs.

Space Analyst for Sybase ASE - Databases Tab

The Databases Tab lets you analyze databases.

> **TIP:** You can toggle between graph and grid views.

The table below describes the sections on the Databases Tab of Space Analyst:

Section	Description
Database Summary	Contains chart that displays database used and free space, along with log used and free space information for databases. The summary section totals up the total database, log, table, and index space. Also displays a summary count of all user tables and indexes.
Database Details	Displays detailed information for selected databases. Lets you select whether the graph shows space utilization or object fragmentation scores.

Available Functionality

The following functionality is available in the Command menu, shortcut menu, and toolbar:

Analysis Report	Disk Resize	Extract
Index Details	Move Log	Database Wizard (Sybase)
Database Device Wizard (Sybase)	Segment Wizard (Sybase)	Open Editor
Open Wizard	DBCC	Reinitialize
Refit	Table Details	Threshold Options

For more information, see Space Analyst Tabs.

Space Analyst for IBM DB2 for Linux, Unix, and Windows - Tablespaces Tab

The Tablespace Tab lets displays tablespace space allocation and tablespace.

>

Available Functionality

The following functionality is available in the Command menu, shortcut menu, and toolbar:

NOTE: Functionality differ by platform.

Analysis Report	Extract	**Open Editor** - see Modifying objects using editors
Open Wizard	Reorganize	Update Statistics
Tablespace Wizard for IBM DB2 (IBM DB2 LUW)	Threshold Options	

For more information, see Space Analyst Tabs.

Space Analyst for Functionality

Space Analyst offers functionality that help you analyze space. The Space Analyst support the following functionality:

NOTE: Functionality differ by platform.

Analysis Report	Convert to Dictionary Managed	Convert to Locally Managed
Disk Resize	DBCC	Drop
Estimate Size	Extract	Index Details
Map	Move Log	Database Wizard (Sybase)
Database Device Wizard (Sybase)	Segment Wizard (Sybase)	Table Wizard for IBM DB2 (IBM DB2 LUW)
Tablespace Wizard for IBM DB2 (IBM DB2 LUW)	Open Editor	Open Wizard
Options	Place	Quick Reorganize
DBCC	Refit	Reinitialize
Reorganize	Shrink	Table Details
Update Statistics		

Space Analyst - Analysis Report

The Analysis Report dialog box lets you build an HTML analysis report for the selected tablespace. You can run reports on a recurring basis and dynamically alter any saved thresholds that flag space problems in the tablespace being analyzed.

For Oracle, the report includes sections for particular reorganization needs, including:

- Fragmentation index for the tablespace
- Tables with high block inefficiencies
- Tables or indexes with next extent space deficit problems
- Tables or indexes with maximum extent limit problems
- Tables or indexes with percent increase greater than zero

- Tables with chained/migrated row problems
- Indexes with B-tree depth problems

For Microsoft, the report contains different sections that highlight particular reorganization needs, including:

- Free space levels for the database and log
- Tables with forwarded record problems
- Indexes with average fragmentation problems
- Tables or indexes with extent fragmentation problems
- Tables or indexes with logical fragmentation problems
- Tables or indexes with scan density problems
- Tables or indexes with average page density problems

For Sybase, the report contains different sections that highlight particular reorganization needs, including:

- Free space levels for the database and log
- Tables with forwarded row problems
- Tables with empty page problems
- Tables with deleted row problems
- Tables with low fragmentation scores
- Indexes with unused space problems
- Indexes with empty leaf problems
- Indexes with height/level problems
- Indexes with low fragmentation scores

For IBM DB2 for Linux, Unix, and Windows, the report contains different sections that highlight particular reorganization needs, including:

- Tables with overflow problems
- Tables with space utilization problems
- Tables with page utilization problems
- Tables with index space contention
- Indexes with cluster ratio problems
- Indexes with used space% problems
- Indexes with balance problems
- Indexes with deleted RID problems
- Indexes with empty leaf problems
- Indexes with low uniqueness
- Free levels of tablespace and log space

To open the Analysis Report dialog box:

1 Open the **Space Analyst.**

2 On the **Command** menu, click the **Analysis Report**.

For all platforms, browse to or type a report title and make your desired selections from the drop-down lists. After the report executes, you can print it or save it.

To print the report:

1 Click **Print** from the Analysis Report toolbar.

To save the report:

1 Select **Save** from the **File** menu on the main toolbar.

Space Analyst - Convert to Locally-Managed

NOTE: The Convert to Locally-Managed dialog box is for Oracle 8.1.6 or later.

The Convert to Locally-Managed dialog box lets you convert a dictionary-managed tablespace to locally-managed. Objects in locally-managed tablespaces have unlimited extents; they will never fail to extend due to an imposed maximum extent limit.

Important Notes

There are a few restrictions on these conversion procedures (for example, Oracle9i UNDO tablespaces currently can't be converted, etc.), so you should check the Oracle documentation for the specifics on performing a conversion.

The Convert to Locally-Managed option is only available if the selected tablespace(s) are:

- ONLINE
- In Read/Write mode
- Dictionary managed
- Not UNDO tablespaces
- Not Temporary tablespaces
- Not the SYSTEM tablespace

The table below describes the options and functionality on the Convert to Locally-Managed dialog box:

Option	Description
Tablespaces to be Converted	Displays the tablespaces to be converted to locally-managed.
Conversion Options	OPTIONAL: Lets you type allocation unit and relative file number information.

Completing the Convert to Locally-Managed Dialog Box for Space Analyst

To complete this dialog box, do the following:

1 Select **Convert to Locally-Managed**.

OR

Right-click the target object, and then select **Convert to Locally-Managed**.

The Space Analyst opens the dialog box.

2 Do one of the following:

Button	Description
	Opens the Preview dialog box.
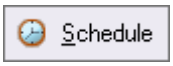	Opens a scheduling application.
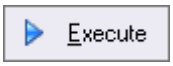	Executes the task.

For more information, see Convert to Locally-Managed.

Space Analyst - Convert to Dictionary-Managed

NOTE: The Convert to Dictionary-Managed functionality is for Oracle 9i only.

The Convert to Dictionary-Managed dialog box lets you convert a locally-managed tablespace to dictionary.

Important Notes

The Convert to Dictionary-Managed dialog box is only available if the selected tablespace(s) are:

- ONLINE
- In Read/Write mode
- Locally managed
- Not UNDO tablespaces
- Not Temporary tablespaces
- Not the SYSTEM tablespace

The table below describes the options and functionality on the Convert to Dictionary-Managed dialog box:

Option	Description
Tablespaces to be Converted	Displays the tablespaces to be converted to dictionary-managed.

Completing the Dictionary-Managed Dialog Box for Space Analyst

To complete this dialog box, do the following:

1 Select **Convert to Dictionary-Managed**.

OR

Right-click the target object, and then select **Convert to Dictionary-Managed**.

The Space Analyst opens the dialog box.

2 Do one of the following:

Button	Description
	Opens the Preview dialog box.
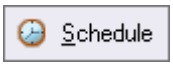	Opens a scheduling application.
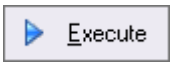	Executes the task.

For more information, see Convert to Dictionary-Managed.

Space Analyst - Extract

The Extract functionality lets you extract the database DDL in an ISQL editor.

To extract the database DDL in an ISQL editor, do the following:

1 Click the **Extract DDL** button.

OR

Select **Extract**.

OR

Right-click the target object, and then select **Extract**.

Space Analyst extract the database DDL in an ISQL editor.

Space Analyst - Map

The Map functionality opens the Tablespace Map Tab or Map Tab of the Space Analyst for the selected index.

Important Notes

None

For more information, see Opening the Tablespace Map Tab.

Opening the Tablespace Map Tab for Space Analyst

To open the Tablespace Map or Map Tab, do the following:

1 Select **Map**.

OR

Right-click the target object, and then select **Map**.

Space Analyst open the Tablespace Map Tab or Map Tab.

For more information, see Map.

Space Analyst - Open Editor

The Open functionality lets you modify objects in the object editors.

To open the object editor, do the following:

1 In any grid, select an object.

2 Click the **Open Editor** button.

OR

Select **Open**.

OR

Right-click the target object, and then select **Open**.

Space Analyst opens the object editor.

Space Management Wizard

The Space Management Wizard can create two specialized space management jobs that allow you to proactively eliminate space threats before they affect your database:

You can create a "reorganization analyst job" that allows you to specifically select various databases and/or objects and then set reorganization thresholds that determine when a reorganization will occur. Once constructed, you can schedule these jobs to run as often as you'd like. During execution, these jobs will interrogate your chosen database/objects and determine if a reorganization needs to occur by applying your customized reorganization thresholds. If any objects are found to need reorganization, they are dynamically reorganized using any customized options you specify.

You can build space maintenance jobs that will perform functions such as check the validity of your databases and filegroups, repair any found minor corruption, and eliminate unused space in your databases. You can also set the job up to notify you if any problems are found during job execution.

Creating or Editing a Space Management Job

The Space Management Wizard lets you create space reorganization or maintenance jobs without knowing the code. After creation, jobs can be reused as is, or edited.

>

To open the Space Management Wizard

1 Start the Space Analyst and click the **Open Wizard** button on the Space Analyst toolbar.

OR

Select **Open Wizard** from the Space Analyst menu.

The sections below display the fields you may find creating or editing a space management job. Be aware that the sections you see, and the order in which they appear in the wizard vary depending on the selections you make and the database platform.

New or Used Task	Type of Space Management Task	Database/Objects to Be Part of the Reorganization or Maintenance Task
Custom Reorganization	Save the OperationExecution Preferences for a Reorganization Operation	Global Customizations
Specific Object Customization	Threshold Criteria for Object Reorganization	Integrity Check Options
Integrity Check for Specific Database Objects	Unused Space Options	Problem Notification

New or Used Task

Here you choose whether to create a new space management task or whether to use a previously saved task. If you want to reuse a previously created job, you can type or browse to find the previously saved task definition file (*.sptsk).

Type of Space Management Task

Required Information	Description
Create standard reorganization job	Select to build a reorganization job for a database or the objects you specify.
Create reorganization analyst job	Select to create and deploy a reorganization job that runs periodically to detect and automatically correct space inefficiencies.
Create space maintenance job	Select to build a space maintenance job that validates database structures and performs storage maintenance tasks.

Database/Objects to Be Part of the Reorganization or Maintenance Task

Choose the databases and/or objects you want to include in the job you are creating. Simply move items from the tree in the right pane to the left pane by clicking the right arrow. You can amend your selections using the left arrow. The selection is limited to entire databases for a space maintenance job.

Custom Reorganization

Required Information	Description
Perform no customizations	When you choose this option, all default object and database attributes are used.
Perform global customizations	This option gives you the chance to define options that will be globally applied to all objects being reorganized.
Perform specific object customizations	Here you individualize options for specific objects being reorganized.

Save the Operation

If you want to save the operation, you can use the default name supplied or type another in the upper panel. If you don't want to save the operation, remove the default name. Browse to the location where you want to save the file. The job as it was created appears in the read-only preview panel. To make changes, you must go back and change your selections or open the job in an ISQL window prior to execution.

Execution Preferences for a Reorganization Operation

Option	Description
Open the reorganization script in an ISQL editor.	The reorganization script will be generated and then opened in a new ISQL editing session.
Perform reorganization now	The reorganization is run immediately with results presented afterward in an ISQL window.
Schedule the reorganization script for later execution.	The reorganization script will be generated and then scheduled to run at a specified date/time.

>

Global Customizations

Option	Description
Relocate to file group	Select to perform an extended alter and move the table(s) to the requested file group using the ON keyword in the table creation DDL.
Perform sorted reorg using table's clustered index	Select to rebuild of any table's clustered index.
Perform online reorganization where possible	**MICROSOFT SQL SERVER 2000 OR LATER ONLY:** Select to use the DBCC INDEXDEFRAG command to perform an online rebuild of every table's index(es).
View reorganization output messages	Select for the WITH NO_INFOMSGS clause of any DBCC command to not be used.
Defragment tables in DROP/CREATE fashion	Select to defragment tables in DROP/CREATE fashion.
Relocate to file group	Select for DROP_EXISTING clause instead of doing a drop and then create.
Fill Factor	Select to specify a new fill factor for the DBCC DBREINDEX command or a create/drop of an existing index.
Index Reorganization	Rebuild: This is the default (ALTER INDEX...REBUILD for MS SQL 2005; MS SQL Server 2000 and 2005 display DBCC DBREINDEX). Drop/Create **For MS SQL Server 2005 or later:** Reorganize
Perform online reorganization where possible	**MICROSOFT SQL SERVER 2000 OR LATER ONLY:** Select to use the DBCC INDEXDEFRAG command to perform an online rebuild of every index.
View reorganization output messages	Select for the WITH NO_INFOMSGS clause of any DBCC command to not be used.

Specific Object Customization

Option	Description
Select	Drop-down list lets you switch between indexes and tables and see the particular objects for each.
Relocate to file group	Select a clause instead of doing a drop and then create.
Fill Factor	For indexes, select to specify a new fill factor.
Perform sorted reorganization using table's clustered index	For tables, select to perform sorted reorganization using table's clustered index.
Index reorganization	Rebuild: This is the default. Drop/Create Reorganize
Perform online reorganization where possible	**MICROSOFT SQL SERVER 2000 OR LATER ONLY:** Select to perform an online rebuild of every index.
View reorganization output messages	Select for the WITH NO_INFOMSGS clause of any DBCC command to not be used.
Defragment tables in DROP/CREATE fashion	For tables, select to defragment tables in DROP/CREATE fashion.

Threshold Criteria for Object Reorganization

NOTE: The table combines criteria for MS SQL Server 2000 and 2005.

Option	Description
Table: Average Page Density	Default is less than or equal to 50.
Scan Density Percent	Default is less than or equal to 0.
Logical Fragmentation	Greater than or equal to 25 is the default value.
Extent Fragmentation	Default is greater than or equal to 50.
Forwarded Record Percent	Default is greater than or equal to 50.
Index: Average Fragmentation	Default is greater than or equal to 25.
Scan Density Percent	Default is less than or equal to 0.
Logical Fragmentation	Greater than or equal to 25 is the default value.
Extent Fragmentation	Default is greater than or equal to 50.
Index Tree Depth	Default is greater than or equal to 3.
Average Page Density	Default is less than or equal to 50.

NOTE: All defaults are editable. Also, all current metrics for these criteria can be found on the Objects tab >Space Diagnostics view. The default thresholds are visible when you choose Threshold Options from the Space Analyst drop-down menu.

Integrity Check Options

Option	Description
Do not perform any integrity checks	Self-explanatory.
Check databases	Checks the database integrity for database(s) you identified earlier.
Check file groups	Self-explanatory.
Check specific database objects	In a later step you are able to select the particular objects.
Include index checks for all tables	This is the default and isn't actually optional.
Attempt to repair minor problems	Enabled as a default for database or database objects integrity checks. You can unselect.
Include check for database catalogs	Enabled as a default for database integrity checks. You can unselect.
Correct data dictionary inaccuracies	Enabled as a default for database integrity checks. You can unselect.
Check data purity	Verifies that columns only have data in the correct ranges, and so on.

Integrity Check for Specific Database Objects

Here you select objects for the database(s) you have included in this maintenance job.

>

Unused Space Options

Option	Description
Do not eliminate unused space	Self-explanatory
Eliminate unused space for databases	Select to only truncate space for databases larger than a specific size (in MB).
Only eliminate space when database is larger than	Available if you selected Eliminate unused space for databases option. Lets you specify maximum database size in MB.
Amount of free space to leave for database	Lets you specify how much free space is to remain after a space truncation.
Eliminate unused space for specific database tables	Self-explanatory.

Problem Notification

Option	Description
No, do not notify me	Select to not have Space Analyst notify you of any errors.
Yes, notify me in the following ways.	Select to have Space Analyst notify you of any errors and then type an e-mail or net send address.

Space Analyst - Threshold Options

The Space Analyst Reorganization Thresholds Options dialog box lets you customize the various reorganization thresholds that can be used to flag storage and data objects that need to be reorganized. The thresholds you specify will apply to the Space Analysis Reports.

Important Notes

None

The table below describes the options and functionality on the Space Analyst Thresholds Options dialog box for IBM DB2 for Linux, Unix, and Windows:

Option	Description
Table Reorganization Thresholds	The table thresholds apply to the displays used in tab 4 of the Space Analyst and the queries that are run to present the various customized displays. They are also used for the tablespace analysis report.
Index Reorganization Thresholds	The index thresholds apply to the displays used in tab 5 of the Space Analyst and the queries that are run to present the various customized displays. They are also used for the tablespace analysis report.

The table below describes the options and functionality on the Space Analyst Thresholds Options dialog box for Oracle:

Option	Description
Tablespace and Datafile Fragmentation Thresholds	The tablespace and datafile thresholds apply to the tablespace and datafile fragmentation index. The default for each is 50 - meaning, flag any tablespace or datafile whose fragmentation index is less than 50. They are also used for the tablespace analysis report.
Table Reorganization Thresholds	The table thresholds apply to the displays used in tab 4 of the Space Analyst and the queries that are run to present the various customized displays. They are also used for the tablespace analysis report.

Option	Description
Index Reorganization Thresholds	The index thresholds apply to the displays used in tab 5 of the Space Analyst and the queries that are run to present the various customized displays. They are also used for the tablespace analysis report.

The table below describes the options and functionality on the Space Analyst Thresholds Options dialog box for Microsoft SQL Server (excluding SQL Server 2005) and Sybase:

Option	Description
Database and Log Space Thresholds	The database and log space thresholds apply to the database and log fragmentation index. The default for each is 80 - meaning, flag any database or log whose fragmentation index is less than 80. The database and log space thresholds apply to any grid that shows free space percent for those entities.
Table Reorganization Thresholds	The table thresholds apply to the displays used in tab 5 of the Space Analyst and the procedures that are run to present the various customized displays. They are also used for the database analysis report.
Index Reorganization Thresholds	The index thresholds apply to the displays used in tab 5 of the Space Analyst and the queries that are run to present the various customized displays. They are also used for the database analysis report.

The table below describes the options and functionality on the Space Analyst Thresholds Options dialog box for SQL Server 2005:

Option	Description
Database and Log Space Thresholds	The database and log space thresholds apply to the database and log fragmentation index. The default for each is 80 - meaning, flag any database or log whose fragmentation index is less than 80. The database and log space thresholds apply to any grid that shows free space percent for those entities.
Table Reorganization Thresholds	The table thresholds apply to the displays used in tab 5 of the Space Analyst and the procedures that are run to present the various customized displays. They are also used for the database analysis report.
Index Reorganization Thresholds	The index thresholds apply to the displays used in tab 5 of the Space Analyst and the queries that are run to present the various customized displays. They are also used for the database analysis report.
Fragmentation Analysis Mode	Detailed is the default. The other options are Sampled and Limited. This option serves as a parameter that's passed to the space diagnostic procedures use on the Objects tab for the reorganization summary and the detail views for indexes and tables.

Completing the Space Analyst Thresholds Options Dialog Box for Space Analyst

To complete this dialog box, do the following:

1 On the **Analysts Toolbar**, click the **Options** button.

The Space Analyst opens the dialog box.

2 Review thresholds.

3 Click **OK**.

>

Space Analyst for Oracle Tutorial

Embarcadero Space Analyst provides sophisticated diagnostic capabilities to troubleshoot bottlenecks and performance inefficiencies that result from poor space management. Running inside the Embarcadero DBArtisan console, Space Analyst helps you identify and correct storage problems by supplying powerful analysis facilities and corrective utilities and that will satisfy advanced experts, yet are easy enough to use for database novices.

About This Space Analyst for Oracle Tutorial

This tutorial helps you get started with Embarcadero Space Analyst.

After completion of this tutorial, you will have the foundation you need to explore the many features and benefits of Space Analyst. You will have learned how to get a summary view of storage for your database as well as how to identify potential space-related threats. Finally, you'll use the reorganization wizard to correct any space problems that you find.

This tutorial is divided into 5 sections.

Session 1: Getting Started with Space Analyst

Session 2: Obtaining a Global View of Storage

Session 3: Identifying Storage Structure Problems

Session 4: Pinpointing Object Performance Inefficiencies

Session 5: Correcting Space-Related Problems

You can use this basic tutorial as a roadmap of product highlights.

After starting this tutorial, you can select **Help** from the menu to find many additional resources that complement and build on many of the activities shown in this brief tutorial.

For more information, see Additional Evaluation Resources.

Session 1: Getting Started with Space Analyst for Oracle

Download and Install

You can obtain the latest version of Space Analyst by downloading DBArtisan from the Embarcadero Web site at http://www.embarcadero.com/downloads/downloaddbartisan.asp. Space Analyst runs fully contained within DBArtisan so you will need the DBArtisan product to use Space Analyst.

Click **Download**, and follow the steps indicated. Save the file on your computer and then double-click to launch the self-extracting file that will guide you through the installation process.

When you first install an evaluation copy of DBArtisan, you can use the tool for 14 days. After that time, a permanent license is needed.

Proceed to Space Analyst for Oracle Overview

Space Analyst for Oracle Overview

Once you are connected to an Oracle database, you can open Space Analyst from within DBArtisan two ways:

1 On the **Analyst** menu, select **Space Analyst**.

NOTE: You do not need to install anything for Space Analyst to work.

Proceed to Session 2: Obtaining a Global View of Storage Using Space Analyst for Oracle

Session 2: Obtaining a Global View of Storage Using Space Analyst for Oracle

When it opens, Space Analyst shows a complete overview of the storage situation of a particular database or database server. The first tab of Space Analyst displays all pertinent global statistics regarding storage structure usage throughout a database.

There are four distinct sections displayed on the first **Database Overview Tab** of Space Analyst:

- Database Overview
- Tablespace Space Overview
- Tablespace Fragmentation Overview
- Tablespace Objects Overview

Proceed to Space Analyst for Oracle - Database Overview

Space Analyst for Oracle - Database Overview

The Database Overview displays:

- The total count of all tablespaces, datafiles and temporary files.
- The percentage breakdown of the various object types in a database. Positioning the mouse pointer over the graph causes it to display the actual percentage amounts for each object type.
- The amount of space in kilobytes that each object type consumes is also shown. This section helps you get a clear picture of exactly how much space the various object types in your database are taking up.

> **TIP:** If you want to see a grid-style display of the information shown in the graph, position your mouse over the graph, click the right mouse button and select the ‘toggle display’ option.

Proceed to Space Analyst for Oracle - Tablespace Space Overview

Space Analyst for Oracle - Tablespace Space Overview

The Tablespace Space Overview section displays summary information regarding every tablespace in the database. When invoked in graph form, this section shows how much total space is currently allocated to each tablespace. When shown in grid form, this section shows not only the total space allocated for each tablespace, but also includes the used, free, and percent free amounts, along with an indicator that shows whether any of the tablespace’s underlying data or temp files can automatically grow if necessary.

> **TIP:** To see a grid-style display of the information shown in the graph, position your mouse over the graph, click the right mouse button and select the ‘toggle display’ option. To return to the graph, just perform the same actions over the grid.

When this section is in grid form, you can also perform a number of actions on a tablespace like open it in a DBArtisan editor, extract it's DDL, convert it to either local or dictionary-managed, and more. For example, to open a tablespace in one of DBArtisan's object editors, do the following:

1 In the grid, select a tablespace.

2 Right-click and select **Open**.

OR

Select a tablespace and click the **Open** toolbar button.

DBArtisan opens a Tablespaces Editor.

Proceed to Space Analyst for Oracle - Fragmentation Overview

Space Analyst for Oracle - Fragmentation Overview

The fragmentation overview provides a quick glance into how badly fragmented a tablespace is. Space Analyst computes a fragmentation index for every tablespace that can be used to quickly tell if fragmentation problems are plaguing a tablespace. A 100 percent reading indicates no fragmentation, while lower scores indicate the presence of fragmentation.

To see the percentage of fragmentation for each tablespace in the graph, do the following:

1 Position the pointer over each bar in the chart.

To see a grid-style display of the information shown in the graph, do the following:

1 Position the pointer over the graph.

2 Right-click and select **Toggle Display**.

Proceed to Space Analyst for Oracle - Tablespace Objects Overview

Space Analyst for Oracle - Tablespace Objects Overview

The tablespace objects overview provides a window into the types of objects that each tablespace contains. This view is useful because it helps you quickly understand if undesired objects are present in a tablespace.

NOTE: There is no graphical view available for this section.

Proceed to Session 3: Identifying Storage Structure Problems Using Space Analyst for Oracle

Session 3: Identifying Storage Structure Problems Using Space Analyst for Oracle

Space Analyst can help you quickly find storage structure problems that threaten the availability of your database.

Proceed to Pinpointing 'Out of Space' Problems Using Space Analyst for Oracle

Pinpointing 'Out of Space' Problems Using Space Analyst for Oracle

To quickly tell if any of the tablespaces in your database are about to run out of free space, do the following:

1. Select the **Database Overview Tab**.
2. In the **Tablespace Space Overview** section, position the pointer over the graph.
3. Right-click and select **Toggle Display**.
4. Examine the Percent Free columns of each tablespace along with the Autoextend column. If a tablespace is nearing the end of free space and cannot automatically grow, take action to prevent an out-of-space error.

Proceed to Locating Fragmentation Issues Using Space Analyst for Oracle

Locating Fragmentation Issues Using Space Analyst for Oracle

There are a number of facilities within Space Analyst that helps you understand storage structure fragmentation.

To get summary information regarding tablespace fragmentation, do the following:

1. Select the **Database Overview Tab**.
2. In the **Tablespace Fragmentation** section, position the pointer over the graph.
3. Right-click and select **Toggle Display**.

The tablespace fragmentation grid scores each tablespace in terms of fragmentation levels and also present data regarding the number of free chunks in a tablespace along with the largest free chunk size. Tablespaces with poor fragmentation ratings are highlighted in red. Note that you can sort the grid from worst to best tablespaces by clicking on the Fragmentation Index column header in the grid.

To get the same information for datafiles, do the following:

1. Select the **Data Files Tab**.

   The datafile fragmentation grid scores each datafile in terms of fragmentation levels and also present data regarding the number of free chunks in the datafile along with the largest free chunk size. Note that temp files are not shown in this display. Datafiles with poor fragmentation ratings are highlighted in red.

To determine the type of fragmentation for a tablespace (bubble or honeycomb), do the following:

1. Select a tablespace or datafile that appears to have fragmentation issues.
2. Right-click and select **Map**.

   Space Analyst displays a graphical rendering of the tablespace on the Tablespace Map Tab.

Bubble fragmentation can be confirmed by the presence of many free space fragments being found throughout the tablespace. Honeycomb fragmentation can be identified by seeing free space fragments adjacent to one another in the map.

> **NOTE:** Large tablespace maps can be read easier if the "Full Screen" option is chosen. The standard view may be returned from full screen by clicking on the small full screen dialog box that appears at the upper left side of the display.

Note that object extent fragmentation problems can also be identified in Space Analyst's tablespace map. To see if any table or index in a designated tablespace suffers from extent fragmentation, do the following:

1. Select the **Tablespace Map Tab**.
2. Select a tablespace.

3 In the **View By** section, select the **Extent Problems** option.

Space Analyst reconstructs the tablespace map and accompanying left-hand object grid to reveal any objects that have extent problems. All objects with extent problems are color-coded in red for easy identification. Objects colored in green indicate that they doe not suffer from any extent problems.

Proceed to Using Analysis Reports in Space Analyst for Oracle

Using Analysis Reports in Space Analyst for Oracle

You can run HTML-formatted reports that pinpoint all major storage-related problems for a particular tablespace and the objects that are contained within it. To produce an Analysis Report, do the following:

1 On the toolbar, click **Analysis Report**.

Space Analyst opens the Analysis Report dialog box.

2 In **Report File Name** and **Report Title** you can change the location and report title of the report.

3 In **Category**, select the types of space problems you would like to investigate. The default option is **All**.

4 In **Tablespace**, select the tablespace you would like to analyze.

5 Click **Execute**.

Space Analyst creates a frame-based Analysis report for the tablespace.

Proceed to Session 4: Pinpointing Object Performance Inefficiencies Using Space Analyst for Oracle

Session 4: Pinpointing Object Performance Inefficiencies Using Space Analyst for Oracle

Space Analyst can greatly accelerate the process of locating objects in your database that are degrading overall response times because of their current storage/space situation.

Proceed to Finding Space-Related Table Problems with Space Analyst for Oracle

Finding Space-Related Table Problems with Space Analyst for Oracle

To quickly determine if any tables in your database may be contributing to performance slowdowns, do the following:

1 In **DBArtisan** connect to the database you would like to check.

2 Open **Space Analyst**.

3 Select the **Tables Tab**.

Space Analyst runs diagnostics against your database and graphically display counts of tables that are found to contain various types of space problems.

The graph at the top of the Tables Tab shows the overall count of problem tables, but you can also view the details for each type of problem by using the Details grid at the bottom of the tab. For example, if Space Analyst identifies tables that have block inefficiencies, to find the actual tables and view details regarding the identified space problem, do the following in the Details grid:

1 In the **Tablespace** list, select **All**.

1 In the **Table Owner** list, select **All**.

2 In the **Display** list, select **Tables with inefficient block usage**.

3 Click **Execute**.

Space Analyst displays the actual tables that suffer from block inefficiencies along with their demographic details. To take corrective action on any object, select one or more tables and click the right mouse button to bring up a menu of possible actions.

Proceed to Finding Space-Related Index Problems with Space Analyst for Oracle

Finding Space-Related Index Problems with Space Analyst for Oracle

To quickly determine if any indexes in your database may be contributing to performance slowdowns by following these steps:

1 In **DBArtisan** connect to the database you would like to check.

2 Open **Space Analyst**.

3 Select the **Indexes Tab**.

The graph at the top of the Indexes Tab shows the overall counts of problem indexes, but you can also view the details for each type of problem by using the Details grid at the bottom of the tab. For example, if Space Analyst identifies indexes that have next extent problems, then to find the actual indexes and view details regarding the space problem, do the following in the Details grid:

1 In the **Tablespace** list, select **All**.

2 In the **Index Owner** list, select **All**.

3 In the **Display** list, select **Indexes with NEXT_EXTENT space deficits**.

4 Click **Execute**.

Space Analyst displays the actual indexes that suffer from next extent space problems along with their demographic details. If you wanted to take corrective action on any object, select one or more indexes and click the right mouse button to bring up a menu of possible actions.

NOTE: Another critical statistic in determining index problems is deleted leaf rows. To display deleted leaf row statistics in the Details grid, do the following:

1 On the **Command** toolbar, select **Get Deleted Leaf Row Statistics**.

Proceed to Session 5: Correcting Space-Related Problems Using Space Analyst for Oracle

Session 5: Correcting Space-Related Problems Using Space Analyst for Oracle

After you have identified space-related bottlenecks in your database with Space Analyst, you can use a number of different tools and facilities within the product to eliminate each storage problems.

Proceed to Correcting Storage Structure 'Out of Space' Problems for Space Analyst for Oracle

Correcting Storage Structure 'Out of Space' Problems for Space Analyst for Oracle

Space Analyst can help you quickly remedy tablespaces that are running out of room by:

- Resizing an existing tablespace to have more room.
- Adding a new datafile to an existing tablespace.
- Setting the Autoextend attribute for one or more of a tablespace's datafiles to allow for automatic growth.

To manage these functions, do the following:

1. Select the **Database Overview Tab**.
2. In the **Tablespace Space Overview** section, right-click the graph, and then select **Toggle Display**.
3. Select the tablespace that is running low on free space.
4. Select the **Open** toolbar button or the **Open** shortcut menu option.

   DBArtisan opens a Tablespace Editor.

In the Tablespace Editor, you can remedy the tablespace problem three ways:

- To resize an existing tablespace, open the tablespace's datafile in the editor (using the **Edit** button) and change the datafile's size to a new, larger size. Then close the editor and select the **Alter** toolbar button to make the change permanent to the database.
- To add a new datafile to the tablespace, select the **Add** button in the Datafile Tab of the tablespace editor, and specify the new datafile's attributes. Once complete, close the Add Datafile dialog and select the **Alter** toolbar button to make the change permanent to the database.
- To set Autoextend for one or more of a tablespace's datafiles, open each datafile in the datafile editor by selecting the **Edit** button. Once in the Modify Datafile dialog, select the option for **Auto Extend** and specify other attributes as needed. When complete, close the dialog and select the **Alter** toolbar button to make the change permanent to the database.

Proceed to Correcting Storage Structure Fragmentation Issues Using Space Analyst for Oracle

Correcting Storage Structure Fragmentation Issues Using Space Analyst for Oracle

Space Analyst can help you quickly fix tablespaces that are suffering from fragmentation by:

- Coalescing adjacent free space fragments (honeycombs) in a tablespace
- Reorganizing a tablespace to eliminate free space bubbles
- Converting a dictionary-managed tablespace to locally-managed

To coalesce free space fragments in a tablespace, follow these steps:

1. On the **Database Overview Tab**, select any tablespace that appears to be suffering from fragmentation problems.
2. Select the **Open** toolbar button or the **Open** shortcut menu option.

   DBArtisan opens a Tablespaces Editor.
3. To view fragmentation details for the tablespace, click the **Space Tab** of the Tablespaces Editor.

4 To combine adjacent free space pockets into singular free space chunks, click the **Coalesce** button.

NOTE: The fragmentation score of the tablespace does not change after a coalesce operation, then the tablespace likely suffers from bubble fragmentation. This can be verified by viewing the tablespace map for the tablespace.

To reorganize any tablespace suffering from bubble fragmentation, do the following:

1 On the toolbar, click the **Open Wizard** button.

Space Analyst opens the Reorganization Wizard.

2 Select the **New Reorganization** option, and then click **Next** to continue.

Space Analyst opens the next panel of the Reorganization Wizard.

3 Select **Reorganize a Tablespace**, and then click **Next** to continue.

Space Analyst opens the next panel of the Reorganization Wizard.

4 Select one or more tablespaces you would like to reorganize, and then click **Next** to continue.

Space Analyst opens the next panel of the Reorganization Wizard.

5 Select a scratch tablespace that temporarily holds the reorganized tablespace's objects.

6 Select the option to perform a space analysis so you can be assured that the scratch tablespace has enough room to temporarily hold all of the reorganized tablespace's objects.

7 If the selected scratch tablespace does not contain enough free space, select another tablespace to serve as the scratch tablespace.

8 Click **Next** to continue.

Space Analyst opens the next panel of the Reorganization Wizard.

9 You can choose to perform no customizations for the reorganization, or you can select various global or object-level options to customize your reorganization.

10 Click **Finish**.

Space Analyst builds a reorganization script to reorganize your selected tablespace. Note that while the script is being constructed, you can still perform other work in Space Analyst or DBArtisan. Just minimize the script build window and continue to do other work in DBArtisan. You can check back whenever you'd like to see if the script is ready.

After the reorganization script is complete, you can choose to view/run the script in a DBArtisan ISQL window, schedule the script for later execution, or immediately execute the reorganization.

If you would like to convert a dictionary-managed tablespace to locally-managed, you can do one of the following:

- Create a new locally-managed tablespace in DBArtisan and then use the Reorganization Wizard to permanently relocate all objects from the dictionary-managed tablespace to the new locally-managed tablespace. This is the cleanest method for converting a dictionary-managed tablespace to locally-managed.
- Choose the option in Space Analyst that allows a quick conversion of an existing dictionary-managed tablespace to a locally-managed tablespace. This option is only available for Oracle versions 8.1.6 or later. This option does not defragment an existing dictionary-managed tablespace, but converts it so that all space management is conducted through the tablespace's datafiles instead of the data dictionary and ensure that any new objects added to the tablespace conform to the tablespace's storage attributes.

Proceed to Correcting Space-Related Object Problems for Space Analyst for Oracle

Correcting Space-Related Object Problems for Space Analyst for Oracle

Space Analyst can help you quickly fix objects that are experiencing space problems by:

- Modifying various storage-related properties for an object.
- Reorganizing the object to eliminate space inefficiencies.

Storage properties that are affecting the availability of an object (for example, the next extent or maximum extent values for a table or index) can easily be altered by using the integration that exists between Space Analyst and DBArtisan. For example, to modify a table's next extent value, do the following:

1. On the **Tablespace Map Tab** or the **Table Details** grid on the **Tables Tab**, select a table you would like to alter.
2. Select the **Open** toolbar button or the **Open** shortcut menu option.

   DBArtisan opens a Tablespaces Editor.
3. Select the **Storage Tab**.
4. Change the numeric value for the table's next extent.

   NOTE: You can not change next extent values for tables that exist in locally-managed tablespaces.
5. On the toolbar, click the **Alter** button.

   DBArtisan displays the DDL that change the table's next extent value.

You can easily reorganize one or more objects with the Reorganization Wizard. To reorganize any table or index that in any grid, do the following:

1. Select the object(s).
2. On the toolbar, click the **Reorganization Wizard** button, or right-click and select **Reorganize**.

   Space Analyst opens the Reorganization Wizard.

You can customize object reorganizations by using either the global or specific object customization options in the Reorganization Wizard. For example, you can perform online reorganizations of indexes or index-organized tables.

Space Analyst's Reorganization Wizard can also intelligently identify objects that need to be reorganized. For example, to reorganize tables in your database that have chained/migrated rows, do the following:

1. On the toolbar, click the **Reorg Wizard** button.

   Space Analyst opens the Reorganization Wizard.
2. Select the **New Reorganization** option, and then click **Next** to continue.

   Space Analyst opens the next panel of the Reorganization Wizard.
3. Select **Repair tables with chained/migrate rows**, and then click **Next** to continue.

   Space Analyst runs diagnostics against your database to find any table or table partition that suffers from chained rows. Space Analyst displays the objects under either the Tables or Table Partition(s) node in the Reorg Wizard explorer tree. You can then select one or all table objects and continue with the reorganization process.

>

Space Analyst for Microsoft SQL Server Tutorial

Embarcadero Space Analyst provides sophisticated diagnostic capabilities to troubleshoot bottlenecks and performance inefficiencies that result from poor space management. Running inside the Embarcadero DBArtisan console, Space Analyst helps you identify and correct storage problems by supplying powerful analysis facilities and corrective utilities and that will satisfy advanced experts, yet are easy enough to use for database novices.

Space Analyst for Microsoft Product Benefits

Embarcadero Space Analyst allows the database administrator to quickly ascertain if space problems are threatening the availability or performance of their key databases. If storage problems are found, Space Analyst provides powerful but easy to use space management wizards and utilities that can accomplish simple or complex reorganizations of all or parts of a database.

About this Space Analyst for Microsoft Tutorial

This tutorial helps you get started with Embarcadero Space Analyst.

After completion of this tutorial, you will have the foundation you need to explore the many features and benefits of Space Analyst. You will have learned how to get a summary view of storage for your database as well as how to identify potential space-related threats. Finally, you'll use the reorganization wizard to correct any space problems that you find.

This tutorial is divided into 5 sections.

Session 1: Getting Started with Space Analyst

Session 2: Obtaining a Global View of Storage

Session 3: Identifying Storage Structure Problems

Session 4: Pinpointing Object Performance Inefficiencies

Session 5: Correcting Space-Related Problems

You can use this basic tutorial as a roadmap of product highlights.

After starting this tutorial, you can select **Help** from the menu to find many additional resources that complement and build on many of the activities shown in this brief tutorial.

For more information, see Additional Evaluation Resources.

Session 1: Getting Started with Space Analyst for Microsoft

Download and Install

You can obtain the latest version of Space Analyst by downloading DBArtisan from the Embarcadero Web site at http://www.embarcadero.com/downloads/downloaddbartisan.asp. Space Analyst runs fully contained within DBArtisan so you will need the DBArtisan product to use Space Analyst.

Click **Download**, and follow the steps indicated. Save the file on your computer and then double-click to launch the self-extracting file that will guide you through the installation process.

When you first install an evaluation copy of DBArtisan, you can use the tool for 14 days. After that time, a permanent license is needed.

Proceed to Space Analyst for Microsoft Overview

>

Space Analyst for Microsoft Overview

Once you are connected to an Microsoft SQL Server database, you can open Space Analyst from within DBArtisan two ways:

1 On the **Analyst** menu, select **Space Analyst**.

NOTE: You do not need to install anything for Space Analyst to work.

Proceed to Proceed to Session 2: Obtaining a Global View of Storage Using Space Analyst for Microsoft

Session 2: Obtaining a Global View of Storage Using Space Analyst for Microsoft

When first invoked, Space Analyst displays a complete overview of the storage situation of a particular SQL Server. The first tab of Space Analyst can be used to see all pertinent global statistics regarding storage structure usage throughout a server.

There are four distinct sections displayed on the first tab (labeled Storage Overview) of Space Analyst:

Server Overview

Database Overview

Log Overview

Server Disk Overview

Space Analyst for Microsoft - Server Overview

The total count of all databases, file groups, and files are shown in this section. You can also see the percentage breakdown of the database vs. log space on the server. Positioning the mouse pointer over the graph causes it to display the actual percentage amounts for storage type.

TIP: If you want to see a grid-style display of the information shown in most graphs, position your mouse over the graph, click the right mouse button and select the 'toggle display' option.

Proceed to Space Analyst for Microsoft - Database Overview

Space Analyst for Microsoft - Database Overview

Summary information regarding every database on the server is present in the database overview section. When invoked in graph form, this section shows how much total space is used by each database. When shown in grid form, this section shows not only the total space used for each database, but also includes the total allocated, free, and percent free amounts, along with an indicator that shows whether any of the database's underlying datafiles can automatically grow if necessary.

TIP: To see a grid-style display of the information shown in the graph, position your mouse over the graph, click the right mouse button and select the ‘toggle display’ option. To return to the graph, just perform the same actions over the grid.

>

When this section is in grid form, you can also perform a number of actions on a database like open it in a DBArtisan editor, extract its DDL, shrink it, and more. For example, if you want to open a database in DBArtisan's database manager, do the following:

1 In the grid, select a database.

2 Right-click and select **Open**.

OR

3 Select a database and click the **Open** toolbar button.

DBArtisan opens the Database Manager.

Proceed to Space Analyst for Microsoft - Log Overview

Space Analyst for Microsoft - Log Overview

The log overview provides a quick glance into database transaction log usage. When invoked in graph form, this section shows the percentage of used log space per database. When shown in grid form, this section shows not only the total percentage of transaction log space used for each database, but also includes the total allocated and free amounts (in MB), along with an indicator that shows whether any of the database's transaction logs underlying log files can automatically grow if necessary.

Proceed to Space Analyst for Microsoft - Server Disk Overview

Space Analyst for Microsoft - Server Disk Overview

The server disk overview provides a window into how much free space exists on the actual server machine. In graph form, this section shows the amount of free space (in MB) across each of the server's drives. When in grid form, it shows free space per drive, but also includes the amount of allocated SQL Server data and log space per drive.

Proceed to Session 3: Identifying Storage Structure Problems in Space Analyst for Microsoft

Session 3: Identifying Storage Structure Problems in Space Analyst for Microsoft

Space Analyst can help you quickly find storage structure problems that threaten the availability of your database.

Proceed to Pinpointing ‘Out of Space’ Problems Using Space Analyst for Microsoft

Pinpointing ‘Out of Space’ Problems Using Space Analyst for Microsoft

To quickly tell if any of the tablespaces in your database are about to run out of free space, do the following:

1 Select the **Storage Overview Tab**.

2 In the **Database Overview** and **Log Overview** sections, position the pointer over the graph.

3 Right-click and select **Toggle Display**.

Examine the Percent DB Used columns of each database and transaction log along with the DB Can Grow column. If a database or transaction log is nearing the end of free space and cannot automatically grow, then action should be taken to prevent an out-of-space error.

You can also obtain graphical views of databases and pertinent storage property information by viewing the Databases Tab of Space Analyst. You can visually get a handle on the amounts of table, index, and free space of your databases and quickly view critical storage property facts such as the number of databases that cannot automatically grow in size if needed.

All the graphs can be toggled to grids by using the right mouse menu option of Toggle Display. When in grid mode, very detailed information is available concerning database and transaction log space. Note that the graph is scrollable so you may have to scroll to see all your databases.

You can perform a number of actions on a database, such as shrinking it or invoking database-related DBCC operations by using either the right mouse menu or the Command toolbar button (the yellow arrow).

Proceed to Viewing Filegroup and File Details Using Space Analyst for Microsoft

Viewing Filegroup and File Details Using Space Analyst for Microsoft

You can also view space usage at the filegroup and file levels. To get summary and detailed information for filegroups and files, follow these steps:

1. Invoke Space Analyst and view the **File Groups/Files Tab**.

   The default presentation is a graphical overview of filegroups. Note that the graph is scrollable so you may have to scroll to see all your filegroups.

2. To view detailed information for filegroups, right-click and select **Toggle Display** to view the grid.

3. To view file information instead of filegroup metrics, select **Files**.

When viewing filegroup and file information, you can perform a number of actions on selected filegroups and files such as extracting their DDL or shrinking the parent database or selected file. All options are available by using either the shortcut menu or Command menu.

Proceed to Using Analysis Reports in Space Analyst for Microsoft

Using Analysis Reports in Space Analyst for Microsoft

You can run HTML-formatted reports that pinpoint all major storage-related problems for a particular tablespace and the objects that are contained within it. To produce an Analysis Report, do the following:

1. On the toolbar, click **Analysis Report**.

   Space Analyst opens the Analysis Report dialog box.

2. In **Report File Name** and **Report Title** you can change the location and report title of the report.

3. In **Category**, select the types of space problems you would like to investigate. The default option is **All**.

4. In **Database**, select the database you would like to analyze.

5. In **File Group**, select the filegroup level.

6. Click **Execute**.

   Space Analyst creates a frame-based Analysis report for the database.

Proceed to Session 4: Pinpointing Object Performance Inefficiencies Using Space Analyst for Microsoft

Session 4: Pinpointing Object Performance Inefficiencies Using Space

Analyst for Microsoft

Space Analyst can greatly accelerate the process of locating objects in your databases that are degrading overall response times because of their current inefficient storage usage.

Proceed to Finding Space-Related Table Problems with Space Analyst for Microsoft

Finding Space-Related Table Problems with Space Analyst for Microsoft

To quickly determine if any tables in your database may be contributing to performance slowdowns, do the following:

1. Open **Space Analyst**.
2. Select the **Objects Tab**.
3. In **Object Details**, select the target object.

By default, the General Space button is selected. This view will show you a graphical layout of table used and free space for the database. Note that the graph is scrollable so you can scroll to see all the tables for your database.

The Object details grid in the lower area of the tab can be used to view detailed storage information for every table in the selected database.

To troubleshoot table storage issues in your database, select the Space Diagnostics radio button. This option will interrogate your database and will provide you with a graphical count of the number of tables that are experiencing specific types of storage inefficiencies. Space Analyst uses predefined thresholds (which you can customize by selecting the Threshold Options toolbar button) to determine how many and what objects are suffering from storage inefficiencies.

The table below describes the types of table-related space problems shown in the bar graph found on the Objects Tab:

Problem	Description
Scan Density	This is computed by dividing the optimum number of extent switches by the actual number of extent switches. This percentage can be misleading, however, if the object spans more than one database file and should not be considered in such an event. Values close to 100% are best.
Logical Fragmentation	This metric represents the percentage of pages that are out of physical order. High values (greater than 25%) for this metric may mean that index scan performance is not what it could be. Note that this statistic is meaningless for tables without clustered indexes (heap tables).
Extent Fragmentation	This highlights any gaps (lack of contiguous order) in extents and indicates the presence of external fragmentation. Percentages of 25-50% or more are not a favorable reading.
Average Page Density	This metric represents the fullness of the data or index pages, and is an indicator of internal fragmentation. Remember that the more full a data or index page is, the less I/O needed by SQL Server when performing scans of tables or indexes. High values are good here, with anything below 50% being suspect. Low values for this metric often indicate the presence of page splitting. Keep in mind, however, that internal fragmentation isn't necessarily bad in OLTP environments where large FILLFACTORs are often specified. Also understand that small objects (those will little or no data) will likely show low readings because they simply don't consume enough space in a database page. You can ignore these objects.

Problem	Description
Forwarded Records	SQL Server will move rows in a table under certain conditions. One situation might arise when you update a row in a table that has a variable-length column to a larger size that will no longer fit on its original page. Another situation would be if SQL Server moves a row when the clustered index column changes. Forwarded records can reduce performance at times because additional I/O is involved to first obtain the record pointer to the relocated row, and then the row itself. But when does row forwarding become a problem? For example, just because a table has one hundred forwarded rows, does that mean a performance problem exists? The answer is: it depends. If the table has one hundred rows in it, then yes, a problem exists - 100% of the table suffers from forwarded rows. If the table has three million rows, then the forwarded row situation is likely not causing much fuss in the database.

The graph at the top of the Objects Tab shows the overall count of problem tables, but you can also view the details for each type of problem by using the Details grid at the bottom of the tab. For example, if Space Analyst identifies tables that have forwarded records, then to find the actual tables and view details regarding the identified space problem, do the following:

1. In the **File Group** list, select **All**.
2. In the **Owner** list, select **All**.
3. In the **Display** list, select **Tables with forwarded records problems**.
4. Click **Refresh**.

You would then see the actual tables that suffer from forwarded records along with their demographic details. If you wanted to take corrective action on any object, you could highlight one or more tables and click the right mouse button to bring up a menu of possible actions.

Proceed to Session 5: Correcting Space-Related Problems Using Space Analyst for Microsoft

Session 5: Correcting Space-Related Problems Using Space Analyst for Microsoft

After you have identified space-related bottlenecks in your database with Space Analyst, you can use a number of different tools and facilities within the product to eliminate each storage problems.

Proceed to Correcting Storage Structure ‘Out of Space’ Problems Using Space Analyst for Microsoft

Correcting Storage Structure ‘Out of Space’ Problems Using Space Analyst for Microsoft

Space Analyst can help you quickly remedy databases and transaction logs that are running out of room by:

- Resizing an existing database or transaction log to have more room.
- Adding a new filegroup or file to an existing database or transaction log.
- Setting the Autogrowth attribute for one or more of a database's or transaction log's files to allow for automatic growth.

To manage these functions, do the following:

1. Select the **Databases Tab**.
2. In the **Database Summary** section, right-click the graph, and then select **Toggle Display**.
3. Select the database that is running low on free space.
4. Select the **Open** toolbar button or the **Open** shortcut menu option.

   DBArtisan opens the Database Manager.

In the Database Manager, you can remedy the database problem in one of three ways:

- To resize an existing database or transaction log, open one of the database's or log's files in the editor (using the **Edit** button) and change the file's size to a new, larger size. Then close the editor and select the **Alter** toolbar button to make the change permanent to the database.
- To add a new file or file group to the database or transaction log (note: you can't add new filegroups to a transaction log, only files), select the **Add** button in the editor, and specify the new filegroup or file's attributes. Once complete, close the Add Datafile dialog and select the **Alter** toolbar button to make the change permanent to the database.
- To set Autogrowth for one or more of a database's or transaction log's files, open each file in the file editor by selecting the file and then pressing the **Edit** button. Once in the Modify File dialog, select the option for **Allow Growth** and specify other attributes as needed. When complete, close the dialog and select the **Alter** toolbar button to make the change permanent to the database.

Proceed to Removing Wasted Space in Databases or Files Using Space Analyst for Microsoft

Removing Wasted Space in Databases or Files Using Space Analyst for Microsoft

Space Analyst can help you quickly eliminate unused space in a database or database file. The unused space can be released back to the operating system for other use. To remove unused space from a database, do the following:

1. To shrink an entire database, right-click a database in any of Space Analyst's grids, and then select **Shrink Database**.

   DBArtisan opens the Shrink File dialog box.
2. To shrink a specific file in a database, select the **Filegroups/Files Tab** and then view files by selecting the Files radio button.
3. Right-click a file and select **Shrink File**.

   DBArtisan opens the Shrink File dialog box.
4. Select desired options.
5. Click **Execute** to make the change permanent.

Proceed to Correcting Space-Related Object Problems for Space Analyst for Microsoft

Correcting Space-Related Object Problems for Space Analyst for Microsoft

Once you identify specific objects that need storage corrections, you can use Space Analyst to help you quickly resolve your found issues. The Space Management wizard inside of Space Analyst can be used to build reorganization and space maintenance jobs that run right the first time. All jobs can either be immediately run or scheduled for later execution.

To create a standard database or object reorganization job, do the following:

1. Click the **Open Wizard** button.
2. Select **New Space Management Task**.
3. Click **Next**.
4. Select **Create standard reorganization job**.
5. Click **Next**.
6. Select the databases and objects you would like to reorganize from the Explorer tree on the left hand side of the wizard. You can reorganize full databases, selected schema accounts, or selected objects.
7. Click **Next**.
8. The next panel allows you to customize the reorganization. You can either set options at the global level (affects all objects) or specific object level. Using the options, you can relocate objects to different filegroups, perform online index rebuilds (for SQL Server 2000 or later), and specify other customizations.
9. Click **Next**.
10. The next panel builds your reorganization script and allows you to preview the operation. Click Next when ready.
11. Click **Next**.

The final panel allows you to open the reorganization script in a SQL editor, immediately run it, or schedule it for later execution.

Proceed to Proactive Space Management Using Space Analyst for Microsoft

Proactive Space Management Using Space Analyst for Microsoft

If you have many databases to take care of, then you might want to relieve the burden of performing detailed storage analysis on every database, each of which may or may not require attention. Fortunately, Space Analyst makes it easy to create smart space management jobs that can automate this task for you.

You can create two specialized space management jobs that allow you to proactively eliminate space threats before they affect your database:

You can create a "reorganization analyst job" that allows you to specifically select various databases and/or objects and then set reorganization thresholds that determine when a reorganization will occur. Once constructed, you can schedule these jobs to run as often as you'd like. During execution, these jobs will interrogate your chosen database/objects and determine if a reorganization needs to occur by applying your customized reorganization thresholds. If any objects are found to need reorganization, they are dynamically reorganized using any customized options you specify.

You can build space maintenance jobs that will perform functions such as check the validity of your databases and filegroups, repair any found minor corruption, and eliminate unused space in your databases. You can also set the job up to notify you if any problems are found during job execution.

These "smart" space maintenance jobs allow you to truly "set and forget" the task of monitoring and reorganizing the space of your critical databases.

Space Analyst for Sybase ASE Tutorial

Embarcadero Space Analyst provides sophisticated diagnostic capabilities to troubleshoot bottlenecks and performance inefficiencies that result from poor space management. Running inside the Embarcadero DBArtisan console, Space Analyst helps you identify and correct storage problems by supplying powerful analysis facilities and corrective utilities and that will satisfy advanced experts, yet are easy enough to use for database novices.

Space Analyst for Sybase ASE Product Benefits

Embarcadero Space Analyst allows the database administrator to quickly ascertain if space problems are threatening the availability or performance of their key databases. If storage problems are found, Space Analyst provides powerful but easy to use space management wizards and utilities that can accomplish simple or complex reorganizations of all or parts of a database.

About this Space Analyst for Sybase ASE Tutorial

This tutorial helps you get started with Embarcadero Space Analyst.

After completion of this tutorial, you will have the foundation you need to explore the many features and benefits of Space Analyst. You will have learned how to get a summary view of storage for your database as well as how to identify potential space-related threats. Finally, you'll use the reorganization wizard to correct any space problems that you find.

This tutorial is divided into 4 sections.

Session 1: Getting Started with Space Analyst

Session 2: Obtaining a Global View of Storage

Session 3: Identifying Storage Structure Problems

Session 4: Correcting Space-Related Problems

You can use this basic tutorial as a roadmap of product highlights.

After starting this tutorial, you can select **Help** from the menu to find many additional resources that complement and build on many of the activities shown in this brief tutorial.

For more information, see Additional Evaluation Resources.

Session 1: Getting Started with Space Analyst for Sybase ASE

Download and Install

You can obtain the latest version of Space Analyst by downloading DBArtisan from the Embarcadero Web site at http://www.embarcadero.com/downloads/downloaddbartisan.asp. Space Analyst runs fully contained within DBArtisan so you will need the DBArtisan product to use Space Analyst.

Click **Download**, and follow the steps indicated. Save the file on your computer and then double-click to launch the self-extracting file that will guide you through the installation process.

When you first install an evaluation copy of DBArtisan, you can use the tool for 14 days. After that time, a permanent license is needed.

Proceed to Space Analyst for Sybase ASE Overview

Space Analyst for Sybase ASE Overview

The graphic below illustrates all the main elements of the DBArtisan User Interface.

Once you are connected to an Sybase ASE database, you can open Space Analyst from within DBArtisan two ways:

1 On the **Analyst** menu, select **Space Analyst**.

NOTE: You do not need to install anything for Space Analyst to work.

Proceed to Session 2: Obtaining a Global View of Storage Using Space Analyst for Sybase ASE

Session 2: Obtaining a Global View of Storage Using Space Analyst for Sybase ASE

When first invoked, Space Analyst displays a complete overview of the storage situation of a particular server. The first tab of Space Analyst can be used to see all pertinent global statistics regarding storage structure usage throughout a server.

There are four distinct sections displayed on the Storage Overview Tab of Space Analyst:

Server Overview

Log Overview

Database Overview

Device Overview

Space Analyst for Sybase ASE - Server Overview

The total count of all databases, file groups, and files are shown in this section. You can also see the percentage breakdown of the database vs. log space on the server. Positioning the mouse pointer over the graph causes it to display the actual percentage amounts for storage type.

TIP: If you want to see a grid-style display of the information shown in most graphs, position your mouse over the graph, click the right mouse button and select the 'toggle display' option.

Proceed to Space Analyst for Sybase ASE - Log Overview

Space Analyst for Sybase ASE - Log Overview

The log overview provides a quick glance into database transaction log usage. When invoked in graph form, this section shows the percentage of used log space per database. When shown in grid form, this section shows not only the total percentage of transaction log space used for each database, but also includes the total allocated and free amounts (in MB), along with an indicator that shows whether any of the database's transaction logs underlying log files can automatically grow if necessary.

Proceed to Space Analyst for Sybase ASE - Database Overview

Space Analyst for Sybase ASE - Database Overview

Summary information regarding every database on the server is present in the database overview section. When invoked in graph form, this section shows how much total space is used by each database. When shown in grid form, this section shows not only the total space used for each database, but also includes the total allocated, free, and percent free amounts, along with an indicator that shows whether any of the database's underlying datafiles can automatically grow if necessary.

> **TIP:** To see a grid-style display of the information shown in the graph, position your mouse over the graph, click the right mouse button and select the 'toggle display' option. To return to the graph, just perform the same actions over the grid.

When this section is in grid form, you can also perform a number of actions on a database like open it in a DBArtisan editor, extract its DDL, shrink it, and more. For example, if you want to open a database in DBArtisan's database manager, do the following:

1 Select a database, right-click, and select **Open**.

OR

Select a database and click the **Open** toolbar button.

DBArtisan opens the Database Manager.

Proceed to Space Analyst for Sybase ASE - Device Overview

Space Analyst for Sybase ASE - Device Overview

The device overview section contains a graphical and grid control that works in the following way:

- When the graph is displayed, a bar chart is shown that lists the total space used (in MB) of all devices defined to the server.
- When the grid is displayed, detailed storage information for each database device is presented:
  - Device name
  - Physical name
  - Total space
  - Used space
  - Free space
  - Percent free
  - Percent used

Proceed to Session 3: Identifying Storage Structure Problems in Space Analyst for Sybase ASE

Session 3: Identifying Storage Structure Problems in Space Analyst for Sybase ASE

Space Analyst can help you quickly find storage structure problems that threaten the availability of your database.

Proceed to Pinpointing 'Out of Space' Problems Using Space Analyst for Sybase ASE

Pinpointing ‘Out of Space’ Problems Using Space Analyst for Sybase ASE

To quickly tell if any of the tablespaces in your database are about to run out of free space, do the following:

1. Select the **Storage Overview Tab**.
2. In the **Database Overview** and **Log Overview** sections, position the pointer over the graph.
3. Right-click and select **Toggle Display**.

Examine the Database Percent Free columns of each database and Log Percent Free for each transaction log. If a database or transaction log is nearing the end of free space and cannot automatically grow, then action should be taken to prevent an out-of-space error.

You can also obtain graphical views of databases and pertinent storage property information by viewing the Databases Tab of Space Analyst. You can visually get a handle on the amounts of table, index, and free space of your databases and quickly view critical storage property facts such as the number of databases that cannot automatically grow in size if needed.

All the graphs can be toggled to grids by using the right mouse menu option of 'Toggle Display'. When in grid mode, very detailed information is available concerning database and transaction log space. Note that the graph is scrollable so you may have to scroll to see all your databases.

Proceed to Using Analysis Reports in Space Analyst for Sybase ASE

Using Analysis Reports in Space Analyst for Sybase ASE

You can run HTML-formatted reports that pinpoint all major storage-related problems for a particular tablespace and the objects that are contained within it. To produce an Analysis Report, do the following:

1. On the toolbar, click **Analysis Report**.

   Space Analyst opens the Analysis Report dialog box.
2. In **Report File Name** and **Report Title** you can change the location and report title of the report.
3. In **Category**, select the types of space problems you would like to investigate. The default option is **All**.
4. In **Database**, select the database you would like to analyze.
5. In **File Group**, select the filegroup level.
6. Click **Execute**.

   Space Analyst creates a frame-based Analysis report for the database.

Proceed to Session 4: Correcting Space-Related Problems Using Space Analyst for Sybase ASE

Session 4: Correcting Space-Related Problems Using Space Analyst for Sybase ASE

After you have identified space-related bottlenecks in your database with Space Analyst, you can use a number of different tools and facilities within the product to eliminate each storage problems.

Proceed to Correcting Storage Structure ‘Out of Space’ Problems Using Space Analyst for Sybase ASE

Correcting Storage Structure ‘Out of Space’ Problems Using Space Analyst for

Sybase ASE

Space Analyst can help you quickly remedy databases and transaction logs that are running out of room by:

- Resizing an existing database or transaction log to have more room.
- Adding a new filegroup or file to an existing database or transaction log.
- Setting the Autogrowth attribute for one or more of a database's or transaction log's files to allow for automatic growth.

To manage these functions, do the following:

1. Select the **Storage Overview Tab**.
2. In the **Database Overview** or **Log Overview** section, right-click the graph, and then select **Toggle Display**.
3. Select the database that is running low on free space.
4. Select the **Open** toolbar button or the **Open** shortcut menu option.

   DBArtisan opens the Database Manager.

In the Database Manager, you can remedy the database problem in one of three ways:

- To resize an existing database or transaction log, open one of the database's or log's files in the editor (using the **Edit** button) and change the file's size to a new, larger size. Then close the editor and select the **Alter** toolbar button to make the change permanent to the database.
- To add a new file or file group to the database or transaction log (note: you can't add new filegroups to a transaction log, only files), select the **Add** button in the editor, and specify the new filegroup or file's attributes. Once complete, close the Add Datafile dialog and select the **Alter** toolbar button to make the change permanent to the database.
- To set Autogrowth for one or more of a database's or transaction log's files, open each file in the file editor by selecting the file and then pressing the **Edit** button. Once in the Modify File dialog, select the option for **Allow Growth** and specify other attributes as needed. When complete, close the dialog and select the **Alter** toolbar button to make the change permanent to the database.

Proceed to Correcting Space-Related Object Problems for Space Analyst for Sybase ASE

Correcting Space-Related Object Problems for Space Analyst for Sybase ASE

Once you identify specific objects that need storage corrections, you can use Space Analyst to help you quickly resolve your found issues. The Space Management wizard inside of Space Analyst can be used to build reorganization and space maintenance jobs that run right the first time. All jobs can either be immediately run or scheduled for later execution.

To create a standard database or object reorganization job, do the following:

1. Click the **Open Wizard** button.
2. Select **New Space Management Task**.
3. Click **Next**.
4. Select **Create standard reorganization job**.
5. Click **Next**.
6. Select the databases and objects you would like to reorganize from the Explorer tree on the left hand side of the wizard. You can reorganize full databases, selected schema accounts, or selected objects.
7. Click **Next**.

8 The next panel allows you to customize the reorganization. You can either set options at the global level (affects all objects) or specific object level.

9 Click **Next**.

10 The next panel builds your reorganization script and allows you to preview the operation. Click Next when ready.

11 Click **Next**.

The final panel allows you to open the reorganization script in a SQL editor, immediately run it, or schedule it for later execution.

Proceed to Proactive Space Management Using Space Analyst for Sybase ASE

Proactive Space Management Using Space Analyst for Sybase ASE

If you have many databases to take care of, then you might want to relieve the burden of performing detailed storage analysis on every database, each of which may or may not require attention. Fortunately, Space Analyst makes it easy to create smart space management jobs that can automate this task for you.

You can create two specialized space management jobs that allow you to proactively eliminate space threats before they affect your database:

You can create a "reorganization analyst job" that allows you to specifically select various databases and/or objects and then set reorganization thresholds that determine when a reorganization will occur. Once constructed, you can schedule these jobs to run as often as you'd like. During execution, these jobs will interrogate your chosen database/objects and determine if a reorganization needs to occur by applying your customized reorganization thresholds. If any objects are found to need reorganization, they are dynamically reorganized using any customized options you specify.

You can build space maintenance jobs that will perform functions such as check the validity of your databases and filegroups, repair any found minor corruption, and eliminate unused space in your databases. You can also set the job up to notify you if any problems are found during job execution.

These "smart" space maintenance jobs allow you to truly "set and forget" the task of monitoring and reorganizing the space of your critical databases.

Space Analyst for IBM DB2 for Linux, Unix, and Windows Tutorial

Embarcadero Space Analyst provides sophisticated diagnostic capabilities to troubleshoot bottlenecks and performance inefficiencies that result from poor space management. Running inside the Embarcadero DBArtisan console, Space Analyst helps you identify and correct storage problems by supplying powerful analysis facilities and corrective utilities and that will satisfy advanced experts, yet are easy enough to use for database novices.

Space Analyst for IBM DB2 for Linux, Unix, and Windows Product Benefits

Embarcadero Space Analyst allows the database administrator to quickly ascertain if space problems are threatening the availability or performance of their key databases. If storage problems are found, Space Analyst provides powerful but easy to use space management wizards and utilities that can accomplish simple or complex reorganizations of all or parts of a database.

About this Space Analyst for IBM DB2 for Linux, Unix, and Windows Tutorial

This tutorial helps you get started with Embarcadero Space Analyst.

After completion of this tutorial, you will have the foundation you need to explore the many features and benefits of Space Analyst. You will have learned how to get a summary view of storage for your database as well as how to identify potential space-related threats. Finally, you'll use the reorganization wizard to correct any space problems that you find.

This tutorial is divided into 5 sections.

Session 1: Getting Started with Space Analyst

Session 2: Obtaining a Global View of Storage

Session 3: Identifying Storage Structure Problems

Session 4: Pinpointing SQL Server Performance Inefficiencies

Session 5: Correcting Space-Related Problems

You can use this basic tutorial as a roadmap of product highlights.

After starting this tutorial, you can select **Help** from the menu to find many additional resources that complement and build on many of the activities shown in this brief tutorial.

For more information, see Additional Evaluation Resources.

Session 1: Getting Started with Space Analyst for IBM DB2 for Linux, Unix, and Windows

Download and Install

You can obtain the latest version of Space Analyst by downloading DBArtisan from the Embarcadero Web site at http://www.embarcadero.com/downloads/downloaddbartisan.asp. Space Analyst runs fully contained within DBArtisan so you will need the DBArtisan product to use Space Analyst.

Click **Download**, and follow the steps indicated. Save the file on your computer and then double-click to launch the self-extracting file that will guide you through the installation process.

When you first install an evaluation copy of DBArtisan, you can use the tool for 14 days. After that time, a permanent license is needed.

Proceed to Space Analyst for IBM DB2 for Linux, Unix, and Windows Overview

Space Analyst for IBM DB2 for Linux, Unix, and Windows Overview

Once you are connected to an IBM DB2 for Linux, Unix, and Windows database, you can open Space Analyst from within DBArtisan two ways:

1 On the **Analyst** menu, select **Space Analyst**.

NOTE: You do not need to install anything for Space Analyst to work.

Proceed to Session 2: Obtaining a Global View of Storage Using Space Analyst for IBM DB2 for Linux, Unix, and Windows

Session 2: Obtaining a Global View of Storage Using Space Analyst for IBM DB2 for Linux, Unix, and Windows

When it opens, Space Analyst shows a complete overview of the storage situation of a particular database or database server. The first tab of Space Analyst displays all pertinent global statistics regarding storage structure usage throughout a database.

There are four distinct sections displayed on the first **Storage Overview Tab** of Space Analyst:

- Storage Overview
- Tablespace Space Overview
- Log Overview
- Tablespace Objects Overview

Space Analyst for IBM DB2 for Linux, Unix, and Windows - Storage Overview

The breakdown of space usage of the allocated tablespace space, by index, table and temp data is shown in this section. You can also see the total number of tablespaces and the number SMS and DMS spaces here.

NOTE: Make sure your database object statistics are up to date. Space analyst uses these statistics in the various graphs and grids.

Proceed to Space Analyst for IBM DB2 for Linux, Unix, and Windows -Tablespace Space Overview

Space Analyst for IBM DB2 for Linux, Unix, and Windows -Tablespace Space Overview

The Tablespace Space Overview section displays summary information regarding every tablespace on the database. When invoked in graph form, this section shows how much total space is allocated to each tablespace. When shown in grid form, this section shows not only the total allocated for each tablespace, but also includes information such as used space, free space, state, number of containers etc.

TIP: To see a grid-style display of the information shown in the graph, position your mouse over the graph, click the right mouse button and select the ‘toggle display’ option. To return to the graph, just perform the same actions over the grid.

When this section is in grid form, you can also perform a number of actions on a tablespace like open it in a DBArtisan editor, extract its DDL, update its object statistics and more. For example, if you want to update the statistics on all tables in a tablespace, you would perform the following:

1 Select a tablespace in a grid.

2 Right-click the tablespace, and select **Update Statistics**.

OR

3 On the **Space Analyst** toolbar click the **Update Statistics**.

Proceed to Space Analyst for IBM DB2 for Linux, Unix, and Windows - Log Overview

Space Analyst for IBM DB2 for Linux, Unix, and Windows - Log Overview

The log overview provides a quick glance into database transaction log usage. This section shows the percentage of used and free log space for the database, the various database configuration parameters associated with the log space and the high watermarks for the log space usage.

Proceed to Space Analyst for IBM DB2 for Linux, Unix, and Windows -Tablespace Objects Overview

Space Analyst for IBM DB2 for Linux, Unix, and Windows -Tablespace Objects Overview

The tablespace objects overview section provides an overview of how much space is used by the tables and indexes on each tablespace.

Proceed to Session 3: Identifying Storage Structure Problems

Session 3: Identifying Storage Structure Problems

Space Analyst can help you quickly find storage structure problems that threaten the availability of your databases.

Proceed to Pinpointing 'Out of Space' Problems

Pinpointing 'Out of Space' Problems

You can quickly tell if your transaction logs are about to run out of free space by following these steps:

1 Open **Space Analyst**.

2 View the **Database** and **Log Overview** sections on the **Storage Overview Tab**.

3 Examine the free and used space metrics in the pie chart. If the transaction log is nearing the end of free space, then action should be taken to prevent a log full error.

You can also obtain graphical views of tablespaces and pertinent storage property information by viewing the Tablespaces tab of Space Analyst. You can visually get a handle on the amount of used and free space of your tablespaces and quickly view critical storage property facts such as the number of containers for each tablespace, the space allocated for each container, the usage of tablespace space on each partition in a multi-partition database.

All the graphs can be toggled to grids by using the right mouse menu option of 'Toggle Display'. When in grid mode, very detailed information is available concerning the selected tablespace.

NOTE: You can perform further corrective actions on a tablespace, by opening it in the tablespace editor by using either the right mouse menu or the Command toolbar button (the yellow arrow).

Proceed to Using Analysis Reports

Using Analysis Reports

You can run HTML-formatted reports that pinpoint all major storage-related problems for a particular database and the objects that are contained within it. To produce an Analysis Report, follow these steps:

1. On the Space Analyst toolbar, select **Analysis Report**.
2. If you want to change the location and report title of the report, edit the predefined entries for the **Report File Name** and **Report Title**.
3. Select the types of space problems you would like identified on the report by using the **Category** list.
4. Click **Execute**.

   Space Analyst creates a frame-based Analysis report for the database you indicated.

Proceed to Session 4: Pinpointing SQL Server Performance Inefficiencies

Session 4: Pinpointing SQL Server Performance Inefficiencies

Space Analyst can greatly accelerate the process of locating objects in your database that are degrading overall response times because of their current inefficient storage usage.

Proceed to Finding Space-Related Table Problems

Finding Space-Related Table Problems

You can quickly tell if any tables in a tablespace that may be contributing to performance slowdowns by following these steps:

1. Open the **Space Analyst** and select the **Tables Tab**.
2. Select the tablespace you would like to interrogate in the **Tablespace** list.
3. Click the **Refresh** button.

You with see a graphical count of the number of tables that are experiencing specific types of storage inefficiencies in the upper area of the tab. The table reorganization candidate details grid in the lower area of the tab can be used to view detailed storage information for every table in the selected tablespace. Space Analyst uses predefined thresholds (which you can customize by selecting the Threshold Options toolbar button) to determine how many and what tables are suffering from storage inefficiencies.

The types of table-related space problems shown in the bar graph found on the objects tab include:

- **Overflow Problems** -Overflows occur in tables when overflows occur when rows are updated with more bytes than the original row (VARCHAR fields etc.). Overflows also occurs when the new columns are added to existing tables. Since the overflow data will not be stored contiguously with the other data, it causes access inefficiencies.
- **Space Utilization Problems** - Normally, DB2 reclaims any space freed up by deletes, in subsequent insert/update operations. However, if a DB2 table has the 'APPEND' mode set to on and a large amount of data is deleted from the table DB2 will not reclaim freed space. Also when some types of large load operations are performed, the table may be left with a large amount of unused space.
- **Page Utilization Problems** - As the number of empty pages increases, the need for table reorganization also increases. Empty pages waste space and reduce the table scan efficiency, since empty pages are read into the buffer pool for a table scan wasting bufferpool memory.

- **Index Contention** - Unless the user explicitly separates the index and table tablespaces, DB2 tables and indexes are placed on the same tablespace. This can become a bottleneck with heavily accessed tables using index scans since data and index accesses will compete for the IO device resources.

The graph at the top of the Tables tab shows the overall count of problem tables, but you can also view the details for each type of problem by using the Details grid at the bottom of the tab. For example, if Space Analyst identifies tables that have overflow records, then to find the actual tables and view details regarding the identified space problem, you would perform the following:

1. Select **All** from the **Tablespace** list in the **Details** grid.
2. Select **All** from the **Table Schema** list in the **Details** grid.
3. Select **Tables with Overflow Problems** from the **Display** list in the **Details** grid.
4. Click **Refresh**.

You would then see the actual tables that suffer from overflowed records along with their demographic details. If you wanted to take corrective action on any object, you could highlight one or more tables and click the right mouse button to bring up a menu of possible actions.

Proceed to Finding Space-Related Index Problems

Finding Space-Related Index Problems

You can quickly tell if any indexes in your database may be contributing to performance slowdowns by following the exact same steps as outline above for tables with one exception: You want to select the Indexes tab instead of the Tables tab. The overview graph and index details sections of the Indexes tab work in the exact same manner for indexes as they do for tables.

Proceed to Session 5: Correcting Space-Related Problems

Session 5: Correcting Space-Related Problems

Once you have identified space-related bottlenecks in your database with Space Analyst, you can use a number of different tools and facilities within the product to eliminate each storage headache.

Proceed to Correcting Tablespace 'Out of Space' Problems

Correcting Tablespace 'Out of Space' Problems

Space Analyst can help you quickly remedy tablespaces that are running out of room by:

- Resizing existing tablespace containers to have more room.
- Adding new containers file to an existing tablespace.
- Moving Objects from a near full tablespace to less utilized tablespace

Resizing and adding of tablespace containers can be managed by performing the following steps:

1. Open **Space Analyst** and select the **Database Overview Tab**.
2. Select the graph in **Tablespace Overview** section and change to the grid display by selecting the right mouse menu option toggle display.
3. Select the DMS tablespace that is running low on free space and either select the **Open** toolbar option or the **Open** right mouse menu option.

This opens the Tablespace in DBArtisan's tablespace editor. Once you are in the editor, you can remedy the tablespace problem in one of two ways:

- To resize an existing container, select the container from the grid and click the Edit button and extend/ resize the container to a new, larger size. Then close the editor and select the Alter toolbar button to make the change permanent to the tablespace.
- To add a new container to the tablespace, select the Add button in the editor, and specify the new container attributes. Once complete, close the Add Container dialog and select the Alter toolbar button to make the change permanent to the database.

Proceed to Correcting Space-Related Object Problems

Correcting Space-Related Object Problems

Once you identify specific objects that need storage corrections, you can use Space Analyst to help you quickly resolve your found issues. The Space Management wizard inside of Space Analyst can be used to build reorganization and space maintenance jobs that run right the first time. All jobs can either be immediately run or scheduled for later execution.

To create a standard database or object reorganization job, follow these steps:

1. On the toolbar, click the **Space Management**.
2. Select **New Space Management Task**.
3. Click **Next**.
4. Select **Create standard reorganization job**.
5. Click **Next**.
6. Select the objects you would like to reorganize from the Explorer tree on the left hand side of the wizard. Note that you can reorganize full tablespaces, selected schemas or selected tables.
7. Click **Next**.
8. The next few panels let you customize the reorganization. You can choose to either set options at the global level (affects all objects) or specific object level. Using the options, you can reorganize tables and indexes, perform LOB and Long data reorganizations and specify other customizations.
9. The second to last panel builds your reorganization script and allows you to preview the operation.
10. Click **Next**.
11. The final panel allows you to execute the reorganization script immediately, or schedule it for later execution.

Proceed to Proactive Space Management

Proactive Space Management

If you have many tablespace/schemas to take care of, then you might want to relieve the burden of performing detailed storage analysis on every tablespace or schema, each of which may or may not require attention. Fortunately, Space Analyst makes it easy to create smart space management jobs that can automate this task for you.

You can create a "reorganization analyst job" that allows you to specifically select various tablespaces/ schemas or objects and then set reorganization thresholds that determine when reorganization will occur. Once constructed, you can schedule these jobs to run as often as you'd like. During execution, these jobs will interrogate your chosen database/objects and determine if a reorganization needs to occur by applying your customized reorganization thresholds. If any objects are found to need reorganization, they are dynamically reorganized using any customized options you specify. Optionally, this job can also automatically update the statistics for the selected objects before and after reorganization and rebind any packages that depend on the reorganized objects.

These "autonomic" space maintenance jobs allow you to truly "set and forget" the task of monitoring and reorganizing the space of your critical databases.

Space Analyst Expert Guide

IBM DB2 fragmentation can be broken down into two basic types:

- **External Fragmentation** - this situation exists when indexes have a logical order, based on their key value, which does not match the actual physical order inside the database file that houses the index. When an ordered scan is performed on an index, I/O performance is enhanced when the physical order of the database pages matches the logical order of the index. This is because the disk head can scan in one direction instead of moving back and forth to obtain the needed information.
- **Internal Fragmentation** - this situation exists when tables and indexes aren't utilizing space as efficiently as they should. The amount of wasted space in the database pages artificially inflates the size of the table or index and causes more I/O to be performed than would otherwise be needed if the object were compacted.

Keep in mind that fragmentation is not a performance factor at all once data reaches the IBM DB2 memory caches.

Related Topics

Environments That Benefit From Fragmentation Elimination

Diagnosing Fragmentation

"I've Reorganized, But It Didn't Help!"

Forwarded Rows

Environments That Benefit From Fragmentation Elimination

One important point must be understood regarding fragmentation: not every situation benefits from fragmentation removal. Before you invest a lot of time and effort into diagnosing and attempting the removal of fragmentation, you first need to see if jumping through hoops will actually improve your database's availability and performance.

At the highest level, the environments that benefit the most from fragmentation removal are read-intensive databases where large indexes are being frequently scanned. There are a couple of reasons why this is the case.

In most OLTP environments, data retrieval tends to be mostly selective, which negates most of the bad effects of external fragmentation. OLTP environments also often benefit from internal fragmentation because it's smart to leave room on index pages that are being added to and updated frequently. Free index page space helps avoid the dreaded page split where a page is divided in two due to the fact that incoming clustered index data can't logically fit on a needed page, so IBM DB2 makes room by creating two index pages out of one. Administrators oftentimes specify a low FILLFACTOR for their indexes to create internal fragmentation and hopefully avoid page splits.

Data warehousing or databases with many resource-intensive/scanning queries are another matter. These environments will likely benefit from fragmentation removal. One reason is that the indexes tend to be larger, and therefore are not cached by IBM DB2 as easily as small indexes. Remember that fragmentation has no effect on performance once the data is nestled safely in IBM DB2's memory caches.

Another reason why these environments benefit from fragmentation removal is the effect fragmentation has on IBM DB2's read ahead manager. The read ahead manager helps queries that perform large scans by scanning index pages (and data pages) that it believes will be read and placing them into memory before they are actually needed. Naturally, this process can reduce overall scan times because data read in memory is many times faster than when that same data must be retrieved from disk. However, fragmentation affects how well the read ahead manager works. The read ahead manager will dynamically adjust the size of I/O it performs based on the actual physical order of the data and index pages on the server. When little external fragmentation is present, the read ahead manager can work very efficiently because it can read larger blocks at a time, whereas excessive fragmentation causes the read ahead manager to read smaller blocks of data. The end result when fragmentation is present is less overall throughput.

SQL execution patterns are something else to examine. Queries that perform large index scans are the ones that should see performance improvements when fragmentation is removed in a database. In DBArtisan, you can visually set I/O statistics on in the ISQL facility and examine queries to determine the amount of physical and logical read activity. Queries that show physical and read ahead I/O may see their response times drop when fragmentation has been removed.

Diagnosing Fragmentation

If you have an environment that will benefit from reducing fragmentation, how do you go about recognizing it? There are a number of metrics and factors to consider before you can intelligently start creating reorganization plans.

When diagnosing fragmentation, pay particular attention to the following metrics:

- **Extent Fragmentation** - this highlights any gaps (lack of contiguous order) in extents and indicates the presence of external fragmentation. Percentages of 25-50% or more are not a favorable reading.
- **Logical Fragmentation** - this metric represents the percentage of pages that are out of physical order. High values (greater than 25%) for this metric may mean that index scan performance is not what it could be. Note that this statistic is meaningless for tables without clustered indexes (heap tables).
- **Average Page Density** - this metric represents the fullness of the data or index pages, and is an indicator of internal fragmentation. Remember that the more full a data or index page is, the less I/O needed by IBM DB2 when performing scans of tables or indexes. High values are good here, with anything below 50% being suspect. Low values for this metric often indicate the presence of page splitting. Keep in mind, however, that internal fragmentation isn't necessarily bad in OLTP environments where large FILLFACTORs are often specified. Also understand that small objects (those will little or no data) will likely show low readings because they simply don't consume enough space in a database page. You can ignore these objects.
- **Extent Switches** - in a perfect world, this will equal # of object extents - 1. Higher values (values many times higher than an object's extent count) can indicate external fragmentation.
- **Scan Density** - this is computed by dividing the optimum number of extent switches by the actual number of extent switches. This percentage can be misleading, however, if the object spans more than one database file and should not be considered in such an event. Values close to 100% are best.

As you can see, you can't just blindly diagnose fragmentation in IBM DB2, but instead have to evaluate a number of factors.

"I've Reorganized, But It Didn't Help!"

Sometimes, a DBA will reorganize a table or index and no measurable benefits are noticed afterwards. The same or worse fragmentation metrics are observed for the objects in question. What could have gone wrong?

When attempting the reorganization of either tables or indexes, you have to understand that only objects that are of a significant size will show improvement from a reorganization. IBM DB2 for Linux, Unix, and Windows has stated that objects with less than 1,000 pages will oftentimes offer no performance improvements whatsoever from being reorganized.

Other points to keep in mind include:

- Tables without indexes receive zero benefit (naturally) from running either a DBCC DBREINDEX or a DBCC INDEXDEFRAG. The only way to reorganize a table without indexes is to create a clustered index on the table and then immediately drop it.
- Heap tables should not be diagnosed as having fragmentation through the logical scan fragmentation measure. This metric is meaningless for such tables.

- As mentioned above, small objects seldom benefit from reorganizations. For example, the page density for a table with only 10 rows in it will likely not be high, and reorganizing it won't change things one bit. You need to keep such things in mind when determining whether to reorganize an object or not. Space Analyst keeps you informed of not only reorganization metrics but object metrics (extent size, row counts, etc.) as well so you can make smart reorganization choices.
- Lack of contiguous free space can result in all fragmentation not being eliminated. Without large amounts of contiguous free space, reorganization operations may have to reuse other areas of space within the database files. The end result is oftentimes an index being reorganized, but still having an amount of logical scan fragmentation.

Forwarded Rows

IBM DB2 will move rows in a table under certain conditions. One situation might arise when you update a row in a table that has a variable-length column to a larger size that will no longer fit on its original page. Another situation would be if IBM DB2 moves a row when the clustered index column changes. When IBM DB2 creates a forwarding pointer, the pointer remains in place unless one of two things happens. The first is when a row shrinks enough to move back to its original location. The second is when the entire database shrinks. When a database file shrinks, IBM DB2 will reassign the row identifiers, which are used as the row locators, so the shrink process never generates forwarded rows.

Forwarded records can reduce performance at times because additional I/O is involved to first obtain the record pointer to the relocated row, and then the row itself. But when does row forwarding become a problem? For example, just because a table has one hundred forwarded rows, does that mean a performance problem exists? The answer is: it depends. If the table has one hundred rows in it, then yes, a problem exists - 100% of the table suffers from forwarded rows. If the table has three million rows, then the forwarded row situation is likely not causing much fuss in the database.

Backup Analyst

Welcome to the Embarcadero Backup Analyst. Backup Analyst increases the efficiency of performing and managing database backups. Backup Analyst can significantly improve operational efficiency and produce significant savings for your organization. A key capability of Backup Analyst for SQL Server is its ability to compress database backup images significantly. Because it can compress backup images by up to 90 percent or more, it will reduce your storage costs. The Backup Analyst Backup Wizard and Backup Analyst Recovery Wizard make backing up and recovering SQL Server databases straightforward and easy.

NOTE: Backup Analyst does not have to be run from within DBArtisan. It is fully supported via scripting and command line functions so you can use the Backup Operations and Restore Operations anywhere.

The table below describes the major sections of Help for the Backup Analyst.

Section	Description
Introduction	Provides information about installing, opening, and licensing Backup Analyst.
Application Basics	Provides conceptual information about Backup Analyst's product design and navigation.
Using Backup Analyst for Microsoft SQL Server	Shows how to use each component of the application. Using Backup Analyst breaks each process and procedure into simple step-by-step instructions.
Backup Analyst for Microsoft Command Line Reference	Describes the wide range of operations that Backup Analyst supports for backing up and restoring SQL Server databases using the command line interface.

Welcome to Backup Analyst

Establishing a reliable backup and recovery plan for production databases is perhaps the most important responsibility of any database administrator (DBA). Indeed, many DBAs find themselves spending a large share of their time implementing, refining and monitoring backup operations to safeguard the availability of critical database applications.

A number of major trends suggest that managing database backups will become a more pressing issue for professional DBAs over time because corporate databases are growing so quickly. Industry analysts estimate that corporate data is doubling every 2 - 2.5 years. As databases become larger, they become progressively harder to backup and recover as time windows become tighter and the risk of failed backups increases. Every DBA knows that it is far harder to backup a 500 GB database than one that is only 50 GB in size. In addition, DBAs find themselves managing more database applications and, therefore, more database backups. In the quest to compete more effectively, companies are automating more business processes and adopting a wider variety of applications, which means that more production databases get deployed in the enterprise. Finally, organizations have renewed their focus on controlling IT costs and raising staff productivity. The confluence of these factors has thrust DBAs into the uncomfortable position of managing a larger number of critical databases, each of which is growing rapidly in size, with limited incremental resources.

Backup Analyst increases the efficiency of performing and managing database backups. Backup Analyst for SQL Server can significantly improve operational efficiency and produce significant savings for your organization. A key capability of Backup Analyst for SQL Server is its ability to compress database backup images significantly. Because it can compress backup images by up to 90% or even more, it will reduce your storage costs. In addition, Backup Analyst for SQL Server can significantly cut the time required to perform backup and recovery operations and the network bandwidth consumed by them. As a result, DBA teams can achieve substantially greater productivity as well as better operational results in managing backup operations, freeing them to spend time on more strategic tasks.

Reliable - Backup Analyst for SQL Server extends the native functionality for performing backup and recovery in SQL Server to incorporate compression and encryption. Backup Analyst is tightly integrated with SQL Server and leverages native functionality of the database wherever possible in order to maximize product reliability and to minimize its footprint.

Fast and Efficient - Customers will find that they can reduce backup file sizes and total processing time substantially by using Backup Analyst for SQL Server. In fact, they will find that the gains are proportionally better as database backups get bigger. Backup Analyst for SQL Server utilizes highly efficient and reliable compression algorithms in order to achieve the best balance of compression ratio vs. compression processing time.

Scalable - Customers can deploy Backup Analyst for SQL Server to a large number of SQL Server instances very efficiently. The product consumes very little computing resource and works very efficiently with large database servers.

Manageable - Backup Analyst for SQL Server builds on the strength of Embarcadero's DBArtisan family of database administration products. From the DBArtisan Workbench, users will find that it is easy to manage the disaster recovery plans for even a large number of SQL Servers.

Economical - Because of its reliability, efficiency, scalability and ease of management, customers will find that Backup Analyst for SQL Server yields significant and tangible savings in storage costs, bandwidth requirements, processing time and manpower.

For more information, see Backup Analyst Functionality.

Backup Analyst Functionality

The Backup Analyst includes the following features:

- Backup/Recovery Management

>

- Backup Wizard that offers:
  - Increased backup speed for large databases
  - Backup compression
  - Backup encryption
  - Support for full, differential, file/file group, transaction log, table, and indexed view backups
  - Scheduling
- Recovery Wizard that offers:
  - Increased backup speed for large databases
  - Backup compression
  - Backup encryption
  - Support for full, differential, file/file group, transaction log, table, and indexed view backups
  - Scheduling
- Notifications and scheduling via SQL Agent
- Reports
- Alerts
- Basic database/object maintenance functions including:
  - Database Wizard (SQL Server)
  - Modifying objects using editors
  - Extract
  - Update Statistics
  - DBCC

Backup Analyst System Requirements

The basic requirements for installing and using Backup Analyst for SQL Server are quite modest. Essentially, you must be running a supported version of SQL Server and have a small amount of disk space on the SQL Server machine to store the backup agent code.

Requirement	Description
Supported Versions of SQL Server	SQL Server 2005 SQL Server 2000 SQL Server 2000 64-bit (forthcoming) Yukon (forthcoming)
Disk Space	5 MB
Windows Security	Privileges to install software on the SQL Server machine
SQL Server Security	Must login to the SQL Server using a database account that is either system administrator ('sa') or is a member of the 'sysadmin' fixed server role.

For more information, see:

- Backup Analyst Installation

What You Should Know Before Installation

When you install the Analyst, several Embarcadero-generated stored objects are pushed onto the database server:These objects are as follows:

embt_ba_create_alerts	embt_ba_create_embt_backupinfo	embt_ba_create_job_categories
embt_ba_list_backupdevices	embt_ba_list_db_filegroups	sp_BackupRestoreDrop
sp_BackupRestoreInstall	sp_embt_ba_database_backup_history	sp_embt_ba_database_drill_down
sp_embt_ba_database_file_backup_history	sp_embt_ba_database_overview	sp_embt_ba_database_restore_history
sp_embt_ba_databases_needing_backup	sp_embt_ba_file_summary	sp_embt_ba_filegroup_summary
sp_embt_ba_list_database_status	sp_embt_ba_root_db_summary	sp_embt_ba_table_summary
sp_embt_ba_trans_log_summary	sp_embt_ba_version	sp_execBackupRestore

Backup Analyst Installation

Installing Backup Analyst for SQL Server should only take several minutes for each database instance. The installation requires you to follow a simple setup wizard that walks you through the process of installing and registering the backup agents on an instance of SQL Server.

The Backup Analyst for SQL Server requires two components: a console that installs as part of the DBArtisan 8.0 and higher client, and a set of EMBT proprietary extended stored procedures that are installed on the target SQL Server server.

NOTE: The client-side console can only be used against those SQL Servers that are licensed to run the proprietary extended stored procedures. A Backup Analyst license applies to one SQL Server, but it covers all databases running on that server.

CAUTION: Backup Analyst server components must be installed by the SQL Server service account owner who has administrative rights on the server machine.

In order to complete the wizard, you must:

- Know the name of the target SQL Server instance.
- **Have sufficient Windows privileges to install software on the SQL Server machine**. You must have local admin rights on the server machine.

  NOTE: The BackupAnalyst is installed and started on the SQL Server machine. You do not need to install DBArtisan on the server machine, just the Backup Analyst component.

- Obtain a login and password to that SQL Server instance with sufficient privileges to install the agents (see Security Requirements).

Installing Backup Analyst

To install Backup Analyst, do the following:

1. Copy the **BackupAnalyst** folder. This folder is typically located in the main DBArtisan installation directory, typically **C:\Program Files\Embarcadero\DBA**<*version_identifier*>****.

   The BackupAnalyst folder includes the Backup Analyst components necessary for the server-side installation of Backup Analyst.

2. Paste the **BackupAnalyst** folder to any location on the machine that hosts the target, licensed SQL Server instance. You could paste the folder to the desktop for example.

3. Open the **BackupAnalyst** folder and double-click the **setup.exe** icon.

   DBArtisan opens the Embarcadero Backup Analyst InstallShield Wizard that guides you through the installation of Backup Analyst. You

4. Complete the wizard panels, and then click **Install**.

   DBArtisan installs the Backup Analyst engine on the SQL Server machine(s) that you will be testing against.

DBArtisan also performs the following tasks:

- Copies extended stored procedure DLLs and dependencies to target SQL Server directory
- Creates extended stored procedures into any target database
- Creates a standard stored procedures used by the product into any target database
- Creates a backup device if needed
- Creates a new role that controls security
- Grants a new role to all system administrator logins
- Creates a new system job categories (optional)
- Creates new alerts (optional)

NOTE: When you first install an evaluation copy of DBArtisan and the Backup Analyst engine, you can use the tools for 14 days. After that time, a permanent license is needed.

Uninstalling Backup Analyst

To uninstall Backup Analyst, do the following:

1. In **Program Files\Microsoft SQL Server\BackupAnalyst**, double-click the **setup.exe** icon.

   DBArtisan opens the Embarcadero Backup Analyst InstallShield Wizard.

2. Select **Remove**.

3. Click **Next**.

4. Select **Remove**.

   The wizard uninstalls all Backup Analyst dlls, exes, and scripts.

5. Click **Finish**.

Repairing Backup Analyst

The InstallShield Wizard lets you repair installation errors in BackUp Analyst. Fixes missing or corrupt files, shortcuts, and registry entries.

1 In **Program Files\Microsoft SQL Server\BackupAnalyst**, double-click the **setup.exe** icon.

DBArtisan opens the Embarcadero Backup Analyst InstallShield Wizard.

2 Select **Repair**.

3 Click **Next**.

4 Select **Install**.

The wizard reinstalls Backup Analyst.

5 Click **Finish**.

Starting the SQL Agent

To use Backup Analyst, your SQL Server SQL Agent should be running on the machine where you are installing the components. You can start it via the Services window.

1 Select **Start**, **Settings**, **Control Panel**, **Administrative Tools**, and then **Services**.

2 Right-click **SQLSERVERAGENT**, and select **Start**.

Opening Backup Analyst

To open the Backup Analyst, do the following:

1 On the **Analyst** menu, click **Backup Analyst**.

DBArtisan opens the Backup/Recovery Manager.

Licensing Backup Analyst

Backup Analyst includes a 30-day trial period in which you can use the application before you must license it. To use Backup Analyst after 30 days, you must obtain a Workstation license for your machine. To continue using the application without interruption, we recommend that you license it as soon as possible.

You can request Backup Analyst Workstation license from Embarcadero Technologies by e-mail or fax.

To add the new license to Backup Analyst, do the following:

1 Go to **Program Files/Embarcadero/ BackupAnalyst**, and then paste the license.txt into the folder.

TIP: To check the status of your Backup Analyst license, click the Server Information button on the Backup Analyst tool bar.

For more information, see Technical Requirements.

Single-User Mode

Backup Analyst works with single-user mode databases with the following restrictions:

- While in single-user mode, Backup Analyst does not let you Backup or Restore a master database using the UI.

>

- While in single-user mode, use the command line interface to Backup or Restore a master database.

Backup Analyst Application Basics

Backup Analyst includes the following features:

- Backup/Recovery Management
- Backup Wizard that offers:
  - Increased backup speed for large databases
  - Backup compression
  - Backup encryption
  - Support for full, differential, file/file group, transaction log, table, and indexed view backups
  - Scheduling
- Recovery Wizard that offers:
  - Increased backup speed for large databases
  - Backup compression
  - Backup encryption
  - Support for full, differential, file/file group, transaction log, table, and indexed view backups
  - Scheduling
- Notifications and scheduling via SQL Agent
- Reports
- Alerts
- Basic database/object maintenance functions including:
  - Database Wizard (SQL Server)
  - Modifying objects using editors
  - Extract
  - Update Statistics
  - DBCC

Backup Analyst Product Design

The product leverages agent technologies that extend the native backup and restore functionality of SQL Server to include compression and encryption. These agents, which are installed on each supported instance of SQL Server, have a very small footprint, requiring only several megabytes of storage. In order to integrate the extended functionality tightly with SQL Server, Backup Analyst registers these agents as extended stored procedures that can be called from inside the database.

The parameters passed to the extended stored procedures closely follow the patterns established by the native SQL Server backup and restore commands. As a result, users should find the product easy to learn and to integrate with established backup routines. The subsections that follow document the arguments for calling these extended stored procedures.

User Interface

Backup Analyst may be run within DBArtisan, Embarcadero's award-winning, cross-platform database administration tool. Customers using Backup Analyst through DBArtisan will find a rich interface that provides a central console to manage all their backup and recovery needs.

>

The Backup Analyst Backup Wizard and Backup Analyst Recovery Wizard make backing up and recovering SQL Server databases straightforward and easy. The Backup Analyst Alerts Wizard lets you create alerts and notifications so no backup or recovery problem goes overlooked. All backup and recovery jobs can be filtered, managed, and tracked in one place. And informative backup and recovery reports can be created and scheduled to run in an automated fashion so everyone is kept informed of backup operations.

NOTE: Backup Analyst does not have to be run from within DBArtisan. It is fully supported via scripting and command line functions so you can use the Backup Operations and Restore Operations anywhere.

Backup Analyst Toolbar and Menus

Toolbar:

Main Toolbar Menu

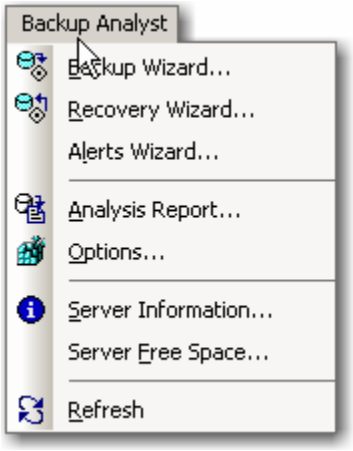

>

Command Menu

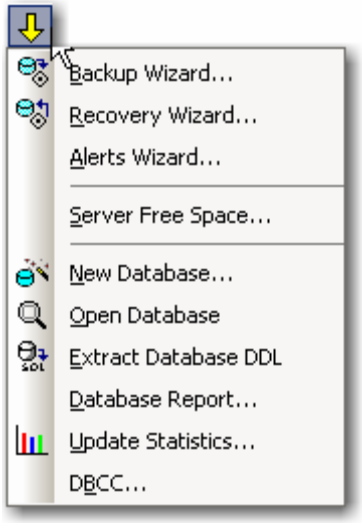

The Backup Analyst Toolbar and menus help you perform the most common operations in Backup Analyst like opening a Backup Wizard and Recovery Wizard, creating reports, and using DBArtisan functionalities including DBCC and Update Statistics.

Backup/Recovery Management

Backup/Recovery management provides instant access to major backup and recovery needs, and key database objects. The Backup/Recovery management facilities also display the current status and events of objects, jobs, and other items.

Backup/Recovery management includes the following components:

Backup Analyst Explorer	Backup Analyst Server Information Panel	Backup Analyst Toolbar

Backup Analyst Explorer

The Backup Analyst Explorer lets you navigate the objects of your SQL Server as they relate to backup and recovery. For example, you can manage databases, files and filegroups, transaction logs, and tables, as well as system objects like backup devices, the SQL Server Agent, jobs, and alerts/notifications.

The Backup Analyst Explorer includes the following nodes:

- Server Node
- Databases Node
- Individual Database Nodes
- File Groups Node (and subnodes)

>

- Backup Recovery Management Node:
  - Alerts Node
  - Backup Devices Node
  - Operators Node
  - Jobs Node

Server Node

The SQL Server name is at the top of the Server Node. The default landing page is displayed when this node is selected.

The upper portion of the landing page displays the current status of various backup/recovery activities including the number of backups and recoveries running, databases needing backup, and job successes and failures. The lower pane displays the summary of the last 'N' backups for all databases on the server (where 'N' is an option defined in the global options dialog box).

Databases Node

On the Explorer tree, green database nodes indicate the database is online. A red node indicates the database is offline. When you select the databases node, the right pane displays the Database Summary of backups. You see when each database was created, its online status, size, the log size, the last backup date and type, compressed and uncompressed sizes, and whether or not backup/restore jobs are actively running.

Individual Database Nodes

When you select a specific database's node in the Explorer, Backup Analyst displays its database properties in the right pane. The top portion displays basic metadata regarding the database in graphic format, while the bottom displays a backup history for the database.

File Groups Node

When you select the Filegroups node in the Explorer, Backup Analyst displays both the File Group Summary and the Filegroup/File Backup Summary in the right pane. The File Group summary identifies each group (for example, Log and Primary groups) and their capacities to grow, their sizes, and other helpful information. The File Backup summary gives you specific information about backups including the backup type, the run time, and so forth. Additional summaries are provided for subordinate nodes as you click them:

> **NOTE:** The list that follows are the default nodes. You can add nodes and subnodes with the Database Wizard (SQL Server).

- Primary,
  - Files,
  - Tables

Transaction Log

The Transaction Log node, when selected, graphically displays transaction log properties in the right pane. The graph enables you to see how much physical space is used and how much is available. You also see a transaction log summary that includes information about when a backup job started and finished, along with its expiration date and other pertinent operation details.

>

Functions available from the Backup Analyst Databases Explorer Nodes include:

Backup Wizard	Recovery Wizard	Database Wizard (SQL Server)
Modifying objects using editors	Extract	Report
Update Statistics	DBCC	

These functions are all available from the Backup Analyst toolbar's command menu.

Backup/Recovery Management Node

When you select the Backup/Recovery Management node in the Explorer, Backup Analyst displays default landing page for the server, which is the same information you see when you select the server node.

Available Functionality

- Start SQL Server Agent - starts the SQL Server Agent
- Stop SQL Server Agent - stops the SQL Server Agent

Select the Explorer Backup/Recovery Management node and use the command menu to effect these functions.

Backup Devices Node

When you select the Backup Devices node in the explorer, Backup Analyst displays the Backup Devices Summary in the right pane.

Available Functionality

Backup Device Wizard (SQL Server)	Modifying objects using editors	Extract
Report	View Contents	

Backup Analyst Server Information Pane

The Backup Analyst Server Information window displays information comparable to what you see when you open the About window for DBArtisan. You see the Backup Analyst version number along with pertinent license information.

You can open this information panel from the Backup Analyst drop-down menu on the main toolbar or by clicking  on the Backup Analyst toolbar.

Using Backup Analyst for Microsoft SQL Server

Using Backup Analyst for Microsoft SQL Server covers how to use each component of the application. The Using chapter breaks each process and procedure into simple step-by-step instructions.

NOTE: Backup Analyst is fully contained within DBArtisan. Backup Analyst can be open while you work with other components of DBArtisan.

Backup Analyst for Microsoft SQL Server Functionality

Backup Analyst lets you create SQL Server backup jobs. Backup Analyst extends standard SQL Server backup functionality by providing specialized features that offer backup file compression that results in storage savings, as well as security encryption that supplies special protection for sensitive database information.

The Backup Analyst supports the following functionality:

- Backup Wizard
- Recovery Wizard
- Alert Wizard
- Options
- Report
- Server Free Space

Backup Analyst Backup Wizard

Implementing a plan to back up critical databases is perhaps the most important job for any database administrator. Without a sound disaster recovery plan, your company has no safeguards against hardware and software failures.

The Backup Analyst's Backup Wizard lets you backup your databases without knowing the underlying syntax. It provides support for backing up databases as well as transaction logs, including commands to stripe the backup across multiple devices.

The Backup Analyst Backup Wizard lets you:

- Create a new backup operation.
- Re-use a previously defined operation.
- Specify the target server, database, and type of backup operation.
- Specify backup device type.
- Specify how to report the progress of the backup operation.

>

To Open the Backup Analyst's Backup Wizard:

1 With the Backup Analyst running:

On the **Backup Analyst Explorer**, right-click the database/node you want to back up and select **Backup Wizard**.

OR

From the Backup Analyst drop-down menu, select **Backup Wizard.**

OR

Select **Backup Wizard** from the Backup Analyst toolbar.

In the following sections, you'll find all the fields you'll encounter as you create or edit a backup operation:

Job Category Wizard

Database Backup Wizard

Job Category Wizard

The Job Category Wizard panels appear ONLY when you are performing a backup operation for a MS SQL datasource. The job categories are relevant to the Backup Analyst and are described below:

Option	Description
Standard Backup	Any backup job which is a full database, differential database, file and file group, transaction log or table backup which does not include the compression or encryption options.
Advanced Backup	Any backup job which is a full database, differential database, file and file group, transaction log or table backup which includes at least one of the compression or encryption options.
Recovery	Any recovery job created by Backup Analyst.
Would you like to create the 16 popular SQL Server Alerts for common Backup and Recovery Errors?	Yes - Opens the Alerts wizard. No - Opens the Database Backup Wizard.

Backup Wizard

The first panel of the Database Backup Wizard welcomes you to the Backup Wizard.

Subsequent panels of the wizard are as follows:

Playback or New Operation	Database Selection	Backup Type
Backup Style	Backup Location	
Backup Options	Backup Operation Summary	Backup Operation Execution Selection

Playback or New Operation

Required Information	Description
New Dump Operation	This initiates the process of creating a new backup operation.
Previous Dump Operation	If you want to replay a backup operation, when you select the radio button, the browse function is enabled so you can find the .DBDump file you want to re-use.

Database Selection

Here you select the database you want to backup from a drop-down list and name the backup operation. You can keep the default name for the operation or replace it with one you type. The job description is optional and you can accept the default or write another.

Backup Type

Required Information	Description
Full Database	Backs up the entire database. The size of the database is listed in parentheses.
Differential Database	Records all modifications made to a database after the last database backup, not just the changes made after the last differential database backup.
File and File Group	Select to specify a file or file group.
Transaction Log	Full Truncate- Backs up the inactive portion of the transaction log. The size of the transaction log is listed in parentheses. With No Log - Removes the inactive portion of the transaction log without making a backup of it. You should only use this command when you run out of room in the database and cannot issue DUMP TRAN WITH TRUNCATE_ONLY. Full No Truncate - Makes it possible to dump a transaction log even if the database is inaccessible. You can use this option when the data portion of your database is damaged and the transaction log resides on a separate device. With Truncate Only - Removes the inactive portion of the transaction log without making a backup of it. You should use this command prior to backing up a database.

Backup Style

Required Information	Description
Backup Analyst Backup	Select to use the Backup Analyst for your backup, assuming the Backup Analyst is installed on the database you are backing up. Also lets you compress the backup and apply an encrypted password. **NOTE**: 128-bit encryption is supported
Standard SQL Server Backup	Select to continue the Database Backup Wizard.

>

Backup Location

Required Information	Description
Backup Device	Depending on what is available to you, you can choose to backup to a disk, floppy, pipe, or tape device.
Disk Directory and Name	Lets you specify a backup location on your machine or click the browse button and select a location.
View Contents	Click to open the Backup Media Contents Dialog Box.

Backup and Media Information

Here, for MS SQL, you specify the backup set name, backup description, backup set password, media name, media description, and media name password.

Backup Options

Option	Description
Expire Date	The date when the backup set expires and can be overwritten. If supplied as a variable (@takeover), this date is specified as either a string constant (@takeover = date), as a variable of character string data type (except for the text or text data types), a small datetime, or datetime variable, and must follow the configured system datetime format.
Retain Days	The number of days that must elapse before this backup media set can be overwritten. If supplied as a variable (@days_var), it must be specified as an integer.
Skip backup set expiration and name checking before overwriting	Microsoft SQL Server will skip the set expiration and name checking before overwriting.
Overwrite all backup sets on the device	Microsoft SQL Server will overwrite the contents of the device.
Format the entire media before backup operation	Will format the entire media before the backup operation starts.
Read and verify the integrity after the backup after finish	After the operation completes, the integrity of the backup is verified.
Backup Threads	0 is the default. Zero enables the server default value.

Backup Operation Summary

Here you indicate whether or not you want to save the operation, and if so, where the file should be saved. Also displayed are the elements of the function you are creating. You can go back to earlier panels if you want to make changes before the job runs.

Backup Operation Execution Selection

Option	Description
Open the backup script in an ISQL editor	The backup script will be opened into a new ISQL session.
Perform the backup now	The backup will be executed immediately.
Schedule the backup for later execution	The backup will be scheduled to run at a specified date/time.

Results Summary

If you selected to execute the job immediately, you see a summary of the results of the operation.

For more information, see:

Executing a Backup Script in the ISQL Window

Executing a Backup Script in the ISQL Window

To execute the actual command DBArtisan uses to backup your database in the ISQL Window, do the following:

1. On Backup Operation Execution Selection, click **Open**.

   DBArtisan opens the ISQL Window, which is used to execute SQL and Transact-SQL commands. The ISQL Window displays the procedure call that Backup Analyst makes to back up SQL Server databases, complete with the parameters needed to perform the backup.

2. To execute your backup, on the **ISQL Editor** toolbar, click **Execute**.

   DBArtisan completes your backup and displays a message "Procedure Succeeded" on the Results tab of the ISQL Editor.

Backup Analyst Recovery Wizard

To both validate your back ups and practice the disaster recovery process, you should regularly restore your database dumps. If possible, every day, you should restore your databases on a separate server to enhance system availability. In this manner, you can maintain a warm backup of your database guarding against downtime from system failure.

The Backup Analyst Recovery Wizard lets you:

- Create a new load operation.
- Reuse a previously defined operation.
- Specify the target server, database, and type of restore operation.
- Specify the recovery completion state.
- Specify how to report the progress of the load operation.
- Specify which file to load from the device.
- Specify whether to perform safety checks before replacing the database.
- Save the definition of the dump operation for future re-use.

Completing the Backup Analyst Recovery Wizard

1. On the **Backup/Recovery Manager**, right-click the target node.
2. Select **Recovery Wizard**.

   DBArtisan opens the Recovery Wizard.
3. Complete the wizard panels, and then click **OK**.

   DBArtisan recovers the database.

For more information, see Database Recovery Wizard.

Backup Analyst Alerts Wizard

The Alert Wizard lets you create and configure popular SQL Server Alerts for backup and recovery operations. Alerts inform the user about actions that need to be taken or about events that are occurring and that do not require any action. For example, there is an alert that indicates a backup is in progress. Once the backup is complete, the alert is cleared.

To create backup alerts:

1. Connect to a SQL Server database.
2. From the **Utilities** menu, select **Database Backup**.

   The Job Categories wizard opens and on the second panel of this wizard, if you elect to create Alerts, the Alert Wizard opens.

   OR

1. Connect to a SQL Server database and open the **Backup Analyst**.
2. From the Backup Analyst drop-down menu on the toolbar, select **Alerts Wizard**.

   OR

   From the Command menu on the toolbar, select **Alerts Wizard**.

   NOTE: If an alert already exists and is mapped to one of the existing SQL Server Error numbers, the alert will not be created.

   NOTE: Be aware that when you work with the Alerts, Operators, or Jobs features and click Help, you connect directly to SQL Server Books online. Embarcadero Technologies is not responsible for the contents of the SQL Server Books materials.

SQL Server Agent reads the application log and compares events to alerts that you defined. When SQL Server Agent finds a match, it fires an alert. DBArtisan lets you define alerts to respond to SQL Server events by notifying operators, forwarding the event to another server, or executing a job. These administrative tasks runs automatically if the SQL Server Agent service has been started.

When you select the Alerts node on the Database Explorer, DBArtisan displays all alerts for the server in the right pane of the Explorer window. The column headings in the right pane of the Explorer window provide important information about alerts. These headings are:

- Name
- Enabled
- Error
- Severity

- Last Occurred
- E-mail
- Panel
- Net Send
- Count

You can use the shortcut menu to access functionality associated with alerts. Shortcut menus let you:

- View Properties
- Create a New Alert
- Drop an Alert
- Find Object
- Select All
- Unselect All

NOTE: On the shortcut menu, when you click New, DBArtisan opens the Microsoft SQL Server Agent Alert Properties dialog box where you can enter the new alert information.

Backup Analyst Alert Wizard - Panel 1

The first panel of the wizard displays a message regarding the purpose of the wizard.

Backup Analyst Alert Wizard - Panel 2

This panel lets you select the alerts you would like to create for backup jobs.

NOTE: If there are alerts already defined using any of the SQL Server Errors, the wizard will not create the alert.

Backup Analyst Alert Wizard - Panel 3

This panel lets you select the alerts you would like to create for recovery jobs.

NOTE: If there are alerts already defined using any of the SQL Server Errors, the wizard will not create the alert.

Backup Analyst Alert Wizard - Panel 4

This panel confirms the creation of the alerts.

Backup Analyst Options

The Backup Analyst options enable you to specify some global settings for queries pertaining to your backup jobs.

To Open the Backup Analyst Options dialog box:

1 With the Backup Analyst open:

Click the **Options icon** on the Backup Analyst toolbar,

OR

Select **Options** from the Backup Analyst drop-down menu.

>

The table below describes the options and functionality for the Backup Analyst Options dialog box:

Option	Description
Grace period (days) for gap in backups	Lets you specify the query limit for when a database needs a backup (for example, database 'A' hasn't been backed up in 7 days).
Show last... backups	Lets you specify how much backup history to show (for example, show the last 100 backups) for all backup display grids.
Days to check for successful/failed jobs	Lets you specify the query limit as to how many days to look back to report on succeeded and failed jobs in the main landing page.

Backup Analyst Report

The Backup Analyst Report Wizard lets you:

- Create a new report or re-execute a previously saved report.
- Select the report to create and select the time frame for report data.
- Select the object(s) to be used for the report.
- Select options for the report.
- Save the report definition.
- Preview and then execute or schedule the report.

Important Notes

None

Completing the Backup Analyst for Microsoft Report Wizard

To complete the Backup Analyst Report Wizard, do the following:

1. On the **Backup/Recovery Manager**, right-click the target node.
2. Select **Report**.

   DBArtisan opens the Report Wizard.
3. Complete the wizard panels, and then click **OK**.

   DBArtisan creates the report(s).

Backup Analyst Report Wizard - Welcome

The first panel of the Backup Analyst Report Wizard welcomes you to the wizard.

For more information, see Backup Analyst Report Wizard.

Backup Analyst Report Wizard - Panel 2

The table below describes the options and functionality on this panel of the Backup Analyst Report Wizard.

Option	Description
New Backup Analysis Report	Select to create a new analysis report.

>

Option	Description
Previous Backup Analyst Report	Select to 'playback' a previously saved report definition. The saved report definition file suffix can end with '.ard'

For more information, see Backup Analyst Report Wizard.

Backup Analyst Report Wizard - Time

The table below describes the options and functionality on this panel of the Backup Analyst Report Wizard.

Option	Description
Report	Lets you select type of report: Database Backup History - Provides information regarding all backups over the specified time period. Database Restore History - Provides information regarding all restores over the specified time period. Database File Backup Details - Provides information regarding the files backed up for a database over the specified time period. Databases in Need of Backup - Provides any database that has not been backed up over the specified time period, along with a number of days since the last backup occurred. Objects Modified in Transaction Log - Displays the objects and page counts that have been modified since the last transaction log backup.
Time period sample	Lets you enter a sampling time period.
Databases	Lets you select databases for the report.

For more information, see Backup Analyst Report Wizard.

Backup Analyst Report Wizard - Save

The table below describes the options and functionality on this panel of the Backup Analyst Report Wizard.

Option	Description
Report File Name	Lets you type a report file name. Or click the browse button and locate an existing file.
Report Title	Lets you type a report title.
Save Operation	Lets you type a definition file name. Or click the browse button and locate an existing file.
Execute Now	Generates report immediately.
Schedule	The report will be scheduled to run at a specified date/time.

For more information, see Backup Analyst Report Wizard.

Server Free Space Monitor

The Server Free Space Monitor displays the total amount of physical free space available on the hard drives of a server. This information helps you ensure there is enough physical space for any new backup devices you may need for local storage of hot backups.

To Open the Free Space Monitor:

1. With the Backup Analyst open, select a datasource from the Explorer.
2. Choose **Server Free Space** from the **Command** drop-down menu.

   OR

   Right-click in the right pane and choose **Server Free Space**.

Backup Analyst Command Line Reference

If you do not want to use the Backup Analyst GUI, or if you want to wrap up your database backup and recovery into larger scripts, all of Backup Analyst's functionality is also available via an ISQL window within DBArtisan or from the command line using ETSBackupConsole.exe.

For specific command line operations and examples, see:

Backup Operations

Restore Operations

ETSBackupConsole.exe

For descriptions of all of the parameters for Backup and Restore operations, see:

Description of Backup Parameters

Description of Restore Parameters

Backup Operations

Backup Analyst supports a wide array of options for implementing a disaster recovery plan on SQL Server 2000 and 2005. Specifically, the product allows you to perform a:

- Full database backup
- Incremental database backup
- Transaction log backup
- Backup of specific filegroups or files
- Dynamically naming a backup job

The following sections detail the syntax for executing each of these operations:

- Backing Up an Entire Database
- Backing Up a Database Incrementally
- Backing Up a Database Transaction Log
- Backing Up Specific File Groups
- Backing Up Specific Files
- Dynamic Naming of a Backup Job
- Description of Backup Parameters

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Backing Up an Entire Database

The syntax for backing up an entire database is as follows:

```
EXEC xp_embt_backup
        @Action = {'database'},
        @Database = {string},
```

@BackupFile = {string},

@Compress = {Boolean},

@DBLoginId = {string},

@DBLoginPassword = {string},

@DBPassword = {string},

@DBInstanceName = {string},

@BackupThreads = {long},

@ExpirationDate = {string | date},

@RetainDays = {long},

@Initialize = {boolean},

@MediaDescription = {string},

@MediaName = {string},

@MediaPassword = {string}, - SQL Server 2000 only

@BackupSetName = {string},

@BackupSetDescription = {string},

@BackupSetPassword = {string}, - SQL Server 2000 only

@Restart = {boolean},

@PercentCompleteNotification = {long},

@LogLevel = {long}

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Example

Backing up the entire pubs database:

```
EXEC xp_embt_backup
        @Action = 'database',
        @Database = 'pubs',
        @BackupFile = 'c:\Program Files\Microsoft SQL Server\MSSQL\BACKUP\pubs.dat'
```

For more information, see Backup Operations.

Backing Up a Database Incrementally

The syntax for backing up a database incrementally is as follows:

```
EXEC xp_embt_backup
        @Action = {'incremental'},
        @Database = {string},
        @BackupFile = {string},
```

>

@Compress = {Boolean},

@DBLoginId = {string},

@DBLoginPassword = {string},

@DBPassword = {string},

@DBInstanceName = {string},

@BackupThreads = {long},

@ExpirationDate = {string | date},

@RetainDays = {long},

@Initialize = {boolean},

@MediaDescription = {string},

@MediaName = {string},

@MediaPassword = {string}, - SQL Server 2000 only

@BackupSetName = {string},

@BackupSetDescription = {string},

@BackupSetPassword = {string}, - SQL Server 2000 only

@Restart = {boolean},

@PercentCompleteNotification = {long},

@LogLevel = {long}

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Example

Performing a differential backup of the pubs database:

```
EXEC xp_embt_backup
    @Action = 'incremental',
    @Database = 'pubs',
    @BackupFile = 'c:\Program Files\Microsoft SQL Server\MSSQL\BACKUP\pubs.dat'
```

For more information, see Backup Operations.

Backing Up a Database Transaction Log

The syntax for backing up a transaction log is as follows:

```
EXEC xp_embt_backup
    @Action = {'log'},
    @Database = {string},
    @BackupFile = {string},
    @Compress = {Boolean},
```

>

```
@DBLoginId = 'sa',
@DBLoginPassword = 'password',
@DBPassword = {string},
@DBInstanceName = {string},
@BackupThreads = {long},
@ExpirationDate = {string | date},
@RetainDays = {long},
@Initialize = {boolean},
@MediaDescription = {string},
@MediaName = {string},
@MediaPassword = {string},   -  SQL Server 2000 only
@BackupSetName = {string},
@BackupSetDescription = {string},
@BackupSetPassword = {string},   -  SQL Server 2000 only
@TruncateLog = {'NoTruncate'},
@Restart = {boolean},
@PercentCompleteNotification = {long},
@LogLevel = {long}
```

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Example

Backing up of the transaction log of the pubs database:

```
EXEC xp_embt_backup
    @Action = 'log',
    @Database = 'pubs',
    @BackupFile = 'c:\Program Files\Microsoft SQL Server\MSSQL\BACKUP\pubs.dat',
    @TruncateLog = 'NoOption'
```

For more information, see Backup Operations.

Backing Up Specific File Groups

The syntax for backing up a database filegroup is as follows:

```
EXEC xp_embt_backup
    @Action = {'files'},
    @Database = {string},
    @BackupFile = {string},
```

>

@**DatabaseFileGroups** = {string},

@Compress = {Boolean},

@DBLoginId = 'sa',

@DBLoginPassword = 'password',

@DBPassword = {string},

@DBInstanceName = {string},

@BackupThreads = {long},

@ExpirationDate = {string | date},

@RetainDays = {long},

@Initialize = {boolean},

@MediaDescription = {string},

@MediaName = {string},

@MediaPassword = {string},- SQL Server 2000 only

@BackupSetName = {string},

@BackupSetDescription = {string},

@BackupSetPassword = {string},- SQL Server 2000 only

@Restart = {boolean},

@PercentCompleteNotification = {long},

@LogLevel = {long}

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Example

Backing up the primary filegroup of the pubs database:

```
EXEC xp_embt_backup
        @Action = 'files',
        @Database = 'pubs',
        @BackupFile = 'c:\Program Files\Microsoft SQL Server\MSSQL\BACKUP\pubs.dat',
        @DatabaseFileGroups = 'PRIMARY'
```

For more information, see Backup Operations.

Backing Up Specific Files

The syntax for backing up a database file is as follows:

```
EXEC xp_embt_backup
        @Action = {'files'},
        @Database = {string},
```

>

@**BackupFile** = {string},

@**DatabaseFiles** = {string},

@Compress = {Boolean},

@DBLoginId = 'sa',

@DBLoginPassword = 'password',

@DBPassword = {string},

@DBInstanceName = {string},

@BackupThreads = {long},

@ExpirationDate = {string | date},

@RetainDays = {long},

@Initialize = {boolean},

@MediaDescription = {string},

@MediaName = {string},

@MediaPassword = {string},- SQL Server 2000 only

@BackupSetName = {string},

@BackupSetDescription = {string},

@BackupSetPassword = {string},- SQL Server 2000 only

@Restart = {boolean},

@PercentCompleteNotification = {long},

@LogLevel = {long}

Example

Backing up the primary data file of the pubs database:

```
EXEC xp_embt_backup
        @Action = 'files',
        @Database = 'pubs',
        @BackupFile = 'c:\Program Files\Microsoft SQL Server\MSSQL\BACKUP\pubs.dat',
        @DatabaseFiles = 'c:\Program Files\Microsoft SQL Server\MSSQL\data\pubs.mdf'
```

For more information, see Backup Operations.

Dynamic Naming of a Backup Job

You can easily give a backup job a dynamic name, that is, a time/date stamp by following the script in the example that follows:

```
declare @sql varchar(500)
```

>

```
select @sql = 'master..xp_embt_backup @DATABASE=N"PUBS",

 @ACTION=N"database",

 @BACKUPFILE=N"d:\PUBS_' + rtrim(convert(char(15), getdate(),112)) + '_' + rtrim(replace(convert(char(15),
getdate(), 08),':',")) + '_embt.bak",

 @DBLOGINID=N"sa",

 @DBLOGINPASSWORD=N"password",

 @JOBNAME=N"EMBT_PUBS_080405_083301",

 @OPERATION_STYLE=N"0",

 @BACKUPSETDESCRIPTION=N"Backup Description",

 @BACKUPSETNAME=N"Backup Set Name",

 @JOBDESCRIPTION=N"PUBS Backup Job on 8/4/2005 8:33:01 AM"'


exec (@sql)
```

Description of Backup Parameters

Parameter	Datatype	Required	Description	Microsoft SQL Server Backup Parameter
@Action	String	Always	Valid values: database | files | incremental | log	N/A
@Database	String	Always	The database that is the target of the backup operation.	Database Name
@BackupFile	String	Always	The location and name of the target backup file.	Backup Device/File
@Compress	Boolean	Optional	Whether or not to apply compression to the backup. The default is set to true (1).	
@DBLoginId	String	Optional	The Login ID for accessing the SQL Server and performing the backup. If omitted, then Backup Analyst will use the login of the owner of the extended stored procedure.	N/A
@DBLoginPassword	String	Optional	The password for the login. If omitted, then Backup Analyst will use the password of the owner of the extended stored procedure.	N/A
@DBPassword	String	Optional	Encryption key used to obfuscate password.	N/A
@DBInstanceName	String	Optional	It there are multiple instances of SQL Server installed on the same machine, then you should specify the instance name using this parameter. If none is specified, then 'local' is assumed.	N/A

Parameter	Datatype	Required	Description	Microsoft SQL Server Backup Parameter
@BackupThreads	Long	Optional	Number of processors to dedicate to the backup process.	N/A
@ExpirationDate	Date or String	Optional	Specifies the date when the backup set expires and can be overwritten.	EXPIREDATE
@RetainDays	Long	Optional	Specifies the number of days that must elapse before this backup media set can be overwritten.	RETAINDAYS
@Initialize	Boolean	Optional	Specifies whether or not a backup should overwrite or be appended to any existing backup sets. If overwritten, the media header is preserved.	INIT | NOINIT
@MediaDescription	String	Optional	Specifies information that helps to identify a backup set.	MEDIADESCRIPTION
@MediaName	String	Optional	Specifies the media name, a maximum of 128 characters, for the entire backup media set. If MEDIANAME is specified, it must match the previously specified media name already existing on the backup volume(s). If not specified or if the SKIP option is specified, there is no verification check of the media name.	MEDIANAME
@MediaPassword	String	Optional SQL 2000 Only	Sets the password for the media set.	PASSWORD
@BackupSetDescription	String	Optional	Specifies the free-form text describing the backup set. The string can have a maximum of 255 characters.	DESCRIPTION
@BackupSetName	String	Optional	Specifies the name of the backup set. Names can have a maximum of 128 characters.	NAME
@BackupSetPassword	String	Optional SQL 2000 Only	Sets the password for the backup set. If a password is defined for the backup set, the password must be supplied to perform any restore operation from the backup set.	PASSWORD
@Restart	Boolean	Optional	Specifies that SQL Server restarts an interrupted backup operation. The RESTART option saves time because it restarts the backup operation at the point it was interrupted.	RESTART
@PercentCompleteNotification	Long	Optional	Provides information on the progress of a backup or restore operation.	STATS
@LogLevel	String	Optional	error | info | debug Indicates the level of logging to occur during execution. The default is set to 'info'.	
@DatabaseFileGroups	String	When backing up file groups	Names one or more filegroups to include in the database backup.	FILEGROUPS

>

Parameter	Datatype	Required	Description	Microsoft SQL Server Backup Parameter
@DatabaseFiles	String	When backing up files	Names one or more files to include in the database backup	FILES
@TruncateLog	String	When backing up transaction logs	Valid values: NoLog | NoOption | NoTruncate | Truncate | TruncateOnly NoLog - Records referencing committed transactions are removed. Transaction log is not backed up. NoOption - Transaction log is backed up. Records referencing committed transactions are removed. NoTruncate - Transaction log is backed up. Records referencing committed transactions are not removed, providing a point-in-time image of the log. Truncate - Transaction log is backed up. Records referencing committed transactions are removed. TruncateOnly - Records referencing committed transactions are removed. Transaction log is not backed up.	Various

For more information, see Backup Operations.

Restore Operations

Backup Analyst supports a wide array of options for implementing a disaster recovery plan on Microsoft SQL 2000 and 2005. Specifically, the product allows you to:

- Restore a full database
- Restore a partial database
- Restore a transaction log
- Restore specific filegroups or files
- Verify the integrity of a backup
- Retrieve a summary of the contents of a backup
- Retrieve the header information for a backup
- Retrieve the label for a backup

The following sections detail the syntax for executing each of these operations.

- Restoring a Complete Database
- Restoring a Database to a Different Location
- Restoring a Database Transaction Log
- Restoring Specific File Groups

- Restoring Specific Files
- Retrieving the List of Databases and Log Files in a Backup Set
- Retrieving Backup Header Information
- Retrieving a Backup Label
- Verifying the Integrity of Backup Media
- Description of Restore Parameters

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Restoring a Complete Database

The syntax for restoring an entire database is as follows:

```
EXEC xp_embt_restore
        @Action = {'database'},
        @Database = {string},
        @BackupFile = {string},
        @DBLoginId = {string},
        @DBLoginPassword = {string},
        @FileNumber = {long},
        @EncryptionPassword = {string},
        @DBInstanceName = {string},
        @MediaName = {string},
        @MediaPassword = {string},   - SQL Server 2000 only
        @RelocateFiles = {multi-string},
        @KeepReplication = {boolean},   - SQL Server 2000 only
        @BackupSetName = {string},
        @BackupSetPassword = {string},- SQL Server 2000 only
        @ReplaceDatabase = {string},
        @Restart = {boolean},
        @PercentCompleteNotification = {long},
        @LogLevel = {long},
        @StandbyFiles = {string}
```

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Example

Restoring the entire pubs database:

```
EXEC xp_embt_restore
```

@Action = 'database',

@Database = 'pubs',

@BackupFile = 'c:\Program Files\Microsoft SQL Server\MSSQL\BACKUP\pubs.dat'

For more information, see Restore Operations.

Restoring a Database to a Different Location

The syntax for restoring and moving a database is as follows:

EXEC xp_embt_restore

@Action = {'database'},

@Database = {string},

@BackupFile = {string},

@ReplaceDatabase = {string},

@RelocateFiles = {multi-string},

@DBLoginId = {string},

@DBLoginPassword = {string},

@FileNumber = {long},

@EncryptionPassword = {string},

@DBInstanceName = {string},

@MediaName = {string},

@MediaPassword = {string}, - SQL Server 2000 only

@BackupSetName = {string},

@BackupSetPassword = {string}, - SQL Server 2000 only

@Restart = {boolean},

@PercentCompleteNotification = {long},

@LogLevel = {long},

@StandbyFiles = {string}

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Example

Restoring and moving the pubs database with different file locations (located in drive d, not c):

EXEC xp_embt_restore

@Action = 'database',

@Database = 'pubs',

@BackupFile = 'c:\Program Files\Microsoft SQL Server\MSSQL\BACKUP\pubs.dat',

@ReplaceDatabase = 1,

>

@RelocateFiles = '[pubs],[d:\mssql\data\pubs.mdf],[pubs_log],[d:\mssql\data\pubs_log.ldf]'

For more information, see Restore Operations.

Restoring Part of a Database

The syntax for restoring part of a database is as follows:

```
EXEC xp_embt_restore
        @Action = {'database'},
        @Database = {string},
        @BackupFile = {string},
        @DBLoginId = {string},
        @DBLoginPassword = {string},
        @FileNumber = {long},
        @EncryptionPassword = {string},
        @DBInstanceName = {string},
        @MediaName = {string},
        @MediaPassword = {string}, - SQL Server 2000 only
        @RelocateFiles = {multi-string},
        @BackupSetName = {string},
        @BackupSetPassword = {string}, - SQL Server 2000 only
        @ReplaceDatabase = {string},
        @Restart = {boolean},
        @PercentCompleteNotification = {long},
        @LogLevel = {long},
        @StandbyFiles = {string}
```

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Example

Performing a differential restore of the pubs database:

```
EXEC xp_embt_restore
        @Action = 'database',
        @Database = 'pubs',
        @BackupFile = 'c:\Program Files\Microsoft SQL Server\MSSQL\BACKUP\pubs.dat'
```

For more information, see Restore Operations.

Restoring a Database Transaction Log

The syntax for restoring a database transaction log is as follows:

```
EXEC xp_embt_restore
        @Action = {'log'},
        @Database = {string},
        @BackupFile = {string},
        @DBLoginId = {string},
        @DBLoginPassword = {string},
        @LastRestore = {boolean},
        @ToPointInTime = {string},
        @FileNumber = {long},
        @EncryptionPassword = {string},
        @DBInstanceName = {string},
        @MediaName = {string},
        @MediaPassword = {string} - SQL Server 2000 only
        @RelocateFiles = {multi-string},
        @KeepReplication = {boolean} ,- SQL Server 2000 only
        @BackupSetName = {string},
        @BackupSetPassword = {string}, - SQL Server 2000 only
        @Restart = {boolean},
        @PercentCompleteNotification = {long},
        @LogLevel = {long},
        @StandbyFiles = {string}
```

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Example

Restoring the transaction log of the pubs database:

```
EXEC xp_embt_restore
        @Action = 'log',
        @Database = 'pubs',
        @BackupFile = 'c:\Program Files\Microsoft SQL Server\MSSQL\BACKUP\pubs.dat',
        @LastRestore = false
```

For more information, see Restore Operations.

Restoring Specific File Groups

The syntax for restoring specific database file groups is as follows:

```
EXEC xp_embt_restore
        @Action = {'files'},
        @Database = {string},
        @BackupFile = {string},
        @DatabaseFileGroups = {string},
        @DBLoginId = {string},
        @DBLoginPassword = {string},
        @FileNumber = {long},
        @EncryptionPassword = {string},
        @DBInstanceName = {string},
        @MediaName = {string},
        @MediaPassword = {string}, - SQL Server 2000 only
        @RelocateFiles = {multi-string},
        @BackupSetName = {string},
        @BackupSetPassword = {string}, - SQL Server 2000 only
        @ReplaceDatabase = {string},
        @Restart = {boolean},
        @PercentCompleteNotification = {long},
        @LogLevel = {long},
        @StandbyFiles = {string}
```

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Example

Restoring the primary filegroup of the pubs database:

```
EXEC xp_embt_restore
        @Action = 'files',
        @Database = 'pubs',
        @BackupFile = 'c:\Program Files\Microsoft SQL Server\MSSQL\BACKUP\pubs.dat',
        @DatabaseFileGroups = 'PRIMARY'
```

For more information, see Restore Operations.

Restoring Specific Files

The syntax for restoring specific database files is as follows:

>

```
EXEC xp_embt_restore
        @Action = {'files'},
        @Database = {string},
        @BackupFile = {string},
        @DatabaseFiles = {string},
        @DBLoginId = {string},
        @DBLoginPassword = {string},
        @FileNumber = {long},
        @EncryptionPassword = {string},
        @DBInstanceName = {string},
        @MediaName = {string},
        @MediaPassword = {string}, - SQL Server 2000 only
        @RelocateFiles = {multi-string},
        @BackupSetName = {string},
        @BackupSetPassword = {string}, - SQL Server 2000 only
        @ReplaceDatabase = {string},
        @Restart = {boolean},
        @PercentCompleteNotification = {long},
        @LogLevel = {long},
        @StandbyFiles = {string}
```

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Example

Restoring the primary data file of the pubs database:

```
EXEC xp_embt_restore
        @Action = 'files',
        @Database = 'pubs',
        @BackupFile = 'c:\Program Files\Microsoft SQL Server\MSSQL\BACKUP\pubs.dat',
        @DatabaseFiles = 'c:\Program Files\Microsoft SQL Server\MSSQL\data\pubs.mdf'
```

For more information, see Restore Operations.

Retrieving the List of Databases and Log Files in a Backup Set

The syntax for retrieving the list of databases and log files contained in a backup set is as follows:

```
EXEC xp_embt_restore_filelistonly
        @BackupFile = {string},
```

@EncryptionPassword = {string}

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Example

EXEC xp_embt_restore_filelistonly

@BackupFile = 'c:\Program Files\Microsoft SQL Server\MSSQL\BACKUP\pubs.dat'

For more information, see Restore Operations.

Retrieving Backup Header Information

The syntax for retrieving the backup header information is as follows:

EXEC xp_embt_restore_headeronly

@BackupFile = {string},

@EncryptionPassword = {string}

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Example

EXEC xp_embt_restore_headeronly

@BackupFile = 'c:\Program Files\Microsoft SQL Server\MSSQL\BACKUP\pubs.dat'

For more information, see Restore Operations.

Retrieving a Backup Label

The syntax for retrieving the backup label is as follows:

EXEC xp_embt_restore_labelonly

@BackupFile = {string},

@EncryptionPassword = {string}

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Example

EXEC xp_embt_restore_labelonly

@BackupFile = 'c:\Program Files\Microsoft SQL Server\MSSQL\BACKUP\pubs.dat'

For more information, see Restore Operations.

Verifying the Integrity of Backup Media

The syntax for verifying the integrity of backup media is as follows:

EXEC xp_embt_restore_verifyonly

@BackupFile = {string},

>

@EncryptionPassword = {string}

NOTE: Mandatory parameters to an extended stored procedure are shown in bold text.

Example

EXEC xp_embt_restore_verifyonly

@BackupFile = 'c:\Program Files\Microsoft SQL Server\MSSQL\BACKUP\pubs.dat'

For more information, see Restore Operations.

Description of Restore Parameters

Xp_embt_restore				
Parameter	**Datatype**	**Required**	**Description**	**Microsoft SQL Server Backup Parameter**
@Action	String	Always	Valid values: database | files | log	N/A
@Database	String	Always	The database that is the target of the backup operation.	Database Name
@BackupFile	String	Always	The location and name of the target backup file.	Backup Device/File
@DBLoginId	String	Optional	The Login ID for accessing the SQL Server and performing the backup. If omitted, then Backup Analyst will use the login of the owner of the extended stored procedure.	N/A
@DBLoginPassword	String	Optional	The password for the login. If omitted, then Backup Analyst will use the password of the owner of the extended stored procedure.	N/A
@DBPassword	String	Optional	Encryption key used to obfuscate password.	N/A
@DBInstanceName	String	Optional	It there are multiple instances of SQL Server installed on the same machine, then you should specify the instance name using this parameter. If none is specified, then 'local' is assumed.	N/A
@FileNumber	Long	Optional	Identifies the backup set to be restored. For example, a file_number of 1 indicates the first backup set on the backup medium and a file_number of 2 indicates the second backup set.	FILE
@MediaName	String	Optional	Specifies the media name, a maximum of 128 characters, for the entire backup media set. If MEDIANAME is specified, it must match the previously specified media name already existing on the backup volume(s). If not specified or if the SKIP option is specified, there is no verification check of the media name.	MEDIANAME

>

Xp_embt_restore				
Parameter	**Datatype**	**Required**	**Description**	**Microsoft SQL Server Backup Parameter**
@MediaPassword	String	Optional SQL 2000 Only	Sets the password for the media set.	PASSWORD
@RelocateFiles	Multi-string	Optional	Specifies that the given logical_file_name should be moved to operating_system_file_name. By default, the logical_file_name is restored to its original location. If the RESTORE statement is used to copy a database to the same or different server, the MOVE option may be needed to relocate the database files and to avoid collisions with existing files. Each logical file in the database can be specified in different MOVE statements.	MOVE
@KeepReplication	Boolean	Optional SQL 2000 Only	Instructs the restore operation to preserve replication settings when restoring a published database to a server other than that on which it was created.	KEEP_REPLICATION
@BackupSetName	String	Optional	Specifies the name of the backup set. Names can have a maximum of 128 characters.	NAME
@BackupSetPassword	String	Optional SQL 2000 Only	Sets the password for the backup set. If a password is defined for the backup set, the password must be supplied to perform any restore operation from the backup set.	PASSWORD
@ReplaceDatabase	Boolean	Optional	Specifies that SQL Server should create the specified database and its related files even if another database already exists with the same name. In such a case, the existing database is deleted. When the REPLACE option is not specified, a safety check occurs (which prevents overwriting a different database by accident).	REPLACE
@Restart	Boolean	Optional	Specifies that SQL Server restarts an interrupted backup operation. The RESTART option saves time because it restarts the backup operation at the point it was interrupted.	RESTART
@PercentCompleteNotification	Long	Optional	Provides information on the progress of a backup or restore operation.	STATS
@LogLevel	String	Optional	error | info | debug Indicates the level of logging to occur during execution. The default is set to 'info'.	
@DatabaseFileGroups	String	When backing up file groups	Names one or more filegroups to include in the database backup.	FILEGROUPS
@DatabaseFiles	String	When backing up files	Names one or more files to include in the database backup	FILES

>

Xp_embt_restore				
Parameter	**Datatype**	**Required**	**Description**	**Microsoft SQL Server Backup Parameter**
@StandbyFiles	String	Optional	Specifies the name of an undo file used as part of the backup strategy	STANDBY
@ToPointInTime	String	When restoring transaction logs	Sets an endpoint for restoring a transaction log	N/A
@LastRestore	Boolean	When restoring transaction logs	Specifies whether the transaction log is the last in a chain of log backups	N/A
@LoadHistory	Boolean	Optional	Specifies whether or not to verify the integrity of the backup	N/A

xp_embt_restore_headeronly				
Parameter	**Datatype**	**Required**	**Description**	**SQL Server Backup Parameter**
@BackupFile	String	Yes	Location and name of the target backup file.	Backup Device/File
@EncryptionPassword	String	Optional	Encryption key used to obfuscate password.	N/A

xp_embt_restore_filelistonly				
Parameter	**Datatype**	**Required**	**Description**	**SQL Server Backup Parameter**
@BackupFile	String	Yes	Location and name of the target backup file.	Backup Device/File
@EncryptionPassword	String	Optional	Encryption key used to obfuscate password.	N/A

xp_embt_restore_labelonly				
Parameter	**Datatype**	**Required**	**Description**	**SQL Server Backup Parameter**
@BackupFile	String	Yes	Location and name of the target backup file.	Backup Device/File
@EncryptionPassword	String	Optional	Encryption key used to obfuscate password.	N/A

>

xp_embt_restore_verifyonly				
Parameter	**Datatype**	**Required**	**Description**	**SQL Server Backup Parameter**
@BackupFile	String	Yes	Location and name of the target backup file.	Backup Device/File
@EncryptionPassword	String	Optional	Encryption key used to obfuscate password.	N/A

For more information, see Restore Operations.

ETSBackupConsole.exe

From the command line, the product allows you to perform a:

- Full database backup
- Incremental database backup
- Transaction log backup
- Backup of specific filegroups or files
- Restore a full database
- Restore a transaction log
- Restore specific filegroups or files
- Verify the integrity of a backup
- Retrieve a summary of the contents of a backup
- Retrieve the header information for a backup
- Retrieve the label for a backup

The syntax for backing up or restoring using ETSBackupConsole.exe follows. Mandatory parameters are shown in bold. All other parameters are optional.

NOTE: In addition to {'backup_database'}, action parameters available for **backup operations** are: {'backup_incremental'}, {'backup_log'}, {'backup_files'}, {'restore_database'}, and for **restore operations** are: {'restore_files'}, {'restore_log'}, {'restore_filelistonly'}, {'restore_labelonly'}, {'restore_headeronly'}, and {'restore_verifyonly'}.

@Action = {'backup_database'},

@Database = {string},

@BackupFile = {string},

@DBLoginId = {string},

@DBLoginPassword = {string},

@DBInstanceName = {string},

@BackupThreads = {long},

>

@BackupSetName = {string},

@BackupSetDescription = {string},

@Log Level = {string},

@Compress = {boolean},

@DatabaseFileGroups = {multi-string},

@DatabaseFiles = {multi-string},

@ExpirationDate = {string},

@FormatMedia = {Boolean},

@Initialize = {Boolean},

@MediaDescription = {string},

@MediaName = {string},

@MediaPassword = {string} - SQL Server 2000 only

@BackupSetPassword = {string}, - SQL Server 2000 only

@Restart = {boolean},

@SkipTapeHeader = {boolean},

@TruncateLog = {string},

@UnloadTapeAfter = {string},

@NoRewind = {boolean},SQL Server 2000 only

@PercentCompleteNotification = {long},

@RetainDays = {long}

Example:

Etsbackupconsole.exe

@action=backup_database

@Database=[pubs]

"@backupfile=c:\temp\pubs.embt.bak"

@EncryptionPassword=PASSWORD

Index

B

C

D

E

F

G

H

J

K

L

M

N

O

R

S

T

U

X